

10.0

Monitoring and Performance for IBM MQ



Note

Before using this information and the product it supports, read the information in [“Notices” on page 411.](#)

This edition applies to version 9 release 4 of IBM® MQ and to all subsequent releases and modifications until otherwise indicated in new editions.

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

© **Copyright International Business Machines Corporation 2007, 2026.**

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

Contents

Monitoring and performance.....	5
Monitoring your IBM MQ network.....	5
OpenTelemetry tracing.....	5
Queue manager health check behavior.....	27
Event monitoring.....	30
Message monitoring.....	82
Accounting and statistics messages.....	158
Application activity trace.....	223
System topics for monitoring and activity trace.....	308
Real-time monitoring.....	321
Monitoring clusters.....	333
Monitoring application balancing.....	336
Monitoring performance and resource usage on z/OS.....	338
Tuning your IBM MQ network.....	398
Tuning client and server connection channels.....	398
Tuning distributed publish/subscribe networks.....	399
Reducing the number of unwanted topics in the topic tree.....	408
Aspera gateway can improve performance over high latency networks.....	410
Notices.....	411
Programming interface information.....	412
Trademarks.....	412

IBM MQ Monitoring and performance

Use the monitoring information and guidance in this section, and the specific tuning tips, to help improve the performance of your queue manager network.

About this task

Depending on the size and complexity of your queue manager network, you can obtain a range of information from monitoring the network. You can use that information, along with the information provided in specific tuning tips, to help you tune your network performance.

Monitoring your IBM MQ network

A number of monitoring techniques are available in IBM MQ to obtain statistics and other specific information about how your queue manager network is running. Use the monitoring information and guidance in this section to help improve the performance of your queue manager network.

The following list provides examples of reasons for monitoring your queue manager network:

- Detect problems in your queue manager network.
- Assist in determining the causes of problems in your queue manager network.
- Improve the efficiency of your queue manager network.
- Familiarize yourself with the running of your queue manager network.
- Confirm that your queue manager network is running correctly.
- Generate messages when certain events occur.
- Record message activity.
- Determine the last known location of a message.
- Check various statistics of a queue manager network in real time.
- Generate an audit trail.
- Account for application resource usage.
- Capacity planning.

OpenTelemetry tracing

OpenTelemetry is a collection of APIs, SDKs, and tools, which can be used to instrument, generate, collect, and export telemetry data (metrics, logs, and traces) to help site reliability engineers (SRE) analyze their software's performance and behavior.

- Metrics provide runtime measurements of a system.
- Logs provide recording of events.
- Traces track the flow of an application through a system, enabling you to observe how applications are behaving. The flow normally incorporates several different application components, potentially running in different environments. Traces can provide insights into each individual application's behavior and make it easier to understand where an application might be encountering issues, or performance problems.

 From IBM MQ 9.4.0, queue managers running on AIX®, Linux®, or Windows can take part in traces, and integrate into an OpenTelemetry environment, by installing the IBM MQ OpenTelemetry tracing service.

 From IBM MQ 10.0.0, queue managers on z/OS® can take part in traces, and integrate into an OpenTelemetry environment, using the capabilities of the z/OS platform.

Related information

OpenTelemetry tracing with IBM MQ on AIX, Linux, and Windows

Queue managers running on AIX, Linux, or Windows can take part in traces, and integrate into an OpenTelemetry environment, by installing the IBM MQ OpenTelemetry tracing service.

The IBM MQ OpenTelemetry tracing service is implemented as an IBM MQ API exit. Your IBM MQ entitlement only permits use of the IBM MQ Tracing exit component of the IBM Instana® supporting program. Separate entitlement is required for the component that collects trace data, such as IBM Instana or an OTel Collector.

Issues arising from the use of the exit must be reported to IBM MQ Support unless you are using the tracing exit with an IBM Instana monitoring system and IBM Instana entitlement.

Installing and configuring the IBM MQ OpenTelemetry tracing service

The exit can be downloaded here: <https://ibm.biz/mqinstanaexit>.

For information on how to install and configure IBM MQ tracing, see [Enabling IBM MQ Tracing user exit for on-premises IBM MQ](#) in the IBM Instana Observability documentation.

Details of the supported platforms for the IBM Instana IBM MQ exit can be found here: [Supported platforms for on-premises IBM MQ](#).

Related information

OpenTelemetry tracing with IBM MQ on z/OS

Queue managers running on z/OS can take part in OpenTelemetry traces initiated either on distributed or z/OS and automatically propagate those traces to other z/OS middleware.

High level overview of OpenTelemetry trace

An OpenTelemetry trace provides information about the processing performed on behalf of a specific application request through the components that make up an application. Each component contributes to the trace by emitting a span. A span contains information, span attributes, about the work performed by the component. Spans build up parent-child relationships which allow their position in the trace to be calculated.

A trace is represented by a globally unique number – a trace ID. Each span in the trace refers to the same trace ID.

The first span in a trace is known as the root span. Each span, apart from the root span, has a span immediately preceding it, this is known as the parent of the span. Each span might have zero, one or more spans immediately following it, these are known as child spans. Every span in a trace has a unique number – a span ID. Span IDs are used to build up the parent – child relationships in a trace. Span IDs must be unique inside a given trace.

Information about a trace is stored in an environment specific manner as trace context. At a minimum trace context consists of the trace ID and the most recent span ID. Trace context is propagated between the components that make up an application. As those component emit spans, the trace ID in the trace context stays constant, but the span ID in the trace context is changed to represent the most recent span that has been emitted.

Spans are emitted to an observability backend which stores the spans, and allows them to be visualized as a trace. Example observability backends include IBM Instana as well as open source and third-party tools.

This process is shown in [Figure 1](#):




Figure 1. OpenTelemetry trace overview

High level overview of OpenTelemetry trace in IBM MQ for z/OS

In IBM MQ, trace context is represented by message properties. One of these message properties, **traceparent** contains the trace ID and span ID. Applications, or the environment that the applications run in, can set these properties to enable IBM MQ for z/OS to take part in a trace.

If IBM MQ is enabled for OpenTelemetry, then a queue manager generates a span when a message is put to, or destructively got from, a queue or a topic. The span generated when a message is put is known as a producer span, the span generated when a message is got is known as a consumer span. Additionally, client spans can be generated when the IMS bridge sends a message to IMS via the OTMA protocol.

A queue manager emits spans to the z/OS operating system and the spans are then sent to any observability backend that support the OTLP protocol. The observability backend, may, or may not, run on z/OS.

Emission of spans is controlled by the **OTELTRAC** attribute, on the queue manager, queue objects, and topic objects. For more information, see [“Span emission” on page 10](#).

Trace context can be propagated from an application component into a message when it is put, or from a message into an application component when it is got. This propagation can be either manual or automatic.

- Manual propagation is performed by the application using the IBM MQ message property APIs.
- Automatic propagation is performed by the environment the application component is running in. Automatic propagation is available between IBM MQ for z/OS and CICS® using either the IBM MQ-CICS adapter, or the IBM MQ-CICS bridge. Automatic propagation is available between IBM MQ for z/OS and IMS using either the IBM MQ-IMS adapter, or the IBM MQ-IMS bridge.

Third party instrumentation provides automatic propagation in other environments, for example JMS.

Trace context automatically flows between queue managers using channels. For example trace context will flow between z/OS queue managers, or between z/OS and distributed queue managers.

Trace context propagation is controlled by the **OTELPCTL** attribute, on the queue manager and queue objects. For more information on trace context propagation, see “Trace propagation” on page 13.

Figure 2 shows a queue manager taking part in a trace between two z/OS applications. Spans 2 and 3 are emitted by the queue manager.




Figure 2. A queue manager taking part in a trace

Related information

z/OS Trace context

IBM MQ will take part in an OpenTelemetry trace if there is trace context in the message. Trace context consists of message properties which can be set manually by an application, or by the environment the application is running in.

Two message properties make up trace context: **traceparent** and **tracestate**. These properties follow the W3C specification for trace context (<https://www.w3.org/TR/trace-context/>). You only need to understand the format of these properties if you are setting them, or manipulating them yourself.

If a distributed queue manager using the IBM MQ OpenTelemetry tracing service is sending messages to IBM MQ on z/OS, then Instana specific properties are also set. You do not need to understand the format of these properties.

traceparent

The **traceparent** property is required for IBM MQ to take part in a trace. If this property is set, then IBM MQ optionally emits spans, and performs automatic propagation. IBM MQ manipulates the content of the **traceparent** property as it emits spans.

The format of the **traceparent** property is defined by the W3C specification. An example **traceparent** is:

```
00-0af7651916cd43dd8448eb211c80319c-b7ad6b7169203331-01
```

A **traceparent** consists of four fields separated by a hyphen:

version:

The version of the W3C specification used to form the **traceparent** and **tracestate** properties. Currently the only available version is 00.

trace-id

This is the globally unique value of the current trace in the form of 32 lower case hexadecimal characters. IBM MQ never changes the value of the **trace-id**, it is constant for all spans in a trace. In the example, the **trace-id** is: 0af7651916cd43dd8448eb211c80319c.

parent-id

The ID for a span. This value must be unique for the trace and is in the form of 16 lower case hexadecimal properties. When sending a message to IBM MQ, this is the span ID associated with the application component that sends the message. When receiving messages, this value will have been changed to a new unique value by the queue manager if **OTELTRAC(ON)** is specified on the queue. In the example, the **parent-id** is: b7ad6b7169203331

flags

For a version 00 **traceparent** this can have one of two values: 00 or 01. These two values are used for sampling, as described in the W3C specification. If the flags value is 00, then IBM MQ does not emit spans, even if **OTELTRAC(ON)** has been specified on the relevant queue. If the flags value is 01, then IBM MQ does emit spans, if **OTELTRAC(ON)** has been specified on the relevant queue.

If a message is sent or received from a queue with the **traceparent** property set and that queue has **OTELTRAC(ON)** or **OTELPCTL(AUTO)** set, then the format of the **traceparent** property is validated according to the rules described in the W3C specification. These rules are as follows, where HD is a lower-case hexadecimal digit, 0-9a-f:

- Header must start with 2HD-
- If version is 00, then header must be in the form 2HD-32HD-16HD-2HD
- version cannot be ff
- trace-id cannot be all zeroes
- parent-id cannot be all zeroes
- If version > 0 then after version expect 32HD-
- If version > 0 then after trace-id expect 16HD-
- If version > 0 then after parent-id expect 2HD
- If version > 0 and the above rules apply treat as a version 0 **traceparent**

If validation fails, then both the **traceparent** and **tracestate** properties are removed from the message and no span is emitted. This is tracked using SMF statistics QMSTOCPD, QMSTOCPGD, QQSTOCPD and QQSTOCPGD. For more information, see [“Message manager data records” on page 358](#) and [“Queue data records” on page 377](#).

tracestate

The **tracestate** property is an optional extension to **traceparent** and contains vendor specific data in name value pairs. For example: rojo=00f067aa0ba902b7,congo=t61rcWkgMzE.

IBM MQ does not perform any validation on **tracestate**, and never alters its contents, unless automatic context propagation is enabled. For more information, see [“Trace propagation” on page 13](#).

Instana trace context

If messages are sent to IBM MQ for z/OS with trace context provided by the IBM MQ OpenTelemetry tracing service, then five properties make up trace context: **traceparent**, **tracestate**, and three beginning with **X_INSTANA**. If IBM MQ for z/OS has been configured to emit spans or perform automatic trace propagation, then any **X_INSTANA** properties are deleted from the message.

Span emission

If a message contains trace context then the **OTELTRAC** attribute on the queue manager, or on the queue, or on a topic, controls whether IBM MQ emits a span when a message is sent to, or destructively received from, a queue or published to a topic.

Spans can be emitted both when application, and queue manager processes, send messages to or receive messages from a queue or publish messages to a topic. For example, spans can be generated for:

- CICS, IMS, RRS, JMS, or batch applications.
- Applications connected over a server-connection channel.
- The channel initiator moving messages over a sender-receiver channel, or via a cluster.
- The intra-group queuing agent.
- JMS delivery delay.
- When report messages are generated.
- The IMS bridge receiving a message and sending it to IMS.

The OTELTRAC attribute

The **OTELTRAC** attribute can be configured on the queue manager, or on local, model or remote queues and on topics. By default all local, model or remote queues are configured with **OTELTRAC (QMGR)** which means that changing the **OTELTRAC** on the queue manager can enable, or disable trace for all queues. Topics are configured by default with **OTELTRAC (ASPARANT)**, which delegates its value to the parent topic's configuration. The SYSTEM.BASE.TOPIC uses **OTELTRAC (QMGR)** by default, so if no **OTELTRAC** values have been changed in a topic tree, changing the value on the queue manager enables, or disables trace for all topics.

OTELTRAC on local, model or remote queues can be set to:

- QMGR: the default, use the value of **OTELTRAC** from the queue manager.
- OFF: do not emit spans, even if messages contain trace context.
- ON: emit spans if messages contain trace context.

OTELTRAC on the queue manager can be set to:

- OFF: do not emit spans, even if messages contain trace context.
- ON: emit spans if messages contain trace context.
- NONE: the queue manager never emits spans, even if **OTELTRAC (ON)** is specified on individual queues. Setting **OTELTRAC (ON)** on a queue results in the fast put to waiting getter optimization being disabled.

OTELTRAC on topics can be set to:

- ASPARENT: The default for everything other than SYSTEM.BASE.TOPIC, delegate to the parent topic's configuration.
- QMGR: Use the value of **OTELTRAC** from the queue manager.
- OFF: Do not emit spans, even if messages contain trace context.
- ON: Emit spans if messages contain trace context.

For topics, SYSTEM.BASE.TOPIC uses the default:

- **OTELTRAC (QMGR)**

How span emission works - when sending a message

When a message is sent to a queue which has **OTELTRAC(ON)** configured, then the queue manager checks the message to see if it contains trace context. If it does, and the message was successfully put to the queue then the queue manager generates a producer span containing information about the sending application, and the message that was sent. For a list of attributes that can be in a span, see [“Span attributes” on page 17](#).

Spans for messages sent out of sync point are emitted by the queue manager immediately. Spans for messages that are sent in sync point are emitted when the current unit of work commits, or rolls back.

How span emission works - when receiving a message

A queue manager only emits a span when a message is received, if a span was emitted when the message was sent. A queue manager does not check the value of **OTELTRAC** on a queue when receiving messages from it. The criteria for emitting a span are as follows:

1. A span was emitted when the message was sent to a queue.
2. **OTELTRAC(NONE)** is not set on the queue manager.

If both these criteria are met, then the queue manager generates a consumer span containing information about the receiving application, and the message that was received. For a list of attributes that can be in a span, see [Span attributes](#).

Spans for messages received out of sync point are emitted by the queue manager immediately. Spans for messages that are received in sync point are emitted when the current unit of work commits, or rolls back.

How span emission works – the IMS bridge

When the IMS bridge receives a message from a bridge queue, and sends it to IMS, it will generate a client span if the message contains trace context, and the bridge queue has been configured with **OTELTRAC(ON)**. The client span contains information about the IMS bridge and the IMS region it is interacting with, see [Table 4 in “Span attributes” on page 17](#).

Distributed queuing and clustering considerations

OTELTRAC can be configured on remote queue definitions, and transmission queues. The queue manager applies queue name resolution rules which affect whether **OTELTRAC** comes from a remote queue definition or a transmission queue:

- If a remote queue definition is used for queue name resolution then the value of **OTELTRAC** on that queue applies, regardless of the value of **OTELTRAC** on the transmission queue.
- Otherwise the value of **OTELTRAC** on the transmission queue applies.

OTELTRAC is not a clustered attribute, to emit spans when messages are sent to a remote clustered queue requires that **OTELTRAC(ON)** is specified on the relevant cluster transmission queue.

Publish/subscribe considerations

OTELTRAC on TOPIC objects controls span emission for publish/subscribe. The resolved value of **OTELTRAC** on the topic determines whether spans are emitted for publications on that topic.

Subscriber queue **OTELTRAC** values are not considered when deciding whether spans are emitted for a publication.

If the topic's **OTELTRAC** value resolves to ON, the queue manager emits a publish producer span and for each matching subscription, emits producer child spans as the publication is fanned out to the subscriber destination(s).

Example configuration one

By default, the queue manager does not emit any spans even if messages contain trace context:

- The queue manager defaults to **OTELTRAC(OFF)**.
- Queues default to **OTELTRAC(OFF)** for certain system queues, or **OTELTRAC(QMGR)** for user defined queues or certain system queues.

Example configuration two

Altering the queue manager to **OTELTRAC(ON)** means that any messages sent to the queue manager with trace context result in the queue manager emitting spans.

Example configuration three

If a local queue called APPLICATION.QUEUE.ONE was configured with **OTELTRAC(ON)**, then any message sent to, or received from, that queue with trace context generates a span.

Example configuration four

In Figure 1, APP1 and APP2 are exchanging messages between a pair of z/OS queue managers connected by a sender-receiver channel. Setting **OTELTRAC(ON)** on both the remote and local definition of APP2.REQUEST means that four spans are emitted if APP1 sends a message with trace context:

- A producer span is emitted by MQA when APP1 sends the message to APP2.REQUEST.
- A consumer span is emitted by MQA when the channel initiator takes the message off the transmission queue prior to sending it over the channel.
- A producer span is emitted by MQB when the channel initiator receives the message and puts it to APP2.REQUEST.
- A consumer span is emitted by MQB when APP2 receives the message.




Figure 3. Example configuration 4

Example configuration five

In Figure 2, APP3 and APP4 are exchanging messages between a distributed queue manager and a z/OS queue manager connected by a sender-receiver channel. The distributed queue manager has the IBM MQ OpenTelemetry tracing service (see “[OpenTelemetry tracing with IBM MQ on AIX, Linux, and Windows](#)” on page 6) installed and configured so that spans are emitted when messages are sent to and received from APP4.REQUEST. On the z/OS queue manager, setting **OTELTRAC(ON)** on APP4.REQUEST means that four spans are emitted when APP3 sends a message. The IBM MQ OpenTelemetry tracing service automatically adds trace context when messages are sent to queues that are enabled for OpenTelemetry trace.

- A producer span is emitted by MQC when APP3 sends the message to APP4.REQUEST.

- A consumer span is emitted by MQC when the MCA process takes the message off the transmission queue prior to sending it over the channel.
- A producer span is emitted by MQD when the channel initiator receives the message and puts it to APP4.REQUEST.
- A consumer span is emitted by MQD when APP4 receives the message.




Figure 4. Example configuration 5

When spans are not emitted

IBM MQ prioritizes the integrity of the queue manager over the emission of span data. If a span cannot be emitted for any reason, the span data is discarded. SMF data, either in the [message manager](#), or [per queue](#) is collected to indicate the reasons for discarding span data.

Span data might be discarded for one of the following reasons:

- z/OS has not been configured to collect span data for IBM MQ. Tracked in SMF via QMSTOSFR and QMSTOTSR.
- z/OS has been configured to collect span data for IBM MQ, but has not got capacity to collect span data. Tracked in SMF via QMSTOSFC and QMSTOTSC.
- IBM MQ has reached the limit of how much span data can be stored before it is emitted. Tracked in SMF via QMSTOSFS.
- A transaction has generated more than 1000 spans. Tracked in SMF via QMSTOSPL, QMSTOSCL, QQSTOSPL and QQSTOSCL.
- The queue manager is shutting down with MODE FORCE or RESTART, span data that has not been emitted is discarded.
- The queue manager is shutting down with MODE QUIESCE, after all applications have disconnected the queue manager waits for a maximum of one second for remaining span data to be emitted. Any remaining data is discarded.

The IMS bridge will remove trace context from the data it sends to IMS if that trace context would result in too much data being sent to IMS. In this case the trace context is not propagated to IMS, meaning that IMS cannot emit any spans, and the trace context is therefore not available to the IMS bridge for any reply message, preventing the queue manager from emitting any further spans in that trace. Both of these cases are tracked in SMF via QMSTOBCD.

Related reference

[“Message manager data records” on page 358](#)

Use this topic as a reference for message manager data records.

Trace propagation

IBM MQ propagates trace by sending trace context in message properties. Trace context can manually be propagated by applications interacting with trace context using IBM MQ’s message property APIs. In

certain environments, IBM MQ can automatically propagate trace context to and from the application's environment without the application being involved. Whether trace propagation is manual or automatic is controlled by the **OTELPCTL** attribute.

For more information about trace context, see [“Trace context” on page 8](#).

In order for distributed tracing to be of value, the trace context needs to be propagated between all of the components that make up an application. If trace context is not propagated at any point in the application, then the trace only shows a subset of the components that make up an application. Trace context propagation is a collaboration between the application, and the environment that the application is part of.

Configuring trace propagation

Trace propagation is configured using the **OTELPCTL** attribute on the queue manager and local, model and remote queues or on topics.

OTELPCTL on local, model and remote queues can be set to:

- QMGR: the default, use the value of **OTELPCTL** from the queue manager.
- MANUAL: manual propagation is possible.
- AUTO: automatic propagation is available if the environment supports it.

OTELPCTL on the queue manager can be set to:

- MANUAL: manual propagation is possible
- AUTO: automatic propagation is available if the environment supports it

OTELPCTL on topics can be set to:

- ASPARENT: The default for everything other than SYSTEM.BASE.TOPIC, delegate to the parent topic's configuration.
- QMGR: Use the value of **OTELPCTL** from the queue manager.
- MANUAL: Manual propagation is possible.
- AUTO: Automatic propagation is available if the environment supports it.

For topics, SYSTEM.BASE.TOPIC uses the default:

- **OTELPCTL (QMGR)**

Manual propagation

Applications that cannot take advantage of automatic trace propagation can perform propagation manually by writing suitable code. This is the default.

Applications that send messages can manually create, or propagate, trace context from the application environment into the message by setting the **traceparent** and optional **tracestate** properties in the message using the IBM MQ message property APIs, and then sending the message.

Applications that receive messages can use the IBM MQ message property APIs to extract the trace context properties from the message and copy them into the environment, or into other messages.

Automatic propagation

Applications running in certain environments on z/OS can make use of automatic propagation. This allows trace context propagation to be performed on behalf of the application, without the application needing to be changed, and without the application needing to make use of, or tolerate, message properties.

Automatic propagation is enabled on IBM MQ by setting **OTELPCTL (AUTO)** on the queues used by the application or on topics used by the application. Environment specific configuration, for example in CICS or IMS, might also be required. Setting **OTELPCTL (AUTO)** on a queue results in the fast put to waiting getter optimization being disabled.

Automatic propagation is supported in the following environments:

- IBM CICS Transaction Server for z/OS 6.3 or later (see [CICS and OpenTelemetry](#) in the IBM CICS Transaction Server for z/OS documentation).
- IBM IMS 15.5 or later, with enablement PTFs (see [OpenTelemetry tracing in IMS](#) in the IBM Information Management System documentation) for the IBM MQ IMS bridge.
- IBM IMS 15.5 or later, with PTFs for APAR PH69252 and PH70962 applied for the IBM MQ IMS adapter.

Automatic propagation when messages are sent

If a message is sent to a queue or a topic, configured with **OTELPCTL (AUTO)** from an environment which supports automatic propagation, then the queue manager will take trace context from the environment and set it in the **traceparent** and **tracestate** message properties before sending the message. If the message already contains these properties, then trace context will not be automatically propagated.

If a message is sent to a queue or a topic, configured with **OTELPCTL (AUTO)** from an environment which does not support automatic propagation, then the queue manager does not manipulate the trace context properties. Applications can manually set the properties if required.

Automatic propagation when messages are received

If messages containing trace context are received from a queue configured with **OTELPCTL (AUTO)** and the environment supports automatic propagation, then the queue manager removes the trace context properties from the message and passes them to the environment. If the environment is CICS, then trace context might be associated with the current transaction.

The application does not see the trace context message properties. If the only message properties in the message are the trace context properties, then no message properties are returned to the application. This is useful if the application has not been written to use message properties, or cannot tolerate RFH2 headers.

If a CICS or IMS application, running in a version of CICS or IMS that does not support automatic propagation, or has not been configured for automatic propagation, receives a message from a queue configured with **OTELPCTL (AUTO)**, then any trace context properties are removed from the message and discarded. The application does not see the trace context message properties. If the only message properties in the message are the trace context properties, then no message properties are returned to the application. This is useful if the application has not been written to use message properties, or cannot tolerate RFH2 headers.

If messages containing trace context are received from a queue configured with **OTELPCTL (AUTO)** in other environments that do not support automatic propagation, trace context properties are removed before the message is returned to the application if the application sets **MQGMO_PROPERTIES_AS_Q_DEF** or **MQGMO_NO_PROPERTIES**. If the only message properties in the message are the trace context properties, then no message properties are returned to the application. This prevents the application from unexpectedly receiving an RFH2 header that it might not be able to process.

Automatic propagation when the IMS bridge sends messages to IMS

The IMS bridge automatically propagates trace context to IMS from messages on bridge queues if all the following conditions are met:

- The message contains valid trace context.
- **OTELPCTL (AUTO)** is configured on the bridge queue.
- At the time the IMS bridge connects to IMS, the IMS region supports OpenTelemetry trace and has been enabled to emit and propagate traces (see [OpenTelemetry tracing support](#) in the IBM Information Management System documentation).

Automatic propagation using distributed queuing or queue manager processes

Messages sent or received over channels have the trace context properties flow with the message regardless of the value of the **OTELPCTL** attribute.

Messages that are moved via internal queue manager processes, for example the JMS delivery delay processor or the intra-group queuing agent, automatically have trace context propagated regardless of value of the **OTELPCTL** attribute.

Automatic propagation using OpenTelemetry instrumentation

In some environments, there are zero-code approaches available that automatically start a trace for an application and propagate that context between application components. For example, there is tooling that automatically sets trace context on a JMS message, this tooling can be used with IBM MQ for JMS. It is also possible to use similar tooling with other environments such as Node.js (<https://marketaylor.synology.me/?p=1626>). Trace context set up by these approaches can be used with IBM MQ on z/OS for as long as the trace context message properties are correctly set.

Automatic propagation using OpenTelemetry instrumentation can be used with IBM MQ on z/OS regardless of the value of the **OTELPCTL** attribute.

PROPCTL

Automatic propagation takes place before the processing controlled by the **PROPCTL** attribute on a queue. Therefore, automatic propagation occurs even if **OTELPCTL(NONE)**, or **MQGMO_NO_PROPERTIES** are used.

In some scenarios, messages with trace context might flow from a queue manager which supports OpenTelemetry to one that does not. If the applications that consume messages from the queue manager which does not support OpenTelemetry cannot tolerate message properties, then configure the queues that they consume from with **PROPCTL(NONE)**, which removes all message properties including the trace context properties. If this is done, the OpenTelemetry trace ends at the point where the messages left the queue manager that supported OpenTelemetry.

Related reference

[Change, Copy, and Create Queue DEFINE queues](#)

Configuring z/OS for OTel trace

This topic describes how to configure z/OS to collect spans from IBM MQ for z/OS using the z/OS OpenTelemetry Emitter.

About this task

IBM MQ emits spans in the form of SMF 1158 records. The z/OS OpenTelemetry Emitter needs to be configured to collect SMF 1158 records using an SMF in memory resource.

Create an in-memory resource for use with IBM MQ, using the following instructions. This in-memory resource will be used by any queue manager running on the LPAR.

Procedure

- Determine the name of the in-memory resource.
For example: IFASMF.MQOTEL.
- Ask a security administrator to define a new SAF resource to the FACILITY class to protect the IBM MQ resource.
The resource is the name of the in-memory resource prefixed with IFA. For example: IFA.IFASMF.MQOTEL.
- Give each IBM MQ MSTR started task ID READ access to the SAF resource.
- Update the SMFPRMxx parmlib member to specify the **INMEM** parameter, resource name, and values for the **RESSIZMAX** and **TYPE** sub-parameters for each in-memory resource that you want to define.
For example:


```
SETSMF INMEM(IFASMF.MQOTEL,RESSIZMAX(128M),TYPE(1158))
```

- Issue the SET SMF=xx command to activate the updated SMFPRMxx member.
- Set up and configure the z/OS OpenTelemetry Emitter as described in the z/OS documentation. For more information, see [Setting up z/OS OpenTelemetry Emitter](#).

Spans are from the queue manager's perspective, not the application's perspective

The queue manager generates spans using the information it receives about the API call. This information might not always exactly match the operations that the application has performed. This is particularly likely to be the case when applications make use of IBM MQ classes for JMS, or are connected using a server connection channel. This topic describes the main differences.

- IBM MQ classes for JMS use transactions when applications receive messages, even for non-transacted sessions/JMSContexts. This is to ensure message integrity.
- Applications connected using a server connection channel that has a non-zero value for **MAXSHCNV** in **DIS CHSTATUS** use asynchronous consume, which means that the span has a messaging.operation.type of process, and messaging.ibmmq.consumer.original_wait_interval and messaging.ibmmq.consumer.used_wait_interval are not set even if a get-wait was used.
- Applications issuing receiveNoWait using the IBM MQ classes for JMS in bindings mode see a messaging.ibmmq.consumer.original_wait_interval of 0, as get-wait with a wait interval of zero is used in this scenario.
- Using IBM MQ classes for JMS in migration mode and issuing a receive with a wait-interval results in a maximum messaging.ibmmq.consumer.original_wait_interval of 5000 (five seconds).
- Some internal queue manager tasks use an internal user ID of SYSOPR, which is visible in messaging.ibmmq.application.user_id.
- Applications using IBM MQ classes for JMS and issuing receive with a wait interval do not see messaging.ibmmq.consumer.original_wait_interval and messaging.ibmmq.consumer.used_wait_interval if JMS needs to increase the size of its internal receive buffer.

Span attributes

This topic lists the attributes that can be in spans emitted by IBM MQ for z/OS.

Spans beginning with messaging, that do not contain ibmmq, are defined by the OpenTelemetry semantic conventions for messaging (<https://opentelemetry.io/docs/specs/semconv/attributes-registry/messaging/>).

Common attributes for both consumer and producer spans

Table 1. Common attributes for both consumer and producer spans		
Attribute name	Example value	Description
end_time	"2022-04-29T18:52:58.114304Z"	String When the messaging operation returned to the caller.
mainframe.lpar.name	"MV4C"	String LPAR name.
messaging.client.id	"DF2A711CF2D90001"	String 8 byte CONN value from DIS CONN encoded as 16 upper case characters.

Table 1. Common attributes for both consumer and producer spans (continued)

Attribute name	Example value	Description
messaging.destination.name	"QUEUE1"	String Destination name. If the application performs MQOPEN on an object of type queue this is the ObjectName passed back after MQOPEN/MQPUT1. Therefore, it is the name of the queue before name resolution occurs. If a model queue is opened this will be the name of the generated dynamic queue. The destination name will be a transmission queue if the chinit is taking a message from the queue prior to sending the message over the network. If the application performs MQOPEN on an object of type topic then this will be a topic string. Topic strings can be up to 10240 characters in length so this value might be very long. The value will not be truncated. For producer spans in a publish/subscribe fan out this will be the name of the subscriber's queue.
messaging.destination.temporary	true	Boolean true if the message destination is temp-dynamic and might not exist anymore after messages are processed. Not included otherwise, that is, never set to false.
messaging.ibmmq.application.channel_name	"MY.APP.SVRCONN"	String CHANNEL from DIS CONN . Only included for client or queue manager to queue manager connections.
messaging.ibmmq.application.cics_task_id	"0000034"	String TASKNO from DIS CONN . Only included for CICS applications.
messaging.ibmmq.application.cics_trans_id	"TRAN1"	String TRANSID from DIS CONN . Only included for CICS applications.
messaging.ibmmq.application.completion_code	"MQCC_OK" "MQCC_WARNING"	String IBM MQ completion code in string format.
messaging.ibmmq.application.conname	"9.140.109.78"	String CONNAME from DIS CONN . Only included for client or queue manager to queue manager connections.
messaging.ibmmq.application.description	"IBM MQ Channel Initiator"	String APPLDESC from DIS CONN . Only included if there is a non-blank description.

Table 1. Common attributes for both consumer and producer spans (continued)		
Attribute name	Example value	Description
messaging.ibmmq.application.external_transaction_id	"000000000" "BC8571519B60222D"	String EXTURID from DIS CONN .
messaging.ibmmq.application.ims_psb_name	"CSQQTRMN"	String PSBNAME from DIS CONN . Only included for IMS applications.
messaging.ibmmq.application.ims_pst_id	"0004"	String PSBTID from DIS CONN . Only included for IMS applications.
messaging.ibmmq.application.name	"MQ21CHIN"	String APPLTAG from DIS CONN .
messaging.ibmmq.application.pid	1234	Integer ASID from DIS CONN in decimal.
messaging.ibmmq.application.qmgr_transaction_id	"0000002B4BEFAEFC"	String QMGRURID from DIS CONN .
messaging.ibmmq.application.reason_code	"MQRC_NONE" "MQRC_TRUNCATED_MSG_ACCEPTED"	String IBM MQ reason code in string format.
messaging.ibmmq.application.smf_correlator	"E09D83F4B8DE52527 C63A5A07C658600"	String 32 character string formed from 16 byte value of WTASCORR field. Only included if SMF 116 data is being captured. Allows correlation between accounting data and OpenTelemetry trace.
messaging.ibmmq.application.transaction_type	"QMGR"	String URTYPE from DIS CONN .
messaging.ibmmq.application.type	"CHIN"	String APPLTYPE from DIS CONN .
messaging.ibmmq.application.user_id	"USER1"	String USERID from DIS CONN .
messaging.ibmmq.destination.resolved_name	"A.REMOTE.Q"	String The value of <code>MQOD.ResolvedQName</code> or if type = <code>topic</code> , the topic string from <code>ResObjectString</code> but truncated to 128 characters if necessary.
messaging.ibmmq.destination.resolved_qmgr_name	"REMOTE.QMGR"	String The value of <code>MQOD.ResolvedQMGRName</code> . Might be the local queue manager, or a different one.
messaging.ibmmq.destination.resolved_type	"queue"	String If the application opens an alias queue, the type of object that the alias resolves to. Always "queue".

Table 1. Common attributes for both consumer and producer spans (continued)

Attribute name	Example value	Description
messaging.ibmmq.destination.type	"queue"	String The type of destination the application is interacting with.
messaging.ibmmq.message.accounting_token	"00000000000000000000" 00000000000000000000"	String Value of MQMD.AccountingToken .
messaging.ibmmq.message.appl_identity_data	""	String Value of MQMD.ApplIdentityData .
messaging.ibmmq.message.appl_origin_data	""	String Value of MQMD.ApplOriginData .
messaging.ibmmq.message.expiry	8000 -1	Integer Message expiry converted to milliseconds. -1 if the message will not expire.
messaging.ibmmq.message.feedback	"MQFB_EXPIRATION" "MQRQ_Q_FULL"	String Value of MQMD.Feedback converted to a string. If a MQFB_* or MQRQ_* value then this will be converted to its text equivalent. If the value is an unknown integer, then that integer will be represented as a string.
messaging.ibmmq.message.persistent	true	Boolean true if the message is persistent, false otherwise.
messaging.ibmmq.message.priority	1	Integer Value of MQMD.Priority .
messaging.ibmmq.message.put_appl_name	"USER8"	String Value of MQMD.PutAppName .
messaging.ibmmq.message.put_appl_type	"MQAT_QMGR"	String Value of MQMD.PutAppType converted to a string.
messaging.ibmmq.message.put_date	"20240618"	String Value of MQMD.PutDate .
messaging.ibmmq.message.put_time	"16244201"	String Value of MQMD.PutTime .
messaging.ibmmq.message.reply_destination_name	"Q1"	String Value of MQMD.ReplyToQ .
messaging.ibmmq.message.reply_qmgr_name	"QMGR1"	String Value of MQMD.ReplyToQMGr .
messaging.ibmmq.message.report	"MQRO_NONE" "MQRO_EXCEPTION, MQRQ_COA"	String MQMD.Report converted to a comma separated string. If not a value that the queue manager understands, then will be an integer in string format.

Table 1. Common attributes for both consumer and producer spans (continued)

Attribute name	Example value	Description
messaging.ibmmq.message.type	"MQMT_DATAGRAM"	String MQMD.MsgType converted to a string. If not a value that the queue manager understands, then will be an integer in string format.
messaging.ibmmq.message.user_id	"MLEMING"	String Value of MQMD.UserID .
messaging.ibmmq.transaction.committed	true	Boolean. Whether the transaction was committed (true) or rolledback (false). Only set if the message was put or got in a transaction. This includes scenarios where internal transactions are used, for example with MCA channels, or with JMS.
messaging.ibmmq.transaction.end_time		String UTC timestamp indicating when the transaction completed. Only set if the message was put or got in a transaction.
messaging.ibmmq.transaction.start_time		String UTC timestamp indicating when the transaction started. Only set if the message was put or got in a transaction.
messaging.message.body.size	4096	Integer Message body size in bytes. On MQPUT/MQPUT1: this is the value of the BufferLength passed in by the application. This might include the size of the RFH2 header if the application uses one, for example, with JMS. On MQGET: the value of DataLength returned to the application. If the application requested properties in RFH2 format then those count towards the length
messaging.message.conversation_id	"C3E2D840D4D8F2F14040404040404040DED163210A4A88A1"	String Correlation ID represented as a hex string of 48 upper case characters.

Table 1. Common attributes for both consumer and producer spans (continued)

Attribute name	Example value	Description
messaging.message.envelope.size	9002	Integer Message and meta data size in bytes. Includes message properties and MQMD. On MQPUT: the size of the body, plus the size of either a V1 or V2 MQMD, and if message properties are used the size of the merged properties from OriginalMsgHandle and NewMsgHandle before any trace context property manipulation occurs On MQGET: the size of the body, plus the size of either a V1 or V2 MQMD, and the size of any properties returned via a handle, after any trace context property manipulation occurs
messaging.message.id	"C3E2D840D4D8F2F14040404040404040DED163210A4A88A1"	String Message ID represented as a hex string of 48 upper case characters.
messaging.operation.name	"MQPUT"	String Messaging operation. One of MQPUT, MQPUT1, MQGET.
messaging.operation.type	"send"	String Messaging operation type. One of: • send – when sending a message. • receive – when receiving a message using MQGET. • process – when receiving a message using async consume.
messaging.system	"ibmmq"	String Indicates that IBM MQ was the messaging system that generated the span.
name	"MQPUT Q1"	String The messaging operation and the destination it was performed on. Version 8.0.0.0 If messaging.destination.name is longer than 128 bytes, which might be the case for topics, the value used in the name attribute will be truncated at exactly 128 bytes.
os.type	"z_os"	String Type of operating system the queue manager is running on.
service.name	"MQ21"	String The name of the queue manager emitting the span. For MQGET the queue being consumed from is local to this queue manager. For MQPUT the ultimate destination queue might, or might not be local to this queue manager.

Table 1. Common attributes for both consumer and producer spans (continued)		
Attribute name	Example value	Description
service.version	"09.04.03.00"	String IBM MQ version, encoded as VV.RR.MM.FF.
span_kind	"producer"	String "producer" on MQPUT. "consumer" on MQGET or async consume.
start_time	"2022-04-29T18:52:58.114304Z"	String For MQPUT, the time the application called MQPUT. For MQGET, or async consume, the time when the message was taken off the queue.
W3C.tracestate	"rojo=00f067aa0ba902b7,congo=t61rcWkgMzE."	String. Value of tracestate header if in message. Not included otherwise.

Producer span attributes

Table 2. Producer span attributes		
Attribute name	Example value	Description
messaging.ibmmq.destination.transmission_queue	"MQ21"	String The name of the transmission queue that the message will be put on next. Only included if the message will be put to a transmission queue.
messaging.ibmmq.message.delay	5000	Integer JMS 2 delivery delay set on the message, in milliseconds. Only included if delivery delay is used.
messaging.ibmmq.message.original_conversation_id		String In scenarios where a queue manager generated MQMD is used, this is the CorrelID from the application provided MQMD. Not set otherwise.
messaging.ibmmq.message.original_message_id		String In scenarios where a queue manager generated MQMD is used, this is the MessageID from the application provided MQMD. Not set otherwise.
messaging.ibmmq.message.retained_publication	true	Boolean Set to true if the message being published is a retained publication. Not set otherwise.

Table 2. Producer span attributes (continued)		
Attribute name	Example value	Description
messaging.ibmmq.message.streamed_from	true	Boolean Set to <code>true</code> if a copy of this message has been sent to a different queue as a result of STREAMQ being set on this queue. Not set otherwise. If set, there might be another span with the same message ID and correlation ID.
messaging.ibmmq.message.streamed_to	true	Boolean Set to <code>true</code> if this message has been put to this queue as a result of STREAMQ being set on another queue. Not set otherwise. If set, there might be a span with the same message ID and correlation ID.
messaging.ibmmq.producer.publication_fail_count	0	Integer. Number of delivery failures to subscribers. Can only be non-zero if topic is configured with (N)PMSGDLV(ALLAVAIL ALLDUR). Only included in the “publish” span when publishing to a topic, either directly, or via an alias queue.
messaging.ibmmq.producer.subscriber_count	5	Integer. Number of matching subscribers, which may be zero. Only included in the “publish” span when publishing to a topic, either directly, or via an alias queue.

Consumer span attributes

Table 3. Consumer span attributes		
Attribute name	Example value	Description
messaging.ibmmq.consumer.administrative_subscription	false	Boolean Whether this is an administrative subscription (<code>true</code>) or not (<code>false</code>).
messaging.ibmmq.consumer.durable_subscription	false	Boolean Whether this is a durable subscription (<code>true</code>) or not (<code>false</code>).
messaging.ibmmq.consumer.managed_subscription	true	Boolean Whether this is a managed subscription (<code>true</code>) or not (<code>false</code>).

Table 3. Consumer span attributes (continued)

Attribute name	Example value	Description
messaging.ibmmq.consumer.no_child_spans	true	Boolean If set to true, this indicates that there can be no child spans of this consumer span as trace context cannot be propagated to the getting application. That is: • OTELPCTL(AUTO) and the application is in a version of CICS or IMS that does not support automatic propagation, or in which automatic propagation has been disabled. • OTELPCTL(AUTO) and the application is not running in CICS or IMS and has specified MQGMO_PROPERTIES_AS_Q_DEF or MQGMO_NO_PROPERTIES. • OTELPCTL(MANUAL) and the getting application specifies MQGMO_PROPERTIES_AS_Q_DEF and the queue is configured with PROPCTL(NONE). • OTELPCTL(MANUAL) and the getting application specifies MQGMO_NO_PROPERTIES. Never set to false.
messaging.ibmmq.consumer.original_wait_interval	1000	Integer The wait interval, in milliseconds, that the caller originally specified on the MQGET. Only set if get with wait used. -1 indicates wait forever. Not set when the application is using async-consume.
messaging.ibmmq.consumer.subscription_id	C3E2D840D4D8F2F14040404040404040DED163210A4A88A1	String Subscription ID from DIS SBSTATUS.
messaging.ibmmq.consumer.used_wait_interval	400	Integer How long the caller had to wait in milliseconds before a message was returned. Only in span if get with wait used. Only set if get with wait used. -1 indicates wait forever. Not set when the application is using async-consume.
messaging.ibmmq.message.delivery_count	1	Integer Value of MQMD.BackoutCount + 1.
messaging.ibmmq.message.time_on_queue	400000	Integer How long the message was on the queue, in microseconds. 0 means that put with waiting getter was used and the queue was bypassed. If delivery delay was used, this value does not include the delivery delay time.

IMS bridge client span attributes

Table 4. IMS bridge client span attributes		
Attribute name	Example value	Description
end_time	"2025-10-29T18:52:58.114304Z"	String The time when the IMS bridge received an acknowledgement from IMS that the message data has been received. The acknowledgement will indicate whether the message data can, or cannot, be processed. If the message data cannot be processed then the error.type , otel.status_code and ibm.mq.ims_bridge.reason_code attributes will be set indicating why.
error.type	"00F2020A"	String An 8 character upper case hexadecimal code. Either an OTMA sense code with four leading zeros, or the value 00F2020A indicating that the message could not be converted to OTMA format. Only set if the IMS bridge detects an error, or receives an error back from IMS during initial message transfer.
ibm.mq.ims_bridge.cant_propagate	true	Boolean Set to true if trace context cannot be propagated through the IMS bridge. That is, there was not space to propagate trace context to IMS, or to any reply or report messages generated by the IMS bridge. Not set otherwise. If this value is set in a span, the QMSTOBCD attribute is incremented.
ibm.mq.ims_bridge.ims_member_name	"IYEAZI5A"	String The name of the IMS member that the IMS bridge is connected to.
ibm.mq.ims_bridge.ims_not_otel_enabled	true	Boolean Set to true if the IMS bridge detects that IMS is not enabled for OTel. Not set otherwise.
ibm.mq.ims_bridge.reason_code	"00001B25"	String An 8 character upper case hexadecimal reason code associated with error.type . For more information, see OTMA sense codes for NAK messages in the IBM Information Management System (IMS) documentation .
ibm.mq.ims_bridge.tpipe_name	"M210004C"	String The name of the IMS TPIPE being used by the IMS bridge.
ibm.mq.ims_bridge.transaction_name	"TRAN1"	The name of the IMS transaction that will process the message data. Not set if the message data is going to be sent to an IMS command.

Table 4. IMS bridge client span attributes (continued)		
Attribute name	Example value	Description
name	"MQ IMS Bridge: IYEAZI5A"	String The name of the span. Formed by "MQ IMS Bridge: " with the IMS member name appended.
os.type	"z_os"	String Always z_os.
otel.status_code	"ERROR"	String Indicates that an error has occurred. If set, then the error.type and ibm.mq.ims_bridge.reason_code attributes provide information about the error.
service.name	"MQ21"	String The name of the queue manager emitting the span.
service.version	"10.00.00.00"	String IBM MQ version, encoded as VV.RR.MM.FF.
span_kind	"client"	String Always "client".
start_time	2025-10-29T18:52:58.114004Z"	String The time the IMS bridge issued the MQGET call.
W3C.tracestate	"rojo=00f067aa0ba902b7,congo=t61rcWkgMzE."	String. The value of the tracestate header if in the message. Not included otherwise.

Multi Queue manager health check behavior

The queue manager carries out periodic health checks to ensure stable and reliable performance. This topic describes some of the health checks that the queue manager makes and explains how they can be configured based on environmental requirements.

In most environments, the default configuration is suitable and there is no need to change how frequently these checks are made. Even using the default settings, it can be useful to understand how the queue manager behaves when a problem is detected and the environmental problems that might cause a check to fail. This topic is intended to explain some of those behaviors.



CAUTION: Do not make any changes to the frequency of these checks unless you are advised to do so by IBM Support.

The different components of the queue manager use various means to detect and resolve inconsistencies and this topic is not intended to describe all such mechanisms. For example, IBM MQ processes use various mechanisms to ensure that other processes that they depend on are still running. The behaviors that are described are those behaviors that are made periodically by the execution controller to spot environmental or other unexpected situations. (The execution controller is the primary IBM MQ process that starts and manages most other queue manager processes.) Because they are periodic checks, they are made at specific intervals that can be modified to some extent by setting the appropriate tuning parameters.

Some of the checks that are described are made by a dedicated health checking thread. If a problem is detected with the health checking thread itself, a warning message [AMQ5066](#) is written to the queue manager error logs.

The behaviors that are described in this topic are subject to change in future releases, for example if a different default value is observed to be more stable on a particular platform or configuration.

General health checks

The queue manager performs a variety of checks on a regular interval. By default, these checks are performed every 10 seconds (in some cases, the check allows two cycles before reporting an error, resulting in a 20-second interval for such checks). During the checks, the queue manager makes sure that the various processes that run as part of the queue manager are still running. For a Native HA queue manager, it checks that the queue manager is successfully replicating data to the standby instances.

If a critical check fails at this time (for example, if the `amqzmuc0` process is no longer running) the queue manager will be unable to continue running. However, most of the checks are made to tidy up system resources that are no longer needed and might simply result in a message being written to the queue manager error logs.

In most cases, it is not necessary to change the frequency of these general health checks. The majority of events in the queue manager or in the operating environment are detected instantly without the need for the general health check process to detect them. This process serves as a periodic check for anything not detected elsewhere in the queue manager. If necessary, the frequency can be configured using the **ECHearTBeatLen** tuning parameter. The minimum value is 10000 milliseconds (10 seconds). The maximum value is 60000 milliseconds (60 seconds). If set to its maximum value of 60000, this might result in a two-minute delay for certain checks.

Checks that log progress is being made

The queue manager checks that writes to the log are being made at a reasonable rate. This is not a check that performance of the logger is optimal but is designed to spot conditions that might need further attention. For example, if the disk that the log files are being stored on is particularly slow, or if the queue manager is not receiving sufficient CPU time in a containerized environment to perform all of its work.

If this check fails, the action taken by the queue manager depends on the type of queue manager being used:

- On a non-HA queue manager:
 - An `xecL_W_PERFORMANCE_BOTTLENECK FDC` is written. This can be used as an indication that some part of the system might need further attention. The queue manager remains running. If `xecL_W_PERFORMANCE_BOTTLENECK FDCs` are seen in the errors directory, it might be necessary to work with your storage or platform team to understand if the underlying system resources are sufficient for IBM MQ to run. If IBM MQ is being run in containers on over-committed nodes then IBM MQ might not receive enough scheduled CPU time to perform all of its messaging workload.
 - From IBM MQ 9.3.0, a warning message [AMQ5068W](#) is written to the queue manager error logs and no `xecL_W_PERFORMANCE_BOTTLENECK FDC` is written. If [AMQ5068W](#) messages are seen in the logs, it might be necessary to work with your storage or platform team to understand if the underlying system resources are sufficient for IBM MQ to run. If IBM MQ is being run in containers on over-committed nodes then IBM MQ might not receive enough scheduled CPU time to perform all of its messaging workload. If five [AMQ5068W](#) warning messages are written in succession, an `xecL_W_PERFORMANCE_BOTTLENECK FDC` is written.
- On a multi-instance queue manager:
 - If the log progress health check fails then the primary instance ends. If a standby instance is available it will start and become the primary instance.
 - From IBM MQ 9.3.0, the primary instance checks if a standby instance is available before ending. If a standby queue manager is available to fail over to the primary instance ends. Additionally a warning message [AMQ5068W](#) is written to the queue manager error logs.

- On a Native HA queue manager, this check behaves in the same way as a non-HA queue manager.
- On an RDQM (replicated data queue manager), this check behaves in the same way as a non-HA queue manager.

It is possible that a problem with the progression of the IBM MQ log is caused by a performance issue in the queue manager itself.

By default this check is made every 60 seconds, although the queue manager waits for two cycles of the check before taking an action. This means that with the default settings, two minutes would have to pass before the queue manager wrote an error message (or failed over in the case of an HA queue manager).

In most cases, the default behavior is suitable even where the file system is slow or the queue manager is allocated a small amount of CPU time because other checks such as file locking (see [“Checks that file locks are still held”](#) on page 29) and basic file system operation will cause a primary instance to failover before this check is made. If necessary, the frequency of this check can be configured using the **LivenessHeartBeatLen** tuning parameter. The maximum value that it can be configured to is 600 seconds (10 minutes). The minimum value of 0 has the effect of disabling the check altogether. For a non-HA queue manager, the only effect of the check is an extra warning message in the queue manager error logs. For a multi-instance queue manager, you can configure the **LivenessHeartBeatLen** to cause a primary instance of the queue manager to fail over more quickly (by reducing the value) or slowly (by increasing the value). Increasing the value to reduce the frequency of the log progress check can be useful if your environment occasionally experiences very slow file system IO but where you would prefer the primary instance of the queue manager to remain running. This might be useful if you have applications that are not designed to automatically reconnect to the standby instance and require manual intervention to restart them.

Note: If the **ECHeartBeatLen** has been increased, this affects the timing of the **LivenessHeartBeatLen** checks. Log progress checks are made when the general health checks are performed so reducing the frequency of the general health checks (**ECHeartBeatLen**) might result in log progress checks being made up to a 30 seconds after the configured **LivenessHeartBeatLen**.

General logging filesystem performance

From IBM MQ 9.4.0, a warning message [AMQ6729W](#) is issued in the queue manager error log if general read/write operations to this storage appear to be taking longer than expected. You can use the **AMQ_IODELAY** environment variables to fine tune diagnostics and timings to either aid in diagnosing storage performance issues, or to increase tolerance for such delays. For more information, see [AMQ_IODELAY](#), [AMQ_IODELAY_INMS](#) and [AMQ_IODELAY_FFST](#).

Checks that file locks are still held

For a multi-instance queue manager, the execution controller periodically checks to make sure that it still holds the exclusive lock on the primary multi-instance file. In many cases, if the lock is lost due to a problem with the NFS server the primary instance fails over almost immediately (before this check is made). Additional periodic file lock checks are made to ensure that the primary queue manager fails over in the event of an unusual file system problem.

By default, these file lock checks are made every 20 seconds. If necessary, this value can be changed by setting the **FileLockHeartBeatLen** tuning parameter. The default value for the tuning parameter is 10 seconds (the queue manager allows two cycles of the check before taking an action resulting in the default behavior of checking every 20 seconds). The minimum value of the tuning parameter is 10 seconds, the maximum value is 600 seconds (10 minutes).

Note: If the **ECHeartBeatLen** has been increased, this affects the timing of the **FileLockHeartBeatLen** checks. File lock checks are made when the general health checks are performed so reducing the frequency of the general health checks (**ECHeartBeatLen**) might result in file lock checks being made up to 30 seconds after the configured **FileLockHeartBeatLen**.

Checks on user application health

The queue manager periodically checks that any locally bound applications that are no longer running have performed an MQDISC MQI call before terminating. These checks are performed at the same time as the general health checks described in [“General health checks” on page 28](#). The default interval for such checks is therefore 10000 milliseconds (10 seconds) and changing the value of the **ECHearTBeatLen** tuning parameter changes the frequency at which they are made. This check is primarily to ensure any resources that are associated with an application connected are freed, it does not cause an HA or a non-HA queue manager to end or to fail over to an alternative instance.

IBM MQ client applications that have terminated without issuing an MQDISC MQI call are separately detected by the agent process and any resources associated with the connection are released.

Related concepts

[Configuring high availability](#)

Event monitoring

Event monitoring is the process of detecting occurrences of *instrumentation events* in a queue manager network. An instrumentation event is a logical combination of conditions that is detected by a queue manager or channel instance. Such an event causes the queue manager or channel instance to put a special message, called an *event message*, on an event queue.

IBM MQ instrumentation events provide information about errors, warnings, and other significant occurrences in a queue manager. Use these events to monitor the operation of the queue managers in your queue manager network to achieve the following goals:

- Detect problems in your queue manager network.
- Assist in determining the causes of problems in your queue manager network.
- Generate an audit trail.
- React to queue manager state changes

Related reference

[“Event types” on page 33](#)

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

[Event message reference](#)

[Event message format](#)

Publishing your IBM MQ event messages

How you prepare IBM MQ to publish event messages.

About this task

Event messages are written to specially named queues called `SYSTEM.ADMIN.<feature name>.EVENT`.

The important thing to note about these event queues is that it is the name that matters. By default, on a queue manager, all event queues are defined as local queues. However, you can delete these queues and redefine them, perhaps as a remote queue, so that all events are funneled to a dedicated event processing queue manager. Alternatively, you can use an alias queue that is pointing at a topic object.

In either case, any redirection technique requires that your applications reading the event queues have not hard-coded the name of the queue to read from. Therefore, you must be able to configure the queue the applications are reading from.

The following commands show how you can redefine your event queues so that the event messages will be published, using the following assumptions. You have:

- Not started using events, or

- Removed all the messages from the existing event queues and have deleted the local queues prior to these steps.

These steps only show the QMGR and CHANNEL event queues being redefined, but this could be extended for all events.

Note: The topic string is designed so that an application can be subscribed to all events using a wildcard, or to specific events, as required.

Procedure

Issue the following commands:

```
DEFINE TOPIC(ADMIN.QMGR.EVENT) TOPICSTR('Events/QMgr')
DEFINE TOPIC(ADMIN.CHANNEL.EVENT) TOPICSTR('Events/Channel')

DEFINE QALIAS(SYSTEM.ADMIN.QMGR.EVENT) TARGETTYPE(TOPIC) TARGET(ADMIN.QMGR.EVENT)
DEFINE QALIAS(SYSTEM.ADMIN.CHANNEL.EVENT) TARGETTYPE(TOPIC) TARGET(ADMIN.CHANNEL.EVENT)

DEFINE QLOCAL(ADMIN.EVENT)
DEFINE QLOCAL(ADMIN.QMGR.EVENT)

DEFINE SUB(EVENTS.ALL) TOPICSTR('Events/+') PSROP(NONE)
DESTCLAS(PROVIDED) DEST(ADMIN.EVENT)
DEFINE SUB(EVENTS.QMGR) TOPICSTR('Events/QMgr') PSROP(NONE)
DESTCLAS(PROVIDED) DEST(ADMIN.QMGR.EVENT)
```

Assuming that your event reading application is able to read event messages from any queue, that application can be reconfigured to read from one of the queues defined above as required.

The `PSROP(NONE)` configuration on the `DEFINE SUB` commands is to ensure that none of the message properties added by the publish/subscribe engine, for example `MQTopicString`, is added to the event message, ensuring that existing applications can continue to work unchanged.

Additionally, applications can also subscribe directly using the `MQSUB` call to receive the information, as an alternate way instead of using the administrative `DEFINE SUB` command.

Now, multiple applications are able to consume the information emitted in events by the queue manager.

Instrumentation events

An instrumentation event is a logical combination of conditions that a queue manager or channel instance detects and puts a special message, called an *event message*, on an event queue.

IBM MQ instrumentation events provide information about errors, warnings, and other significant occurrences in a queue manager. You can use these events to monitor the operation of queue managers (with other methods such as Tivoli® NetView for z/OS).

Figure 5 on page 32 illustrates the concept of instrumentation events.




Figure 5. Understanding instrumentation events

Event monitoring applications

Applications that use events to monitor queue managers must include the following provisions:

1. Set up channels between the queue managers in your network.

2. Implement the required data conversions. The normal rules of data conversion apply. For example, if you are monitoring events on an UNIX system queue manager from a z/OS queue manager, ensure that you convert EBCDIC to ASCII.

Event notification through event queues

When an event occurs, the queue manager puts an event message on the appropriate event queue, if defined. The event message contains information about the event that you can retrieve by writing a suitable MQI application program that performs the following steps:

- Get the message from the queue.
- Process the message to extract the event data.

The related information describes the format of event messages.

Conditions that cause events

The following list gives examples of conditions that can cause instrumentation events:

- A threshold limit for the number of messages on a queue is reached.
- A channel instance is started or stopped.
- A queue manager becomes active, or is requested to stop.
- An application tries to open a queue specifying a user ID that is not authorized on IBM MQ for IBM i, AIX, Linux, and Windows systems.
- Objects are created, deleted, changed, or refreshed.
- An MQSC or PCF command runs successfully.
- A queue manager starts writing to a new log extent.
- Putting a message on the dead-letter queue, if the event conditions are met.

Related concepts

[“Performance events” on page 47](#)

Performance events relate to conditions that can affect the performance of applications that use a specified queue. The scope of performance events is the queue. **MQPUT** calls and **MQGET** calls on one queue do not affect the generation of performance events on another queue.

[“Sample program to monitor instrumentation events on Multiplatforms” on page 79](#)

The `amqsevt` sample program formats the instrumentation events that a queue manager can create, and is supplied with IBM MQ for Multiplatforms. The program reads messages from event queues, and formats them into readable strings.

Event types

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

IBM MQ instrumentation events have the following types:

- Queue manager events
- Channel and bridge events
- Performance events
- Configuration events
- Command events
- Logger events
- Local events

For each queue manager, each category of event has its own event queue. All events in that category result in an event message being put onto the same queue.

This event queue:

SYSTEM.ADMIN.QMGR.EVENT
SYSTEM.ADMIN.CHANNEL.EVENT
SYSTEM.ADMIN.PERFM.EVENT
SYSTEM.ADMIN.CONFIG.EVENT
SYSTEM.ADMIN.COMMAND.EVENT
SYSTEM.ADMIN.LOGGER.EVENT
SYSTEM.ADMIN.PUBSUB.EVENT

Contains messages from:

Queue manager events
Channel events
Performance events
Configuration events
Command events
Logger events
Gets events related to Publish/Subscribe. Only used with Multicast. For more information see, [Multicast application monitoring](#).

By incorporating instrumentation events into your own system management application, you can monitor the activities across many queue managers, across many different nodes, and for multiple IBM MQ applications. In particular, you can monitor all the nodes in your system from a single node (for those nodes that support IBM MQ events) as shown in [Figure 6 on page 34](#).

Instrumentation events can be reported through a user-written reporting mechanism to an administration application that can present the events to an operator.




Figure 6. Monitoring queue managers across different platforms, on a single node

Instrumentation events also enable applications acting as agents for other administration networks, for example Tivoli NetView for z/OS, to monitor reports and create the appropriate alerts.

Queue manager events

Queue manager events are related to the use of resources within queue managers. For example, a queue manager event is generated if an application tries to put a message on a queue that does not exist.

The following examples show conditions that can cause a queue manager event:

- An application issues an MQI call that fails. The reason code from the call is the same as the reason code in the event message.

A similar condition can occur during the internal operation of a queue manager; for example, when generating a report message. The reason code in an event message might match an MQI reason code, even though it is not associated with any application. Do not assume that, because an event message

reason code looks like an MQI reason code, the event was necessarily caused by an unsuccessful MQI call from an application.

- A command is issued to a queue manager and processing this command causes an event. For example:
 - A queue manager is stopped or started.
 - A command is issued where the associated user ID is not authorized for that command.

IBM MQ puts messages for queue manager events on the SYSTEM.ADMIN.QMGR.EVENT queue, and supports the following queue manager event types:

Authority (on Multiplatforms only)

V 10.0.0 From IBM MQ 10.0.0, authority events can report successful and unsuccessful authorizations. Authority events report an authorization, such as an application trying to open a queue for which it does or does not have the required authority, or a command being issued from a user ID that does not have the required authority. Authority events are controlled through the queue manager **AUTHOREV** and **AUTHEVSC** settings.

The authority event message can contain the following event data.

The following authority events are produced so long as the queue manager authority event setting (**AUTHOREV**) is ENABLED:

- [Not Authorized \(type 1\)](#)
- [Not Authorized \(type 2\)](#)
- [Not Authorized \(type 3\)](#)
- [Not Authorized \(type 4\)](#)
- [Not Authorized \(type 5\)](#)
- [Not Authorized \(type 6\)](#)

V 10.0.0 The following authority events are produced as long as the queue manager authority event setting (**AUTHOREV**) is ENABLED and the queue manager authority event scope setting (**AUTHEVSC**) is at least ALLCONNS:

- [Authorized \(type 1\)](#)

V 10.0.0 The following authority events are produced as long as the queue manager authority event setting (**AUTHOREV**) is ENABLED and the queue manager authority event scope setting (**AUTHEVSC**) is ALLCHECKS:

- [Authorized \(type 2\)](#)
- [Authorized \(type 3\)](#)
- [Authorized \(type 4\)](#)

V 10.0.0 Authorized (type 2) event messages could be produced in large quantities where a lot the event messages produced are duplicates of each other. To reduce the number of duplicates a system to rate limit duplicate event messages is available. For more information see Limiting the rate of positive authority events.

V 10.0.0 Authority events can be linked together by examining the MQBACF_CONN_TAG and MQBACF_CONNECTION_ID fields. These fields are identical for event messages produced by the same application connection.

All authority events are valid on Multiplatforms only.

Inhibit

Inhibit events indicate that an **MQPUT** or **MQGET** operation has been attempted against a queue where the queue is inhibited for puts or gets, or against a topic where the topic is inhibited for publishes. Inhibit events are controlled through the queue manager **INHIBITEV** setting. The inhibit event message can contain the following event data:

- [Get Inhibited](#)
- [Put Inhibited](#)

Local

When an application or the queue manager has not been able to access a local queue or other local object, for example, because the object has not been defined, the queue manager can generate a local event message. Local events are controlled through the queue manager **LOCALEV** setting. The local event message can contain the following event data:

- [Alias Base Queue Type Error](#)
- [Unknown Alias Base Queue](#)
- [Unknown Object Name](#)

Remote

When an application or the queue manager cannot access a remote queue on another queue manager, for example, the transmission queue is not correctly defined, the queue manager can generate a remote event message. Remote events are controlled through the queue manager **REMOTEEV** setting. The remote event message can contain the following event data:

- [Default Transmission Queue Type Error](#)
- [Default Transmission Queue Usage Error](#)
- [Queue Type Error](#)
- [Remote Queue Name Error](#)
- [Transmission Queue Type Error](#)
- [Transmission Queue Usage Error](#)
- [Unknown Default Transmission Queue](#)
- [Unknown Remote Queue Manager](#)
- [Unknown Transmission Queue](#)

Start and stop

Start and stop events indicate that a queue manager has been started or has been requested to stop or quiesce. Start and Stop events are controlled through the queue manager **STRSTPEV** setting.

 z/OS supports only start events.

Stop events are not recorded unless the default message-persistence of the SYSTEM.ADMIN.QMGR.EVENT queue is defined as persistent. The start and stop event message can contain the following event data:

- [Queue Manager Active](#)
- [Queue Manager Not Active](#)

For each event type in this list, you can set a queue manager attribute to enable or disable the event type.

 *Limiting the rate of positive authority events*

To reduce the number of duplicate events, a rate limiting system is available for Authorized (type 2) event messages.

Enabling positive authority events on particularly active queue managers might cause a large number of duplicate authority events. In cases where a connected application is using the MQPUT1 API call or repeatedly calling MQOPEN, an [Authorized \(type 2\)](#) event is generated for each call. These authority events can be identical except for the MQMD fields. On queue managers with large number of active applications, this might cause an excessive number of authority events. A rate limiting system is available to reduce the number of duplicate events.

Overview

When enabled, an IBM MQ queue manager prevents events that can be rate limited from being published when it detects that the event is a duplicate to one previously released. A duplicate event is an event that has the same MQCFH header and Event Data values.

When a duplicate event is attempted to be published, a decision on whether to publish the event is made based on the time since the event was last published and the IBM MQ configuration settings. If sufficient time has passed since a copy of that event was published, a duplicate event is published, otherwise it is discarded. When a duplicate event is published, it is put on the event queue along with the time when it was last seen and the number of duplicate events that were seen since then. The following fields are added to the event message to represent the first seen time and duplicate count:

EventDuplicateCount

Description:	How many duplicate copies of this message have been processed. Range can be determined by examining the MQMD.PutTime , MQMD.PutDate , and EventDuplicateFrom values.
Identifier:	MQIACF_EVENT_DUPLICATE_COUNT
Data type:	MQCFIN

EventDuplicateFrom

Description:	The time of the last duplicate event, in ISO 8601 format.
Identifier:	MQCACF_EVENT_DUPLICATE_FROM
Data type:	MQCFST
Maximum length:	MQ_ISOTIME_LENGTH

To determine the range that the duplicate message applies for compare the **EventDuplicateFrom** value and the **MQMD.PutTime** and **MQMD.PutDate** values.

Configuration

The configuration setting to specify the amount of time between duplicate events is supplied in the queue manager configuration file's [Tuning Parameters](#) stanza:

- **EventMinDuplicationDelayTime** controls the minimum amount of time before a duplicate event message should be published. After this point, a duplicate event message will be published, resetting the time before the next duplicate event. A setting of -1, which is the default, indicates that duplicate event messages should not be released. A setting of 0 indicates that every event message should be released.

Limitations and considerations

The following limitations and considerations should be made for the event rate limiting system:

- If the event queue becomes full, then the further event messages are not able to be put on the event queue. If an attempt to place a message fails, an AMQ7375 error message is produced to alert you, but the message is discarded. As such, it is important to ensure that the event queue is being consumed adequately to prevent it from becoming full and therefore prevent event messages from being lost. For more information, see [Event queues](#).

Example

This example assumes that:

- An application has connected to a queue manager and is putting messages to a queue using MQPUT1 at a rate of 1 message a second.

The application MQPUT1 call is equivalent to a MQOPEN, MQPUT and MQCLOSE call. For this example the MQOPEN call causes an event message to be generated.

- The configuration has authority events enabled and the scope of the authority events set to all checks.
- The rate limiting feature is configured so that duplicate messages should not be repeated until after 10 seconds has elapsed.

1. The first MQPUT1 call by the application causes a OPEN event message that is not a duplicate because it is the first instance, so as such, the message is released to the event queue. The **EventDuplicateFrom** and **EventDuplicateCount** fields are not present.
2. The second MQPUT1 call by the application causes a OPEN event message that is a duplicate of the first and has been put one second after the last. This is before the 10 second time and so the event message is discarded. No message is released to the event queue.
3. Over the next 10 seconds, the application continues to put messages, and each duplicate causes the internal count of the duplicates to increase.
4. At the 10 second mark, the application makes an MQPUT1 call. This causes a duplicate OPEN event message. However, it has been 10 seconds since the last copy of this event message has been released and 10 seconds is what **EventMinDuplicationDelayTime** is set to:
 - a. Before the duplicate message is released, the **EventDuplicateFrom** and **EventDuplicateCount** fields are appended to the event message where the **EventDuplicateFrom** equals the UTC ISO time at step 1 and the **EventDuplicateCount** field equals the number of duplicate messages held back (10).
 - b. The event message is then published to the event queue. The time since the event was last published and count of the event messages are reset.
5. As the application continues to put messages, the rate limiting system continues to increment the count of the held duplicate message and to release the message when the time exceeds or equals the **EventMinDuplicationDelayTime** value.

If the application stops making MQPUT1 calls, no further event messages are released unless the application resumes MQPUT1 calls.

Channel and bridge events

Channels report these events as a result of conditions detected during their operation. For example, when a channel instance is stopped.

Channel events are generated in the following circumstances:

- When a command starts or stops a channel.
- When a channel instance starts or stops.
- When a channel receives a conversion error warning when getting a message.
- When an attempt is made to create a channel automatically; the event is generated whether the attempt succeeds or fails.

Note: Client connections do not cause Channel Started or Channel Stopped events.

When a command is used to start a channel, an event is generated. Another event is generated when the channel instance starts. However, starting a channel by a listener, the **runmqchl** command, or a queue manager trigger message does not generate an event. In these cases, an event is generated only when the channel instance starts.

A successful start or stop channel command generates at least two events. These events are generated for both queue managers connected by the channel (providing they support events).

If a channel event is put on an event queue, an error condition causes the queue manager to create an event.

The event messages for channel and bridge events are put on the SYSTEM.ADMIN.CHANNEL.EVENT queue.

The channel event messages can contain the following event data:

- Channel Activated
- Channel Auto-definition Error
- Channel Auto-definition OK
- Channel Blocked

- [Channel Conversion Error](#)
- [Channel Not Activated](#)
- [Channel Not Available](#)
- [Channel Started](#)
- [Channel Stopped](#)
- [Channel Stopped By User](#)

IMS bridge events (z/OS only)



These events are reported when an IMS bridge starts or stops.

The IMS bridge event messages can contain the following event data:

- [Bridge Started](#)
- [Bridge Stopped](#)

SSL events

The only TLS event is the Channel SSL Error event. This event is reported when a channel using TLS fails to establish a TLS connection.

The SSL event messages can contain the following event data:

- [Channel SSL Error](#)
- [Channel SSL Warning](#)

Performance events

Performance events are notifications that a resource has reached a threshold condition. For example, a queue depth limit has been reached.

Performance events relate to conditions that can affect the performance of applications that use a specified queue. They are not generated for the event queues themselves.

The event type is returned in the command identifier field in the message data.

If a queue manager tries to put a queue manager event or performance event message on an event queue and an error that would typically create an event is detected, another event is not created and no action is taken.

MQGET and MQPUT calls within a unit of work can generate performance events regardless of whether the unit of work is committed or backed out.

The event messages for performance events are put on the SYSTEM.ADMIN.PERFM.EVENT queue.

There are two types of performance event:

Queue depth events

Queue depth events relate to the number of messages on a queue; that is, how full or empty the queue is. These events are supported for shared queues. The queue depth event messages can contain the following event data:

- [Queue Depth High](#)
- [Queue Depth Low](#)
- [Queue Full](#)

Queue service interval events

Queue service interval events relate to whether messages are processed within a user-specified time interval. These events are not supported for shared queues.



IBM MQ for z/OS supports queue depth events for QSGDISP (SHARED) queues, but not service interval events. Queue manager and channel events remain unaffected by shared queues. The queue service event messages can contain the following event data:

- [Queue Service Interval High](#)
- [Queue Service Interval OK](#)

Configuration events

Configuration events are generated when a configuration event is requested explicitly, or automatically when an object is created, modified, or deleted.

A configuration event message contains information about the attributes of an object. For example, a configuration event message is generated if a namelist object is created, and contains information about the attributes of the namelist object.

The event messages for configuration events are put on the SYSTEM.ADMIN.CONFIG.EVENT queue.

Configuration events work in the following way:

- An audit trail of changes to the queue manager configuration is kept by means of event messages, which are written to the SYSTEM.ADMIN.CONFIG.EVENT queue. You can enable these events to be written by means of the **CONFIGEV** parameter on the [ALTER QMGR](#) command.
- These events are generated when a DEFINE, ALTER or DELETE command acts upon an object, or an MQSET call is used.
- You can create a base-line picture of the current queue manager configuration by using the [REFRESH QMGR TYPE\(CONFIGEV\)](#) command, which will create an event message for every object in the queue manager. Note that since this could be a time consuming operation, if you have many objects, you can break the task down into smaller sets of objects using the NAME and OBJECT qualifiers on the command.
- The event message records one of four possible Reasons:
 - MQRC_CONFIG_CHANGE_OBJECT
 - MQRC_CONFIG_CREATE_OBJECT
 - MQRC_CONFIG_DELETE_OBJECT
 - MQRC_CONFIG_REFRESH_OBJECT

MQRC_CONFIG_CHANGE_OBJECT, MQRC_CONFIG_CREATE_OBJECT or MQRC_CONFIG_DELETE_OBJECT occur for the respective MQSC or PCF commands that you might issue upon an object.

MQRC_CONFIG_REFRESH_OBJECT occurs for those event messages written, when creating the base-line picture.

There are four types of configuration event:

Create object events

Create object events are generated when an object is created. The event message contains the following event data: [Create object](#).

Change object events

Change object events are generated when an object is changed. The event message contains the following event data: [Change object](#).

Delete object events

Delete object events are generated when an object is deleted. The event message contains the following event data: [Delete object](#).

Refresh object events

Refresh object events are generated by an explicit request to refresh. The event message contains the following event data: [Refresh object](#).

Command events

Command events are reported when an MQSC or PCF command runs successfully.

A command event message contains information about the origin, context, and content of a command. For example, a command event message is generated with such information if the MQSC command, ALTER QLOCAL, runs successfully.

The event messages for command events are put on the SYSTEM.ADMIN.COMMAND.EVENT queue.

Command events contain the following event data: [Command](#).

Multi Logger events

Logger events are reported when a queue manager that uses linear logging starts writing log records to a new log extent  or, on IBM i, to a new journal receiver.  Logger events are not available with IBM MQ for z/OS.

A logger event message contains information specifying the log extents required by the queue manager to restart the queue manager, or for media recovery.

The event messages for logger events are put on the SYSTEM.ADMIN.LOGGER.EVENT queue.

The logger event message contains the following event data: [Logger](#).

Event message data summary

Use this summary to obtain information about the event data that each type of event message can contain.

Event type	See these topics
Authority events	Not Authorized (type 1)
	Not Authorized (type 2)
	Not Authorized (type 3)
	Not Authorized (type 4)
	Not Authorized (type 5)
	Not Authorized (type 6)
Channel events	Channel Activated
	Channel Auto-definition Error
	Channel Auto-definition OK
	Channel Blocked
	Channel Conversion Error
	Channel Not Activated
	Channel Started
	Channel Stopped
	Channel Stopped By User
Command events	Command

Event type	See these topics
Configuration events	Create object
	Change object
	Delete object
	Refresh object
IMS bridge events	Bridge Started
	Bridge Stopped
Inhibit events	Get Inhibited
	Put Inhibited
Local events	Alias Base Queue Type Error
	Unknown Alias Base Queue
	Unknown Object Name
Logger events	Logger
Performance events	Queue Depth High
	Queue Depth Low
	Queue Full
	Queue Service Interval High
	Queue Service Interval OK
Remote events	Default Transmission Queue Type Error
	Default Transmission Queue Usage Error
	Queue Type Error
	Remote Queue Name Error
	Transmission Queue Type Error
	Transmission Queue Usage Error
	Unknown Default Transmission Queue
	Unknown Remote Queue Manager
	Unknown Transmission Queue
SSL events	Channel SSL Error
Start and stop events	Queue Manager Active
	Queue Manager Not Active

Controlling events

You enable and disable events by specifying the appropriate values for queue manager, queue attributes, or both, depending on the type of event.

You must enable each instrumentation event that you want to be generated. For example, the conditions causing a Queue Full event are:

- Queue Full events are enabled for a specified queue, and
- An application issues an MQPUT request to put a message on that queue, but the request fails because the queue is full.

Enable and disable events by using any of the following techniques:

- IBM MQ script commands (MQSC).
- The corresponding IBM MQ PCF commands.
-  The operations and control panels for queue managers on z/OS.
- IBM MQ Explorer.

Note: You can set attributes related to events for both queues and queue managers only by command. The MQI call MQSET does not support attributes related to events.

Related concepts

[“Instrumentation events” on page 31](#)

An instrumentation event is a logical combination of conditions that a queue manager or channel instance detects and puts a special message, called an *event message*, on an event queue.

[Using the operations and control panels on z/OS](#)

Related tasks

[Automating administration tasks](#)

[Using Programmable Command Formats](#)

Related reference

[“Event types” on page 33](#)

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

[The MQSC commands](#)

Controlling queue manager events

You control queue manager events by using queue manager attributes. To enable queue manager events, set the appropriate queue manager attribute to ENABLED. To disable queue manager events, set the appropriate queue manager attribute to DISABLED.

To enable or disable queue manager events, use the MQSC command **ALTER QMGR**, specifying the appropriate queue manager attribute. Table 5 on page 43 summarizes how to enable queue manager events. To disable a queue manager event, set the appropriate parameter to DISABLED.

 The queue manager attribute **AUTHEVSC** is used in conjunction with the **AUTHOREV** attribute and is used to determine the scope of authority events once enabled.

Table 5. Enabling queue manager events using MQSC commands	
Event	ALTER QMGR parameter
Authority	AUTHOREV (ENABLED)
 Authority Scope	AUTHEVSC(FAILURES ALLCONNS ALLCHECKS)
Inhibit	INHIBITEV (ENABLED)
Local	LOCALEV (ENABLED)
Remote	REMOTEEV (ENABLED)
Start and Stop	STRSTPEV (ENABLED)

Controlling channel and bridge events

You control channel events by using queue manager attributes. To enable channel events, set the appropriate queue manager attribute to **ENABLED**. To disable channel events, set the appropriate queue manager attribute to **DISABLED**.

To enable or disable channels events use the MQSC command **ALTER QMGR**, specifying the appropriate queue manager attribute. Table 6 on page 44 summarizes how you enable channel and bridge events. To disable a queue manager event, set the appropriate parameter to **DISABLED**.

Restriction:  Channel auto-definition events are not available on IBM MQ for z/OS.

Table 6. Enabling channel and bridge events using MQSC commands	
Event	ALTER QMGR parameter
Channel	CHLEV (ENABLED)
Related to channel errors only	CHLEV (EXCEPTION)
IMS bridge	BRIDGEEV (ENABLED)
SSL	SSLEV (ENABLED)
Channel auto-definition	CHADEV(ENABLED)

With CHLEV set to exception, the following return codes, and corresponding reason qualifiers are generated:

- MQRC_CHANNEL_ACTIVATED
- MQRC_CHANNEL_CONV_ERROR
- MQRC_CHANNEL_NOT_ACTIVATED
- MQRC_CHANNEL_STOPPED
 - with the following ReasonQualifiers:
 - MQRQ_CHANNEL_STOPPED_ERROR
 - MQRQ_CHANNEL_STOPPED_RETRY
 - MQRQ_CHANNEL_STOPPED_DISABLED
- MQRC_CHANNEL_STOPPED_BY_USER
- MQRC_CHANNEL_BLOCKED
 - with the following ReasonQualifiers:
 - MQRQ_CHANNEL_BLOCKED_NOACCESS
 - MQRQ_CHANNEL_BLOCKED_USERID
 - MQRQ_CHANNEL_BLOCKED_ADDRESS

Controlling performance events

You control performance events using the PERFMEV queue manager attribute. To enable performance events, set PERFMEV to **ENABLED**. To disable performance events, set the PERFMEV queue manager attribute to **DISABLED**.

To set the PERFMEV queue manager attribute to **ENABLED**, use the following MQSC command:

```
ALTER QMGR PERFMEV (ENABLED)
```

To enable specific performance events, set the appropriate queue attribute. Also, specify the conditions that cause the event.

Queue depth events

By default, all queue depth events are disabled. To configure a queue for any of the queue depth events:

1. Enable performance events on the queue manager.

2. Enable the event on the required queue.
3. Set the limits, if required, to the appropriate levels, expressed as a percentage of the maximum queue depth.

Queue service interval events

To configure a queue for queue service interval events you must:

1. Enable performance events on the queue manager.
2. Set the control attribute for a Queue Service Interval High or OK event on the queue as required.
3. Specify the service interval time by setting the QSVCINT attribute for the queue to the appropriate length of time.

Note: When enabled, a queue service interval event can be generated at any appropriate time, not necessarily waiting until an MQI call for the queue is issued. However, if an MQI call is used on a queue to put or remove a message, any applicable performance event is generated at that time. The event is not generated when the elapsed time becomes equal to the service interval time.

Controlling configuration, command, and logger events

You control configuration, command, and logger events by using the queue manager attributes CONFIGEV, CMDEV, and LOGGEREV. To enable these events, set the appropriate queue manager attribute to ENABLED. To disable these events, set the appropriate queue manager attribute to DISABLED.

Configuration events

To enable configuration events, set CONFIGEV to ENABLED. To disable configuration events, set CONFIGEV to DISABLED. For example, you can enable configuration events by using the following MQSC command:

```
ALTER QMGR CONFIGEV (ENABLED)
```

Command events

To enable command events, set CMDEV to ENABLED. To enable command events for commands except DISPLAY MQSC commands and Inquire PCF commands, set the CMDEV to NODISPLAY. To disable command events, set CMDEV to DISABLED. For example, you can enable command events by using the following MQSC command:

```
ALTER QMGR CMDEV (ENABLED)
```

Logger events

To enable logger events, set LOGGEREV to ENABLED. To disable logger events, set LOGGEREV to DISABLED. For example, you can enable logger events by using the following MQSC command:

```
ALTER QMGR LOGGEREV(ENABLED)
```

Event queues

When an event occurs, the queue manager puts an event message on the defined event queue. The event message contains information about the event.

You can define event queues, either as:

- Local queues
- Alias queues
- Local definitions of remote queues, or as
- Remote cluster queues

If you define all your event queues as local definitions of the same remote queue on one queue manager, you can centralize your monitoring activities.

You must not define event queues as transmission queues, because event messages have formats that are incompatible with the message format that is required for transmission queues.

Shared event queues are local queues defined with the QSGDISP(SHARED) value.

For more information about defining shared queues on z/OS, see [Application programming with shared queues](#).

When an event queue is unavailable

If an event occurs when the event queue is not available, the event message is lost. For example, if you do not define an event queue for a category of event, all event messages for that category are lost. The event messages are not, for example, saved on the dead-letter (undelivered-message) queue.

However, you can define the event queue as a remote queue. Then, if there is a problem on the remote system putting messages to the resolved queue, the event message arrives on the dead-letter queue of the remote system.

An event queue might be unavailable for many different reasons including:

- The queue has not been defined.
- The queue has been deleted.
- The queue is full.
- The queue has been put-inhibited.

The absence of an event queue does not prevent the event from occurring. For example, after a performance event, the queue manager changes the queue attributes and resets the queue statistics. This change happens whether the event message is put on the performance event queue or not. The same is true in the case of configuration and command events.

Using triggered event queues

You can set up the event queues with triggers so that when an event is generated, the event message being put onto the event queue starts a user-written monitoring application. This application can process the event messages and take appropriate action. For example, certain events might require an operator to be informed, other events might start an application that performs some administration tasks automatically.

Event queues can have trigger actions associated with them and can create trigger messages. However, if these trigger messages in turn cause conditions that would normally generate an event, no event is generated. not generating an event in this instance ensures that looping does not occur.

Related concepts

[“Controlling events” on page 42](#)

You enable and disable events by specifying the appropriate values for queue manager, queue attributes, or both, depending on the type of event.

[“Format of event messages” on page 46](#)

Event messages contain information about an event and its cause. Like other IBM MQ messages, an event message has two parts: a message descriptor and the message data.

[Application programming with shared queues](#)

[Conditions for a trigger event](#)

Related reference

[QSGDisp \(MQLONG\)](#)

Format of event messages

Event messages contain information about an event and its cause. Like other IBM MQ messages, an event message has two parts: a message descriptor and the message data.

- The message descriptor is based on the MQMD structure.
- The message data consists of an *event header* and the *event data*. The event header contains the reason code that identifies the event type. Putting the event message, and any subsequent action, does not affect the reason code returned by the MQI call that caused the event. The event data provides further information about the event.

Typically, you process event messages with a system management application tailored to meet the requirements of the enterprise at which it runs.

When the queue managers in a queue sharing group detect the conditions for generating an event message, several queue managers can generate an event message for the shared queue, resulting in several event messages. To ensure that a system can correlate multiple event messages from different queue managers, these event messages have a unique correlation identifier (*CorrelId*) set in the message descriptor (MQMD).

Related reference

[“Activity report MQMD \(message descriptor\)” on page 122](#)

Use this page to view the values contained by the MQMD structure for an activity report

[“Activity report MQEPH \(Embedded PCF header\)” on page 126](#)

Use this page to view the values contained by the MQEPH structure for an activity report

[“Activity report MQCFH \(PCF header\)” on page 127](#)

Use this page to view the PCF values contained by the MQCFH structure for an activity report

[Event message reference](#)

[Event message format](#)

[Event message MQMD \(message descriptor\)](#)

[Event message MQCFH \(PCF header\)](#)

[Event message descriptions](#)

Performance events

Performance events relate to conditions that can affect the performance of applications that use a specified queue. The scope of performance events is the queue. **MQPUT** calls and **MQGET** calls on one queue do not affect the generation of performance events on another queue.

Performance event messages can be generated at any appropriate time, not necessarily waiting until an MQI call for the queue is issued. However, if you use an MQI call on a queue to put or remove a message, any appropriate performance events are generated at that time.

Every performance event message that is generated is placed on the queue, SYSTEM.ADMIN.PERFM.EVENT.

The event data contains a reason code that identifies the cause of the event, a set of performance event statistics, and other data. The types of event data that can be returned in performance event messages are described in the following list:

- [Queue Depth High](#)
- [Queue Depth Low](#)
- [Queue Full](#)
- [Queue Service Interval High](#)
- [Queue Service Interval OK](#)

Examples that illustrate the use of performance events assume that you set queue attributes by using the appropriate IBM MQ commands (MQSC). On z/OS, you can also set queue attributes using the operations and controls panels for queue managers.

Related reference

[“Event types” on page 33](#)

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

Performance event statistics

The performance event data in the event message contains statistics about the event. Use the statistics to analyze the behavior of a specified queue.

The event data in the event message contains information about the event for system management programs. For all performance events, the event data contains the names of the queue manager and the queue associated with the event. The event data also contains statistics related to the event. [Table 7 on page 48](#) summarizes the event statistics that you can use to analyze the behavior of a queue. All the statistics refer to what has happened since the last time the statistics were reset.

<i>Table 7. Performance event statistics</i>	
Parameter	Description
TimeSinceReset	The elapsed time since the statistics were last reset.
HighQDepth	The maximum number of messages on the queue since the statistics were last reset.
MsgEnqCount	The number of messages enqueued (the number of MQPUT calls to the queue), since the statistics were last reset.
MsgDeqCount	The number of messages dequeued (the number of MQGET calls to the queue), since the statistics were last reset.

Performance event statistics are reset when any of the following changes occur:

- A performance event occurs (statistics are reset on all active queue managers).
- A queue manager stops and restarts.
- The PCF command, Reset Queue Statistics, is issued from an application program.
-  On z/OS only, the RESET QSTATS command is issued at the console.

Related concepts

[“Performance events” on page 47](#)

Performance events relate to conditions that can affect the performance of applications that use a specified queue. The scope of performance events is the queue. **MQPUT** calls and **MQGET** calls on one queue do not affect the generation of performance events on another queue.

[“The service timer” on page 50](#)

Queue service interval events use an internal timer, called the *service timer*, which is controlled by the queue manager. The service timer is used only if a queue service interval event is enabled.

[“Rules for queue service interval events” on page 50](#)

Formal rules control when the service timer is set and queue service interval events are generated.

Related tasks

[“Enabling queue service interval events” on page 51](#)

To configure a queue for queue service interval events you set the appropriate queue manager and queue attributes.

Related reference

[Queue Depth High](#)

[Reset Queue Statistics](#)

[RESET QSTATS](#)

Queue service interval events

Queue service interval events indicate whether an operation was performed on a queue within a user-defined time interval called the *service interval*. Depending on your installation, you can use queue service interval events to monitor whether messages are being taken off queues quickly enough.

Queue service interval events are not supported on shared queues.

The following types of queue service interval events can occur, where the term *get operation* refers to an **MQGET** call or an activity that removes a messages from a queue, such as using the **CLEAR QLOCAL** command:

Queue Service Interval OK

Indicates that after one of the following operations:

- An MQPUT call
- A get operation that leaves a non-empty queue

a get operation was performed within a user-defined time period, known as the *service interval*.

Only a get operation can cause the Queue Service Interval OK event message. Queue Service Interval OK events are sometimes described as OK events.

Queue Service Interval High

Indicates that after one of the following operations:

- An MQPUT call
- A get operation that leaves a non-empty queue

a get operation was **not** performed within a user-defined service interval.

Either a get operation or an MQPUT call can cause the Queue Service Interval High event message. Queue Service Interval High events are sometimes described as High events.

To enable both Queue Service Interval OK and Queue Service Interval High events, set the `QServiceIntervalEvent` control attribute to High. Queue Service Interval OK events are automatically enabled when a Queue Service Interval High event is generated. You do not need to enable Queue Service Interval OK events independently.

OK and High events are mutually exclusive, so if one is enabled the other is disabled. However, both events can be simultaneously disabled.

Figure 7 on page 49 shows a graph of queue depth against time. At time P1, an application issues an MQPUT, to put a message on the queue. At time G1, another application issues an MQGET to remove the message from the queue.




Figure 7. Understanding queue service interval events

The possible outcomes of queue service interval events are as follows:

- If the elapsed time between the put and the get is less than or equal to the service interval:

- A *Queue Service Interval OK* event is generated at time G1, if queue service interval events are enabled
- If the elapsed time between the put and get is greater than the service interval:
 - A *Queue Service Interval High* event is generated at time G1, if queue service interval events are enabled.

The algorithm for starting the service timer and generating events is described in [“Rules for queue service interval events”](#) on page 50.

Related reference

[Queue Service Interval OK](#)

[Queue Service Interval High](#)

[QServiceIntervalEvent \(MQLONG\)](#)

[QServiceIntervalEvent \(10-digit signed integer\)](#)

The service timer

Queue service interval events use an internal timer, called the *service timer*, which is controlled by the queue manager. The service timer is used only if a queue service interval event is enabled.

What precisely does the service timer measure?

The service timer measures the elapsed time between an MQPUT call to an empty queue or a get operation, and the next put or get, provided the queue depth is nonzero between these two operations.

When is the service timer active?

The service timer is always active (running), if the queue has messages on it (depth is nonzero) and a queue service interval event is enabled. If the queue becomes empty (queue depth zero), the timer is put into an OFF state, to be restarted on the next put.

When is the service timer reset?

The service timer is always reset after a get operation . It is also reset by an MQPUT call to an empty queue. However, it is not necessarily reset on a queue service interval event.

How is the service timer used?

Following a get operation or an MQPUT call, the queue manager compares the elapsed time as measured by the service timer, with the user-defined service interval. The result of this comparison is that:

- An OK event is generated if there is a get operation and the elapsed time is less than or equal to the service interval, AND this event is enabled.
- A high event is generated if the elapsed time is greater than the service interval, AND this event is enabled.

Can applications read the service timer?

No, the service timer is an internal timer that is not available to applications.

What about the *TimeSinceReset* parameter?

The *TimeSinceReset* parameter is returned as part of the event statistics in the event data. It specifies the time between successive queue service interval events, unless the event statistics are reset.

Rules for queue service interval events

Formal rules control when the service timer is set and queue service interval events are generated.

Rules for the service timer

The service timer is reset to zero and restarted as follows:

- After an MQPUT call to an empty queue.
- After an MQGET call, if the queue is not empty after the MQGET call.

The resetting of the timer does not depend on whether an event has been generated.

At queue manager startup the service timer is set to startup time if the queue depth is greater than zero. If the queue is empty following a get operation, the timer is put into an OFF state.

Queue Service Interval High events

The Queue Service Interval event must be enabled (set to HIGH).

Queue Service Interval High events are automatically enabled when a Queue Service Interval OK event is generated.

If the service time is greater than the service interval, an event is generated on, or before, the next MQPUT or get operation.

Queue Service Interval OK events

Queue Service Interval OK events are automatically enabled when a Queue Service Interval High event is generated.

If the service time (elapsed time) is less than or equal to the service interval, an event is generated on, or before, the next get operation.

Related tasks

[“Enabling queue service interval events” on page 51](#)

To configure a queue for queue service interval events you set the appropriate queue manager and queue attributes.

Enabling queue service interval events

To configure a queue for queue service interval events you set the appropriate queue manager and queue attributes.

About this task

The high and OK events are mutually exclusive; that is, when one is enabled, the other is automatically disabled:

- When a high event is generated on a queue, the queue manager automatically disables high events and enables OK events for that queue.
- When an OK event is generated on a queue, the queue manager automatically disables OK events and enables high events for that queue.

Table 8. Enabling queue service interval events using MQSC	
Queue service interval event	Queue attributes
Queue Service Interval High Queue Service Interval OK No queue service interval events	QSVCI EV (HIGH) QSVCI EV (OK) QSVCI EV (NONE)
Service interval	QSVCI NT (<i>tt</i>) where <i>tt</i> is the service interval time in milliseconds.

Perform the following steps to enable queue service interval events:

Procedure

1. Set the queue manager attribute **PERFMEV** to ENABLED.
Performance events are enabled on the queue manager.

2. Set the control attribute, **QSVCI EV**, for a Queue Service Interval High or OK event on the queue, as required.
3. Set the **QSVCI NT** attribute for the queue to specify the appropriate service interval time.

Example

To enable Queue Service Interval High events with a service interval time of 10 seconds (10 000 milliseconds) use the following MQSC commands:

```
ALTER QMGR PERFMEV(ENABLED)
ALTER QLOCAL('MYQUEUE') QSVCI NT(10000) QSVCI EV(HIGH)
```

Queue service interval events examples

Use the examples in this section to understand the information that you can obtain from queue service interval events.

The three subtopic examples provide progressively more complex illustrations of the use of queue service interval events.

The figures accompanying the examples in each subtopic have the same structure:

- Figure 1 is a graph of queue depth against time, showing individual MQGET calls and MQPUT calls.
- The Commentary section shows a comparison of the time constraints. There are three time periods that you must consider:
 - The user-defined service interval.
 - The time measured by the service timer.
 - The time since event statistics were last reset (TimeSinceReset in the event data).
- The Event statistics summary section shows which events are enabled at any instant and what events are generated.

The examples illustrate the following aspects of queue service interval events:

- How the queue depth varies over time.
- How the elapsed time as measured by the service timer compares with the service interval.
- Which event is enabled.
- Which events are generated.

Remember: Example 1 shows a simple case where the messages are intermittent and each message is removed from the queue before the next one arrives. From the event data, you know that the maximum number of messages on the queue was one. You can, therefore, work out how long each message was on the queue.

However, in the general case, where there is more than one message on the queue and the sequence of MQGET calls and MQPUT calls is not predictable, you cannot use queue service interval events to calculate how long an individual message remains on a queue. The TimeSinceReset parameter, which is returned in the event data, can include a proportion of time when there are no messages on the queue. Therefore any results you derive from these statistics are implicitly averaged to include these times.

Related concepts

[“Queue service interval events” on page 49](#)

Queue service interval events indicate whether an operation was performed on a queue within a user-defined time interval called the *service interval*. Depending on your installation, you can use queue service interval events to monitor whether messages are being taken off queues quickly enough.

[“The service timer” on page 50](#)

Queue service interval events use an internal timer, called the *service timer*, which is controlled by the queue manager. The service timer is used only if a queue service interval event is enabled.

Queue service interval events: example 1

A basic sequence of MQGET calls and MQPUT calls, where the queue depth is always one or zero.




Figure 8. Queue service interval events - example 1

Commentary

- At P1, an application puts a message onto an empty queue. This starts the service timer.
Note that T0 might be queue manager startup time.
- At G1, another application gets the message from the queue. Because the elapsed time between P1 and G1 is greater than the service interval, a Queue Service Interval High event is generated on the MQGET call at G1. When the high event is generated, the queue manager resets the event control attribute so that:
 - The OK event is automatically enabled.
 - The high event is disabled.
 Because the queue is now empty, the service timer is switched to an OFF state.
- At P2, a second message is put onto the queue. This restarts the service timer.

4. At G2, the message is removed from the queue. However, because the elapsed time between P2 and G2 is less than the service interval, a Queue Service Interval OK event is generated on the MQGET call at G2. When the OK event is generated, the queue manager resets the control attribute so that:
 - a. The high event is automatically enabled.
 - b. The OK event is disabled.

Because the queue is empty, the service timer is again switched to an OFF state.

Event statistics summary

Table 9 on page 54 summarizes the event statistics for this example.

Table 9. Event statistics summary for example 1		
Attribute	Event 1	Event 2
Time of event	T(G1)	T(G2)
Type of event	High	OK
TimeSinceReset	T(G1) - T(0)	T(G2) - T(G1)
HighQDepth	1	1
MsgEnqCount	1	1
MsgDeqCount	1	1

The middle part of [Figure 8 on page 53](#) shows the elapsed time as measured by the service timer compared to the service interval for that queue. To see whether a queue service interval event might occur, compare the length of the horizontal line representing the service timer (with arrow) to that of the line representing the service interval. If the service timer line is longer, and the Queue Service Interval High event is enabled, a Queue Service Interval High event occurs on the next get. If the timer line is shorter, and the Queue Service Interval OK event is enabled, a Queue Service Interval OK event occurs on the next get.

Queue service interval events: example 2

A sequence of MQPUT calls and MQGET calls, where the queue depth is not always one or zero.

This example also shows instances of the timer being reset without events being generated, for example, at time P2.




Figure 9. Queue service interval events - example 2

Commentary

In this example, OK events are enabled initially and queue statistics were reset at time T0.

1. At P1, the first put starts the service timer.
2. At P2, the second put does not generate an event because a put cannot cause an OK event.
3. At G1, the service interval has now been exceeded and therefore an OK event is not generated. However, the MQGET call causes the service timer to be reset.
4. At G2, the second get occurs within the service interval and this time an OK event is generated. The queue manager resets the event control attribute so that:
 - a. The high event is automatically enabled.
 - b. The OK event is disabled.

Because the queue is now empty, the service timer is switched to an OFF state.

Event statistics summary

Table 10 on page 56 summarizes the event statistics for this example.

Table 10. Event statistics summary for example 2	
Attribute	Event 2
Time of event	T(G2)
Type of event	OK
TimeSinceReset	T(G2) - T(0)
HighQDepth	2
MsgEnqCount	2
MsgDeqCount	2

Queue service interval events: example 3

A sequence of MQGET calls and MQPUT calls that is more sporadic than the previous examples.




Figure 10. Queue service interval events - example 3

Commentary

1. At time T(0), the queue statistics are reset and Queue Service Interval High events are enabled.
2. At P1, the first put starts the service timer.
3. At P2, the second put increases the queue depth to two. A high event is not generated here because the service interval time has not been exceeded.
4. At P3, the third put causes a high event to be generated. (The timer has exceeded the service interval.) The timer is not reset because the queue depth was not zero before the put. However, OK events are enabled.
5. At G1, the MQGET call does not generate an event because the service interval has been exceeded and OK events are enabled. The MQGET call does, however, reset the service timer.
6. At G2, the MQGET call does not generate an event because the service interval has been exceeded and OK events are enabled. Again, the MQGET call resets the service timer.
7. At G3, the third get empties the queue and the service timer is equal to the service interval. Therefore an OK event is generated. The service timer is reset and high events are enabled. The MQGET call empties the queue, and this puts the timer in the OFF state.

Event statistics summary

Table 11 on page 57 summarizes the event statistics for this example.

Table 11. Event statistics summary for example 3		
Attribute	Event 1	Event 2
Time of event	T(P3)	T(G3)
Type of event	High	OK
TimeSinceReset	T(P3) - T(0)	T(G3) - T(P3)
HighQDepth	3	3
MsgEnqCount	3	0
MsgDeqCount	0	3

Queue depth events

Queue depth events are related to the queue depth, that is, the number of messages on the queue.

In IBM MQ applications, queues must not become full. If they do, applications can no longer put messages on the queue that they specify. Although the message is not lost if this occurs, a full queue can cause considerable inconvenience. The number of messages can build up on a queue if the messages are being put onto the queue faster than the applications that process them can take them off.

The solution to this problem depends on the particular circumstances, but might involve:

- Diverting some messages to another queue.
- Starting new applications to take more messages off the queue.
- Stopping nonessential message traffic.
- Increasing the queue depth to overcome a transient maximum.

Advance warning that problems might be on their way makes it easier to take preventive action. For this purpose, IBM MQ provides the following queue depth events:

Queue Depth High events

Indicate that the queue depth has increased to a predefined threshold called the Queue Depth High limit.

Queue Depth Low events

Indicate that the queue depth has decreased to a predefined threshold called the Queue Depth Low limit.

Queue Full events

Indicate that the queue has reached its maximum depth, that is, the queue is full.

A Queue Full Event is generated when an application attempts to put a message on a queue that has reached its maximum depth. Queue Depth High events give advance warning that a queue is filling up. This means that having received this event, the system administrator needs to take some preventive action. You can configure the queue manager such that, if the preventive action is successful and the queue depth drops to a safer level, the queue manager generates a Queue Depth Low event.

The first queue depth event example illustrates the effect of presumed action preventing the queue becoming full.

Related concepts

[“Queue depth events examples” on page 61](#)

Use these examples to understand the information that you can obtain from queue depth events

Related reference

[Queue Full](#)

[Queue Depth High](#)

[Queue Depth Low](#)

Enabling queue depth events

To configure a queue for any of the queue depth events you set the appropriate queue manager and queue attributes.

About this task

By default, all queue depth events are disabled. When enabled, queue depth events are generated as follows:

- A Queue Depth High event is generated when a message is put on the queue, causing the queue depth to be greater than or equal to the value set for **QDepthHighLimit**.
 - A Queue Depth High event is automatically enabled by a Queue Depth Low event on the same queue.
 - A Queue Depth High event automatically enables both a Queue Depth Low and a Queue Full event on the same queue.
-  A Queue Depth Low event is generated when a message is removed from a queue by a GET operation, causing the queue depth to be less than or equal to the value set for **QDepthLowLimit**.
 -  A Queue Depth Low event is generated when a message is removed from a queue by a GET operation, or would have been removed but has since expired, causing the queue depth to be less than or equal to the value set for **QDepthLowLimit**.
 - A Queue Depth Low event is automatically enabled by a Queue Depth High event or a Queue Full event on the same queue.
 - A Queue Depth Low event automatically enables both a Queue Depth High and a Queue Full event on the same queue.
- A Queue Full event is generated when an application is unable to put a message onto a queue because the queue is full.
 - A Queue Full event is automatically enabled by a Queue Depth High or a Queue Depth Low event on the same queue.
 - A Queue Full event automatically enables a Queue Depth Low event on the same queue.

Perform the following steps to configure a queue for any of the queue depth events:

Procedure

1. Enable performance events on the queue manager, using the queue manager attribute **PERFMEV**.
The events go to the `SYSTEM.ADMIN.PERFM.EVENT` queue.
2. Set one of the following attributes to enable the event on the required queue:
 - *QDepthHighEvent* (**QDPHIEV** in MQSC)
 - *QDepthLowEvent* (**QDPLOEV** in MQSC)
 - *QDepthMaxEvent* (**QDPMAXEV** in MQSC)
3. Optional: To set the limits, assign the following attributes, as a percentage of the maximum queue depth:
 - *QDepthHighLimit* (**QDEPTHHI** in MQSC)
 - *QDepthLowLimit* (**QDEPTHLO** in MQSC)

Restriction: **QDEPTHHI** must not be less than **QDEPTHLO**.

If **QDEPTHHI** equals **QDEPTHLO** an event message is generated every time the queue depth passes the value in either direction, because the high threshold is enabled when the queue depth is below the value and the low threshold is enabled when the depth is above the value.

Results

Note:

Multi A Queue Depth Low event is not generated when expired messages are removed from a queue by a GET operation causing the queue depth to be less than, or equal to, the value set for **QDepthLowLimit**. IBM MQ generates the queue depth low event message only during a successful GET operation. Therefore, when the expired messages are removed from the queue, no queue depth low event message is generated. Additionally, after the removal of these expired messages from the queue, *QDepthHighEvent* and *QDepthLowEvent* are not reset.

z/OS IBM MQ generates the queue depth low event message during either a successful destructive GET operation, or a destructive GET operation that would have succeeded had a matching message not expired. Otherwise, when expired messages are removed from a queue during general background processing, no queue depth low event message is generated. Additionally, after the removal of expired messages from a queue during general background processing, *QDepthHighEvent* and *QDepthLowEvent* are not reset. For more information on expired message processing, see [Tuning your queue manager on IBM MQ for z/OS](#).

Example

To enable Queue Depth High events on the queue MYQUEUE with a limit set at 80%, use the following MQSC commands:

```
ALTER QMGR PERFMEV(ENABLED)
ALTER QLOCAL('MYQUEUE') QDEPTHHI(80) QDPHIEV(ENABLED)
```

To enable Queue Depth Low events on the queue MYQUEUE with a limit set at 20%, use the following MQSC commands:

```
ALTER QMGR PERFMEV(ENABLED)
ALTER QLOCAL('MYQUEUE') QDEPTHLO(20) QDPLOEV(ENABLED)
```

To enable Queue Full events on the queue MYQUEUE, use the following MQSC commands:

```
ALTER QMGR PERFMEV(ENABLED)
ALTER QLOCAL('MYQUEUE') QDPMAXEV(ENABLED)
```

Shared queues and queue depth events on z/OS

On IBM MQ for z/OS, event monitoring is more straightforward for an application that uses shared queues if all the queue managers in the queue sharing group have the same setting for the **PERFMEV** attribute.

When a queue depth event occurs on a shared queue, and the queue manager attribute **PERFMEV** is set to **ENABLED**, the queue managers in the queue sharing group produce an event message. If **PERFMEV** is set to **DISABLED** on some of the queue managers, event messages are not produced by those queue managers, making event monitoring from an application more difficult. For more straightforward monitoring, give each queue manager the same setting for the **PERFMEV** attribute.

This event message that each queue manager generates represents its individual usage of the shared queue. If a queue manager performs no activity on the shared queue, various values in the event message are null or zero. You can use null event messages as follows:

- Ensure that each active queue manager in a queue sharing group generates one event message
- Highlight cases of no activity on a shared queue for the queue manager that produced the event message

Coordinating queue manager

When a queue manager issues a queue depth event, it updates the shared queue object definition to toggle the active performance event attributes. For example, depending on the definition of the queue attributes, a Queue Depth High event enables a Queue Depth Low and a Queue Full event. After updating the shared queue object successfully, the queue manager that detected the performance event initially becomes the *coordinating queue manager*.

If enabled for performance events, the coordinating queue manager performs the following actions:

1. Issues an event message that captures all shared queue performance data it has gathered since the last time an event message was created, or since the queue statistics were last reset. The message descriptor (MQMD) of this message contains a unique correlation identifier (*CorrelId*) created by the coordinating queue manager.
2. Broadcasts to all other *active* queue managers in the same queue sharing group to request the production of an event message for the shared queue. The broadcast contains the correlation identifier created by the coordinating queue manager for the set of event messages.

Having received a request from the coordinating queue manager, if there is an active queue manager in the queue sharing group that is enabled for performance events, that active queue manager issues an event message for the shared queue. The event message that is issued contains information about all the operations performed by the receiving (active) queue manager since the last time an event message was created, or since the statistics were last reset. The message descriptor (MQMD) of this event message contains the unique correlation identifier (*CorrelId*) specified by the coordinating queue manager.

When performance events occur on a shared queue, n event messages are produced, where n is a number from 1 to the number of active queue managers in the queue sharing group. Each event message contains data that relates to the shared queue activity for the queue manager that generated the event message.

Differences between shared and nonshared queues

Enabling queue depth events on shared queues differs from enabling them on nonshared queues. A key difference is that events are enabled for shared queues even if **PERFMEV** is **DISABLED** on the queue manager. This is not the case for nonshared queues.

Consider the following example, which illustrates this difference:

- QM1 is a queue manager with *PerformanceEvent* (**PERFMEV** in MQSC) set to **DISABLED**.
- SQ1 is a shared queue with **QSGDISP** set to (SHARED) QLOCAL in MQSC.
- LQ1 is a nonshared queue with **QSGDISP** set to (QMGR) QLOCAL in MQSC.

Both queues have the following attributes set on their definitions:

- **QDPHIEV (ENABLED)**
- **QDPLOEV (DISABLED)**
- **QDPMAXEV (DISABLED)**

If messages are placed on both queues so that the depth meets or exceeds the **QDEPTHHI** threshold, the **QDPHIEV** value on SQ1 switches to **DISABLED**. Also, **QDPLOEV** and **QDPMAXEV** are switched to **ENABLED**. SQ1's attributes are automatically switched for each performance event at the time the event criteria are met.

In contrast the attributes for LQ1 remain unchanged until **PERFMEV** on the queue manager is **ENABLED**. This means that if, for example, the queue manager's **PERFMEV** attribute is **ENABLED**, **DISABLED** and then set to **ENABLED** again, the performance event settings on shared queues might not be consistent with those of nonshared queues, even though they might have initially been the same.

Queue depth events examples

Use these examples to understand the information that you can obtain from queue depth events

The first example provides a basic illustration of queue depth events. The second example is more extensive, but the principles are the same as for the first example. Both examples use the same queue definition, as follows:

The queue, MYQUEUE1, has a maximum depth of 1000 messages. The high queue depth limit is 80% and the low queue depth limit is 20%. Initially, Queue Depth High events are enabled, while the other queue depth events are disabled.

The IBM MQ commands (MQSC) to configure this queue are:

```
ALTER QMGR PERFMEV(ENABLED)

DEFINE QLOCAL('MYQUEUE1') MAXDEPTH(1000) QDPMAXEV(DISABLED) QDEPTHHI(80)
QDPHIEV(ENABLED) QDEPTHLO(20) QDPLOEV(DISABLED)
```

Related concepts

[“Queue depth events” on page 57](#)

Queue depth events are related to the queue depth, that is, the number of messages on the queue.

Related tasks

[“Enabling queue depth events” on page 58](#)

To configure a queue for any of the queue depth events you set the appropriate queue manager and queue attributes.

Related reference

[The MQSC commands](#)

Queue depth events: example 1

A basic sequence of queue depth events.

[Figure 11 on page 62](#) shows the variation of queue depth over time.




Figure 11. Queue depth events (1)

Commentary

1. At T(1), the queue depth is increasing (more MQPUT calls than MQGET calls) and crosses the Queue Depth Low limit. No event is generated at this time.
2. The queue depth continues to increase until T(2), when the depth high limit (80%) is reached and a Queue Depth High event is generated.
This enables both Queue Full and Queue Depth Low events.
3. The (presumed) preventive actions instigated by the event prevent the queue from becoming full. By time T(3), the Queue Depth High limit has been reached again, this time from above. No event is generated at this time.
4. The queue depth continues to fall until T(4), when it reaches the depth low limit (20%) and a Queue Depth Low event is generated.

This enables both Queue Full and Queue Depth High events.

Event statistics summary

Table 12 on page 63 summarizes the queue event statistics and Table 13 on page 63 summarizes which events are enabled.

<i>Table 12. Event statistics summary for queue depth events (example 1)</i>		
Event statistics	Event 2	Event 4
Time of event	T(2)	T(4)
Type of event	Queue Depth High	Queue Depth Low
TimeSinceReset	T(2) - T(0)	T(4) - T(2)
HighQDepth (Maximum queue depth since reset)	800	900
MsgEnqCount	1157	1220
MsgDeqCount	357	1820

<i>Table 13. Summary showing which events are enabled</i>			
Time period	Queue Depth High event	Queue Depth Low event	Queue Full event
Before T(1)	ENABLED	-	-
T(1) to T(2)	ENABLED	-	-
T(2) to T(3)	-	ENABLED	ENABLED
T(3) to T(4)	-	ENABLED	ENABLED
After T(4)	ENABLED	-	ENABLED

Queue depth events: example 2

A more extensive sequence of queue depth events.

Figure 12 on page 64 shows the variation of queue depth over time.




Figure 12. Queue depth events (2)

Commentary

1. No Queue Depth Low event is generated at the following times:
 - T(1) (Queue depth increasing, and not enabled)
 - T(2) (Not enabled)
 - T(3) (Queue depth increasing, and not enabled)
2. At T(4) a Queue Depth High event occurs. This enables both Queue Full and Queue Depth Low events.
3. At T(9) a Queue Full event occurs **after** the first message that cannot be put on the queue because the queue is full.
4. At T(12) a Queue Depth Low event occurs.

Event statistics summary

Table 14 on page 65 summarizes the queue event statistics and [Table 15 on page 65](#) summarizes which events are enabled at different times for this example.

Table 14. Event statistics summary for queue depth events (example 2)					
Event statistics	Event 4	Event 6	Event 8	Event 9	Event 12
Time of event	T(4)	T(6)	T(8)	T(9)	T(12)
Type of event	Queue Depth High	Queue Depth Low	Queue Depth High	Queue Full	Queue Depth Low
TimeSinceReset	T(4) - T(0)	T(6) - T(4)	T(8) - T(6)	T(9) - T(8)	T(12) - T(9)
HighQDepth	800	855	800	1000	1000
MsgEnqCount	1645	311	1377	324	221
MsgDeqCount	845	911	777	124	1021

Table 15. Summary showing which events are enabled			
Time period	Queue Depth High event	Queue Depth Low event	Queue Full event
T(0) to T(4)	ENABLED	-	-
T(4) to T(6)	-	ENABLED	ENABLED
T(6) to T(8)	ENABLED	-	ENABLED
T(8) to T(9)	-	ENABLED	ENABLED
T(9) to T(12)	-	ENABLED	-
After T(12)	ENABLED	-	ENABLED

Note: Events are out of syncpoint. Therefore you could have an empty queue, then fill it up causing an event, then roll back all of the messages under the control of a syncpoint manager. However, event enabling has been automatically set, so that the next time the queue fills up, no event is generated.

Configuration events

Configuration events are notifications that are generated when an object is created, changed, or deleted, and can also be generated by explicit requests.

See [“Controlling configuration, command, and logger events” on page 45](#) for information on how you turn on events.

Configuration events notify you about changes to the attributes of an object. There are four types of configuration events:

- Create object events
- Change object events
- Delete object events
- Refresh object events

The event data contains the following information:

Origin information

comprises the queue manager from where the change was made, the ID of the user that made the change, and how the change came about, for example by a console command.

Context information

a replica of the context information in the message data from the command message.

Context information is included in the event data only when the command was entered as a message on the SYSTEM.COMMAND.INPUT queue.

Object identity

comprises the name, type and disposition of the object.

Object attributes

comprises the values of all the attributes in the object.

The event is a message in PCF format. For more information, see:

- [Change object](#)
- [Create object](#)
- [Delete object](#)
- [Refresh object](#)

In the case of change object events, a before message and an after message are produced. They have the following fields in common, so that you can identify matching messages:

- The same correlation ID
- In the PCF header:
 - The before message header has `MsgSeqNumber': 1, 'Control': 'NOT_LAST'`
 - The after message header has `MsgSeqNumber': 2, 'Control': 'LAST'`

Every configuration event message that is generated is placed on the queue SYSTEM.ADMIN.CONFIG.EVENT.

Related concepts

[“Configuration events” on page 40](#)

Configuration events are generated when a configuration event is requested explicitly, or automatically when an object is created, modified, or deleted.

Related reference

[Create object](#)

[Change object](#)

[Delete object](#)

[Refresh object](#)

[“Event types” on page 33](#)

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

Configuration event generation

Use this page to view the commands that cause configuration events to be generated and to understand the circumstances in which configuration events are not generated

A configuration event message is put to the configuration event queue when the CONFIGEV queue manager attribute is ENABLED and

- any of the following commands, or their PCF equivalent, are issued:
 - DELETE AUTHINFO
 - DELETE CFSTRUCT
 - DELETE CHANNEL
 - DELETE NAMELIST
 - DELETE PROCESS
 - DELETE QMODEL/QALIAS/QREMOTE

- DELETE STGCLASS
- DELETE TOPIC
- REFRESH QMGR
- any of the following commands, or their PCF equivalent, are issued even if there is no change to the object:
 - DEFINE/ALTER AUTHINFO
 - DEFINE/ALTER CFSTRUCT
 - DEFINE/ALTER CHANNEL
 - DEFINE/ALTER NAMELIST
 - DEFINE/ALTER PROCESS
 - DEFINE/ALTER QMODEL/QALIAS/QREMOTE
 - DEFINE/ALTER STGCLASS
 - DEFINE/ALTER TOPIC
 - DEFINE MAXSMSGS
 - SET CHLAUTH
 - ALTER QMGR, unless the CONFIGEV attribute is DISABLED and is not changed to ENABLED
- any of the following commands, or their PCF equivalent, are issued for a local queue that is not temporary dynamic, even if there is no change to the queue.
 - DELETE QLOCAL
 - DEFINE/ALTER QLOCAL
- an MQSET call is issued, other than for a temporary dynamic queue, even if there is no change to the object.

When configuration events are not generated

Configuration events messages are not generated in the following circumstances:

- When a command or an MQSET call fails
- When a queue manager encounters an error trying to put a configuration event on the event queue, in which case the command or MQSET call completes, but no event message is generated
- For a temporary dynamic queue
- When internal changes are made to the TRIGGER queue attribute
- For the configuration event queue SYSTEM.ADMIN.CONFIG.EVENT, except by the REFRESH QMGR command
- For REFRESH/RESET CLUSTER and RESUME/SUSPEND QMGR commands that cause clustering changes
- When Creating or deleting a queue manager

Related concepts

[Introduction to Programmable Command Formats](#)

“Configuration events” on page 65

Configuration events are notifications that are generated when an object is created, changed, or deleted, and can also be generated by explicit requests.

Related reference

[The MQSC commands](#)

[MQSET - Set object attributes](#)

[MQSET - Set object attributes](#)

Configuration event usage

Use this page to view how you can use configuration events to obtain information about your system, and to understand the factors, such as CMDSCOPE, that can affect your use of configuration events.

You can use configuration events for the following purposes:

1. To produce and maintain a central configuration repository, from which reports can be produced and information about the structure of the system can be generated.
2. To generate an audit trail. For example, if an object is changed unexpectedly, information regarding who made the alteration and when it was done can be stored.

This can be particularly useful when command events are also enabled. If an MQSC or PCF command causes a configuration event and a command event to be generated, both event messages will share the same correlation identifier in their message descriptor.

For an MQSET call or any of the following commands:

- DEFINE object
- ALTER object
- DELETE object

if the queue manager attribute CONFIGEV is enabled, but the configuration event message cannot be put on the configuration event queue, for example the event queue has not been defined, the command or MQSET call is executed regardless.

Effects of CMDSCOPE

For commands where CMDSCOPE is used, the configuration event message or messages will be generated on the queue manager or queue managers where the command is executed, not where the command is entered. However, all the origin and context information in the event data will relate to the original command as entered, even where the command using CMDSCOPE is one that has been generated by the source queue manager.

Where a queue sharing group includes queue managers that are not at the current version, events will be generated for any command that is executed by means of CMDSCOPE on a queue manager that is at the current version, but not on those that are at a previous version. This happens even if the queue manager where the command is entered is at the previous version, although in such a case no context information is included in the event data.

Related concepts

[Introduction to Programmable Command Formats](#)

[“Configuration events” on page 65](#)

Configuration events are notifications that are generated when an object is created, changed, or deleted, and can also be generated by explicit requests.

Related reference

[MQSET - Set object attributes](#)

[MQSET - Set object attributes](#)

Refresh Object configuration event

The Refresh Object configuration event is different from the other configuration events, because it occurs only when explicitly requested.

The create, change, and delete events are generated by an MQSET call or by a command to change an object but the refresh object event occurs only when explicitly requested by the MQSC command, REFRESH QMGR, or its PCF equivalent.

The REFRESH QMGR command is different from all the other commands that generate configuration events. All the other commands apply to a particular object and generate a single configuration event for that object. The REFRESH QMGR command can produce many configuration event messages potentially

representing every object definition stored by a queue manager. One event message is generated for each object that is selected.

The REFRESH QMGR command uses a combination of three selection criteria to filter the number of objects involved:

- Object Name
- Object Type
- Refresh Interval

If you specify none of the selection criteria on the REFRESH QMGR command, the default values are used for each selection criteria and a refresh configuration event message is generated for every object definition stored by the queue manager. This might cause unacceptable processing times and event message generation. Consider specifying some selection criteria.

The REFRESH QMGR command that generates the refresh events can be used in the following situations:

- When configuration data is wanted about all or some of the objects in a system regardless of whether the objects have been recently manipulated, for example, when configuration events are first enabled.

Consider using several commands, each with a different selection of objects, but such that all are included.

- If there has been an error in the SYSTEM.ADMIN.CONFIG.EVENT queue. In this circumstance, no configuration event messages are generated for Create, Change, or Delete events. When the error on the queue has been corrected, the Refresh Queue Manager command can be used to request the generation of event messages, which were lost while there was an error in the queue. In this situation consider setting the refresh interval to the time for which the queue was unavailable.

Related concepts

[“Configuration events” on page 65](#)

Configuration events are notifications that are generated when an object is created, changed, or deleted, and can also be generated by explicit requests.

Related reference

[REFRESH QMGR](#)

[Refresh Queue Manager](#)

Command events

Command events are notifications that an MQSC, or PCF command has run successfully.

The event data contains the following information:

Origin information

comprises the queue manager from where the command was issued, the ID of the user that issued the command, and how the command was issued, for example by a console command.

Context information

a replica of the context information in the message data from the command message. If a command is not entered using a message, context information is omitted.

Context information is included in the event data only when the command was entered as a message on the SYSTEM.COMMAND.INPUT queue.

Command information

the type of command that was issued.

Command data

- for PCF commands, a replica of the command data
- for MQSC commands, the command text

The command data format does not necessarily match the format of the original command. For example, on [Multiplatforms](#) the command data format is always in PCF format, even if the original request was an MQSC command.

Every command event message that is generated is placed on the command event queue, SYSTEM.ADMIN.COMMAND.EVENT.

Related reference

[Command](#)

[“Event types” on page 33](#)

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

Command event generation

Use this page to view the situations that cause command events to be generated and to understand the circumstances in which command events are not generated

A command event message is generated in the following situations:

- When the CMDEV queue manager attribute is specified as ENABLED and an MQSC or PCF command runs successfully.
- When the CMDEV queue manager attribute is specified as NODISPLAY and any command runs successfully, with the exception of DISPLAY commands (MQSC), and Inquire commands (PCF).
- When you run the MQSC command, ALTER QMGR, or the PCF command, Change Queue Manager, and the CMDEV queue manager attribute meets either of the following conditions:
 - CMDEV is not specified as DISABLED after the change
 - CMDEV was not specified as DISABLED before the change

If a command runs against the command event queue, SYSTEM.ADMIN.COMMAND.EVENT, a command event is generated if the queue still exists and it is not put-inhibited.

When command events are not generated

A command event message is not generated in the following circumstances:

- When a command fails
- When a queue manager encounters an error trying to put a command event on the event queue, in which case the command runs regardless, but no event message is generated
- For the MQSC command REFRESH QMGR TYPE (EARLY)
- For the MQSC command START QMGR MQSC
- For the MQSC command SUSPEND QMGR, if the parameter LOG is specified
- For the MQSC command RESUME QMGR, if the parameter LOG is specified

Related concepts

[“Command events” on page 69](#)

Command events are notifications that an MQSC, or PCF command has run successfully.

Related reference

[REFRESH QMGR](#)

[START QMGR](#)

[SUSPEND QMGR](#)

[RESUME QMGR](#)

[SUSPEND QMGR, RESUME QMGR and clusters](#)

Command event usage

Use this page to view how you can use command events to generate an audit trail of the commands that have run

For example, if an object is changed unexpectedly, information regarding who made the alteration and when it was done can be stored. This can be particularly useful when configuration events are also enabled. If an MQSC or PCF command causes a command event and a configuration event to be generated, both event messages will share the same correlation identifier in their message descriptor.

If a command event message is generated, but cannot be put on the command event queue, for example if the command event queue has not been defined, the command for which the command event was generated still runs regardless.

Effects of CMDSCOPE

For commands where CMDSCOPE is used, the command event message or messages will be generated on the queue manager or queue managers where the command runs, not where the command is entered. However, all the origin and context information in the event data will relate to the original command as entered, even where the command using CMDSCOPE is one that has been generated by the source queue manager.

Related concepts

[“Command events” on page 69](#)

Command events are notifications that an MQSC, or PCF command has run successfully.

[“Command event generation” on page 70](#)

Use this page to view the situations that cause command events to be generated and to understand the circumstances in which command events are not generated

Related reference

[The MQSC commands](#)

[PCF commands and responses in groups](#)

Multi Logger events

Logger events are notifications that a queue manager has started writing to a new log extent  or, on IBM i, a journal receiver.  Logger event messages are not available with IBM MQ for z/OS.

The event data contains the following information:

- The name of the current log extent.
- The name of the earliest log extent needed for restart recovery.
- The name of the earliest log extent needed for media recovery.
- The directory in which the log extents are located.
- The name of the earliest log extent for which archive notification is required.

A logger event is generated for archive log management, that is, when **ARCHLOG** changes, and the **ARCHLOG** value is included the logger event message.

See [DISPLAY QMSTATUS](#) for more information on all these parameters.

You can use the **CURRLOG** and **ARCHLOG** values to determine what should be archived. Whenever **CURRLOG** changes, files with numbers less than **CURRLOG** can be sent for archive and, once archiving has been completed for the extent, you should call [SET LOG](#) to notify the queue manager.

ARCHLOG is the oldest extent that needs archiving. When you have archived that extent and called [SET LOG](#) to notify the queue manager that it has been archived, the queue manager moves **ARCHLOG** to the next extent in sequence. The queue manager schedules extents that are older than **ARCHLOG**, to be deleted or reused.

Once **SET LOG** has been called for this extent a new event is issued, as **ARCHLOG** will have changed, and you need to know what needs to be archived after the new event.

If for some reason your archiving process fails, and a big build up of notifications happens, your administrator can manually issue the command **RESET QMGR TYPE(ARCHLOG)**. This notifies the queue manager that it can reuse or delete all extents older than, and including, the extent you specified.

Every logger event message that is generated is placed on the logger event queue, **SYSTEM.ADMIN.LOGGER.EVENT**.

Related reference

[Logger](#)

[“Event types” on page 33](#)

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

Logger event generation

Use this page to view the situations that cause logger events to be generated and to understand the circumstances in which logger events are not generated

A logger event message is generated in the following situations:

- When the LOGGEREV queue manager attribute is specified as ENABLED and the queue manager starts writing to a new log extent or, on IBM i, a journal receiver.
- When the LOGGEREV queue manager attribute is specified as ENABLED and the queue manager starts.
- When the LOGGEREV queue manager attribute is changed from DISABLED to ENABLED.
- When the LOGGEREV queue manager attribute is notified that the name of the oldest log extent, for which the queue manager is waiting for archive notification (ARCHLOG) has been archived.

Tip: You can use the **RESET QMGR MQSC** command to request a queue manager to start writing to a new log extent.

When logger events are not generated

A logger event message is not generated in the following circumstances:

- When a queue manager is configured to use circular logging.

In this case, the LOGGEREV queue manager attribute is set as DISABLED and cannot be altered.


- When a queue manager encounters an error trying to put a logger event on the event queue, in which case the action that caused the event completes, but no event message is generated.

Related concepts

[“Logger events” on page 71](#)

Logger events are notifications that a queue manager has started writing to a new log extent

 **IBM i**

or, on IBM i, a journal receiver.  **z/OS** Logger event messages are not available with IBM MQ for z/OS.

Related reference

[LoggerEvent \(MQLONG\)](#)

[LoggerEvent \(10-digit signed integer\)](#)

[RESET QMGR](#)

Multi **Logger event usage**

Use this page to view how you can use logger events to determine the log extents that are no longer required for queue manager restart, or media recovery.

You can archive superfluous log extents to a medium such as tape for disaster recovery before removing them from the active log directory. Regular removal of superfluous log extents keeps disk space usage to a minimum.

If the LOGGEREV queue manager attribute is enabled, but a logger event message cannot be put on the logger event queue, for example because the event queue has not been defined, the action that caused the event continues regardless.

Related concepts

[“Logger events” on page 71](#)

Logger events are notifications that a queue manager has started writing to a new log extent

IBM i

or, on IBM i, a journal receiver. **z/OS** Logger event messages are not available with IBM MQ for z/OS.

Related reference

[LoggerEvent \(MQLONG\)](#)

[LoggerEvent \(10-digit signed integer\)](#)

[“Logger event generation” on page 72](#)

Use this page to view the situations that cause logger events to be generated and to understand the circumstances in which logger events are not generated

Multi **Sample C program to monitor the logger event queue**

Use this page to view a sample C program that monitors the logger event queue for new event messages, reads those messages, and puts the contents of the message to stdout.

```

/*****
/*
/* Program name: AMQSLOG0.C
/*
/* Description: Sample C program to monitor the logger event queue and
/* display formatted message content to stdout when a logger
/* event occurs
/*
/* <copyright
/* notice="lm-source-program"
/* pids="5724-H72,"
/* years="2005, 2025"
/* crc="186943832" >
/* Licensed Materials - Property of IBM
/*
/* 5724-H72,
/*
/* (C) Copyright IBM Corp. 2005, 2026. All Rights Reserved.
/*
/* US Government Users Restricted Rights - Use, duplication or
/* disclosure restricted by GSA ADP Schedule Contract with
/* IBM Corp.
/* </copyright>
*****/
/*
/* Function: AMQSLOG is a sample program which monitors the logger event
/* queue for new event messages, reads those messages, and displays the
/* formatted contents of the message to stdout.
/*
*****/
/*
/* AMQSLOG has 1 parameter - the queue manager name (optional, if not
/* specified then the default queue manager is implied)
/*
*****/

/*****
/* Includes
*****/
#include <stdlib.h>
#include <stdio.h>
```

```

#include <string.h>

#include <cmqc.h>          /* MQI constants*/
#include <cmqcfc.h>        /* PCF constants*/

/*****
/* Constants
*****/

#define MAX_MESSAGE_LENGTH 8000

typedef struct _ParmTableEntry
{
    MQLONG  ConstVal;
    PMQCHAR Desc;
} ParmTableEntry;

ParmTableEntry ParmTable[] =
{
    {0, "",},
    {MQCA_Q_MGR_NAME, "Queue Manager Name"},
    {MQCMD_LOGGER_EVENT, "Logger Event Command"},
    {MQRC_LOGGER_STATUS, "Logger Status"},
    {MQCACF_ARCHIVE_LOG_EXTENT_NAME, "Archive Log Extent"},
    {MQCACF_CURRENT_LOG_EXTENT_NAME, "Current Log Extent"},
    {MQCACF_RESTART_LOG_EXTENT_NAME, "Restart Log Extent"},
    {MQCACF_MEDIA_LOG_EXTENT_NAME, "Media Log Extent"},
    {MQCACF_LOG_PATH, "Log Path"}
};

#if defined(MQ_64_BIT)
#define Int32
#else
#define Int32 "l"
#endif

/*****
/* Function prototypes
*****/

static void ProcessPCF(MQHCONN hConn,
                      MQHOBJ hEventQueue,
                      PMQCHAR pBuffer);

static PMQCHAR ParmToString(MQLONG Parameter);

/*****
/* Function: main
*****/

int main(int argc, char * argv[])
{
    MQLONG  CompCode;
    MQLONG  Reason;
    MQHCONN hConn = MQHC_UNUSABLE_HCONN;
    MQOD    ObjDesc = { MQOD_DEFAULT };
    MQCHAR  QMName[MQ_Q_MGR_NAME_LENGTH+1] = "";
    MQCHAR  LogEvQ[MQ_Q_NAME_LENGTH] = "SYSTEM.ADMIN.LOGGER.EVENT";
    MQHOBJ  hEventQueue = MQHO_UNUSABLE_HOBJ;
    PMQCHAR pBuffer = NULL;

    printf("\n/*****/\n");
    printf("/* Sample Logger Event Monitor start */\n");
    printf("/*****/\n");

    /*****
    /* Parse any command line options
    *****/
    if (argc > 1)
    {
        strncpy(QMName, argv[1], (size_t)MQ_Q_MGR_NAME_LENGTH);
    }

    pBuffer = (PMQCHAR)malloc(MAX_MESSAGE_LENGTH);
    if (pBuffer == NULL)
    {
        printf("Can't allocate %d bytes\n", MAX_MESSAGE_LENGTH);
        goto MOD_EXIT;
    }

    /*****
    /* Connect to the specified (or default) queue manager
    *****/

```

```

/*****
MQCONN( QMName,
        &hConn,
        &CompCode,
        &Reason);

if (Reason != MQRC_NONE)
{
    printf("MQCONN ended with reason code %" Int32 "d\n", Reason);
    goto MOD_EXIT;
}

/*****
/* Open the logger event queue for input */
/*****
strncpy(ObjDesc.ObjectQMgrName, QMName, MQ_Q_MGR_NAME_LENGTH);
strncpy(ObjDesc.ObjectName, LogEvQ, MQ_Q_NAME_LENGTH);

MQOPEN( hConn,
        &ObjDesc,
        MQOO_INPUT_EXCLUSIVE,
        &hEventQueue,
        &CompCode,
        &Reason );

if (Reason != MQRC_NONE)
{
    printf("MQOPEN failed for queue manager %.48s Queue %.48s Reason: %" Int32 "d\n",
          ObjDesc.ObjectQMgrName,
          ObjDesc.ObjectName,
          Reason);
    goto MOD_EXIT;
}
else
{
    /*****
    /* Start processing event messages */
    /*****
    ProcessPCF(hConn, hEventQueue, pBuffer);
}

MOD_EXIT:
if (pBuffer != NULL)
{
    free(pBuffer);
}

/*****
/* Close the logger event queue */
/*****
if (hEventQueue != MQHO_UNUSABLE_HOBJ)
{
    MQCLOSE(hConn, &hEventQueue, MQCO_NONE, &CompCode, &Reason);
}

/*****
/* Disconnect */
/*****
if (hConn != MQHC_UNUSABLE_HCONN)
{
    MQDISC(&hConn, &CompCode, &Reason);
}

return 0;
}

/*****
/* Function: ProcessPCF */
/*****
/*
/* Input Parameters:  Handle to queue manager connection */
/*                   Handle to the opened logger event queue object */
/*                   Pointer to a memory buffer to store the incoming PCF */
/*                   message */
/*
/* Output Parameters: None */
/*
/* Logic: Wait for messages to appear on the logger event queue and display */
/*        their formatted contents. */
/*
/*****

```

```

static void ProcessPCF(MQHCONN hConn,
                      MQHOBJ hEventQueue,
                      PMQCHAR pBuffer)
{
    MQCFH * pCfh;
    MQCFST * pCfst;
    MQGMO Gmo = { MQGMO_DEFAULT };
    MQMD Mqmd = { MQMD_DEFAULT };
    PMQCHAR pPCFCmd;
    MQLONG CompCode = MQCC_OK;
    MQLONG Reason = MQRC_NONE;
    MQLONG MsgLen;
    PMQCHAR Parm = NULL;

    Gmo.Options = MQGMO_WAIT +
                  MQGMO_CONVERT +
                  MQGMO_FAIL_IF_QUIESCING;
    Gmo.WaitInterval = MQWI_UNLIMITED; /* Set timeout value */

    /******
    /* Process response Queue */
    /******
    while (Reason == MQRC_NONE)
    {
        memcpy(&Mqmd.MsgId, MQMI_NONE, sizeof(Mqmd.MsgId));
        memset(&Mqmd.CorrelId, 0, sizeof(Mqmd.CorrelId));

        MQGET( hConn,
               hEventQueue,
               &Mqmd,
               &Gmo,
               MAX_MESSAGE_LENGTH,
               pBuffer,
               &MsgLen,
               &CompCode,
               &Reason );

        if (Reason != MQRC_NONE)
        {
            switch(Reason)
            {
                case MQRC_NO_MSG_AVAILABLE:
                    printf("Timed out");
                    break;

                default:
                    printf("MQGET ended with reason code %" Int32 "d\n", Reason);
                    break;
            }
            goto MOD_EXIT;
        }

        /******
        /* Only expect PCF event messages on this queue */
        /******
        if (memcmp(Mqmd.Format, MQFMT_EVENT, MQ_FORMAT_LENGTH))
        {
            printf("Unexpected message format '%8.8s' received\n", Mqmd.Format);
            continue;
        }

        /******
        /* Build the output by parsing the received PCF message, first the */
        /* header, then each of the parameters */
        /******
        pCfh = (MQCFH *)pBuffer;

        if (pCfh->Reason != MQRC_NONE)
        {
            printf("-----\n");
            printf("Event Message Received\n");

            Parm = ParmToString(pCfh->Command);
            if (Parm != NULL)
            {
                printf("Command :%s \n",Parm);
            }
            else
            {
                printf("Command :%" Int32 "d \n",pCfh->Command);
            }
        }
    }
}

```

```

printf("CompCode :%" Int32 "d\n" ,pCfh->CompCode);

Parm = ParmToString(pCfh->Reason);
if (Parm != NULL)
{
    printf("Reason   :%s \n",Parm);
}
else
{
    printf("Reason   :%" Int32 "d \n",pCfh->Reason);
}
}

pPCFCmd = (PMQCHAR) (pCfh+1);
printf("-----\n");
while(pCfh->ParameterCount-->0)
{
    pCfst = (MQCFST *) pPCFCmd;
    switch(pCfst->Type)
    {
        case MQCFT_STRING:
            Parm = ParmToString(pCfst->Parameter);
            if (Parm != NULL)
            {
                printf("%-32s",Parm);
            }
            else
            {
                printf("%-32" Int32 "d",pCfst->Parameter);
            }

            fwrite(pCfst->String, pCfst->StringLength, 1, stdout);
            pPCFCmd += pCfst->StrucLength;
            break;

        default:
            printf("Unrecognised datatype %" Int32 "d returned\n", pCfst->Type);
            goto MOD_EXIT;
    }
    putchar('\n');
}

```

Sample output

This application produces the following form of output:

```

/*****
/* Sample Logger Event Monitor start */
*****/
-----
Event Message Received
Command :Logger Event Command
CompCode :0
Reason   :Logger Status
-----
Queue Manager Name      CSIM

Current Log Extent      AMQA000001
Restart Log Extent      AMQA000001
Media Log Extent        AMQA000001
Log Path                QMCSIM
-----

```

Related concepts

[“Logger event usage” on page 73](#)

Use this page to view how you can use logger events to determine the log extents that are no longer required for queue manager restart, or media recovery.

[“Command event usage” on page 71](#)

Use this page to view how you can use command events to generate an audit trail of the commands that have run

Related reference

[“Logger event generation” on page 72](#)

Use this page to view the situations that cause logger events to be generated and to understand the circumstances in which logger events are not generated

Authority configuration events

Authority configuration events are output when a change is made from any of the security control operations through the command line, MQSC, PCF, or corresponding IBM i commands.

The event data contains the following information:

Origin information

comprises the queue manager from where the change was made, the ID of the user that made the change, and how the change came about, for example by a console command.

Context information

a replica of the context information in the message data from the command message.

Context information is included in the event data when the command was entered as a message on the SYSTEM.ADMIN.COMMAND.QUEUE queue.

Authority Record identity

comprises the profile name, and object type of the authority record.

Object attributes

comprises the values of all the attributes in the authority record.

In the case of change authority record events, two messages are generated, one with the information before the change, the other with the information after the change.

Every event message that is generated is placed on the SYSTEM.ADMIN.CONFIG.EVENT queue.

Related reference

[“Event types” on page 33](#)

Use this page to view the types of instrumentation event that a queue manager or channel instance can report

Authority configuration event generation

Use this page to view the situations that cause authority configuration events to be generated, and to understand the circumstances in which authority configuration events are not generated.

Authority configuration events notify you about changes to the attributes of an authority record. There are three types of authority configuration event:

- [Change Authority Record](#)
- [Delete Authority Record](#)
- [Refresh Authority Record](#)

An authority event message is put to the configuration event queue, when the **CONFIGEV** queue manager attribute is set to *ENABLED* and any of the following commands, or their MQSC equivalent, are issued, even if there is no actual change to the authority record:

- [Delete Authority Record PCF command](#)
- [Set Authority Record PCF command](#)
- [setmqaut](#) control command
- [RVKMQMAUT](#) CL command
- [GRTMQMAUT](#) CL command

When authority configuration events are not generated

The authority configuration event messages are not generated in the following circumstances:

- When a command fails

- When a queue manager encounters an error trying to put a message on the event queue, in which case the command completes, but no event message is generated
- When creating or deleting a queue manager
- When an object is deleted regardless of the **AUTHREC** option on the delete command. The corresponding command event shows that operation, which does not apply to the authority record for individual users.

Related concepts

[“Command events” on page 69](#)

Command events are notifications that an MQSC, or PCF command has run successfully.

Related reference

[REFRESH QMGR](#)

Sample program to monitor instrumentation events on Multiplatforms

The amqsevt sample program formats the instrumentation events that a queue manager can create, and is supplied with IBM MQ for Multiplatforms. The program reads messages from event queues, and formats them into readable strings.

As a sample program, both source and binary are provided. The sample is provided on all the Multiplatforms, including IBM i.

The single binary file amqsevt (or amqsevt.exe) is shipped in the samples fileset and is installed in the samples bin (tools\c\samples\bin or bin64) directory.

The source file amqsevt.c is also shipped in the samples fileset, and is installed in the samples directory, that is, tools\c\samples on Windows.

Note that the program can read from multiple event queues, and subscribe to multiple topics, by using [MQCB](#) to retrieve the messages.

When running as a client, the sample can connect to any queue manager including z/OS.




Attention: You can use the program without specifying any parameters, in which case the program attempts to connect to the default queue manager and read messages from the standard set of event queues (SYSTEM.ADMIN.*.EVENT).

In this situation, the program waits forever for messages, until you press the Enter key to end the program.

However, you are more likely to use the program with the various options described.

Note: You do not need a full distributed IBM MQ installation to access the amqsevt sample program. amqsevt is shipped in the IBM MQ redistributable client, which you can extract and install in a Linux or Windows environment. For more information, see [Redistributable IBM MQ clients](#).

Syntax



Notes:

- ¹ Available only when being used as a client
- ² Queues and topics can have multiple entries

Optional parameters

-m QueueManagerName

Specify a specific queue manager for reading events.

-r Reconnection Options

Auto reconnection options when used as a client. The possible values are:

d

Reconnect the client disabled

r

Reconnect the client

m

Reconnect the queue manager

-b

Browse records only, rather than destructively reading the messages

-c

Selects connection as a client.

-d

Selects the printing mode used in the second example. The MQI constants are printed exactly as they appear in the header files.

-u User ID

Specify a specific user and causes a prompt to appear requesting a password

-w Wait

Causes the program to exit if no event messages have arrived within the number of seconds specified.

Note that, if you do not specify a time, the program only ends normally when you press the Enter key.

-t Topic and

-q Queue

Both the **-q** and **-t** options can be given multiple times on the command line.

Therefore, it is possible to read from some standard queues and also from topics (if events are being sent to them) from a single run of the program.

If no queues or topics are named on the command line, the default event queues are opened.

Note: The program detects if it has connected to a z/OS queue manager as a client, and changes the default set of event queues appropriately, as z/OS does not have the SYSTEM.ADMIN.LOGGER.EVENT queue.

When topics are used, the program uses a non-durable subscription with a managed queue so that everything gets cleaned out when it exits.

-o <output format>

The format of the output. Possible values are:

text

Standard text format; this is the default value.

json

Standard JSON format; any application that understands JSON can take this output and process it directly.

.

Sample output

The following three examples show the output from the program.

The first example uses the default formatting option where the program takes the MQI definition of a field and formats the output to make the output more readable.

```
**** Message (320 Bytes) on Queue SYSTEM.ADMIN.QMGR.EVENT ****
Event Type           : Queue Mgr Event
Reason               : Unknown Alias Base Queue
Event created        : 2015/06/17 13:47:07.02 GMT
  Queue Mgr Name     : V8003_A
  Queue Name         : EVT.NO.BASE.QUEUE
  Base Object Name    : EVT.NOT.DEFINED
  Appl Type          : Unix
  Appl Name          : amqspout
  Base Type          : Queue
```

The second example shows the alternative formatting, using the **-d** option, that does not try to translate MQI constants. This might be preferable for some scripting tools that look for specific MQI values.

```
**** Message (320 Bytes) on Queue SYSTEM.ADMIN.QMGR.EVENT ****
Event Type           : MQCMD_Q_MGR_EVENT
Reason               : MQRC_UNKNOWN_ALIAS_BASE_Q
Event created        : 2015/06/17 13:52:48.18 GMT
  MQCA_Q_MGR_NAME    : V8003_A
  MQCA_Q_NAME        : EVT.NO.BASE.QUEUE
  MQCA_BASE_OBJECT_NAME : EVT.NOT.DEFINED
  MQIA_APPL_TYPE     : MQAT_UNIX
  MQCACF_APPL_NAME   : amqspout
  MQIA_BASE_TYPE     : MQOT_Q
```

The third example shows the JSON output from the command

```
amqsevt -m QM1 -q SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE -o json
```

```
{
  "eventSource" : { "objectName": "SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE",
                    "objectType" : "Queue" },
  "eventType" : {
    "name" : "Activity Trace",
    "value" : 209
  },
  "eventReason" : {
    "name" : "None",
    "value" : 0
  },
  "eventCreation" : {
    "timeStamp" : "2018-07-10T12:44:26Z",
    "epoch" : 1531226666
  },
  "eventData" : {
    "queueMgrName" : "QM1",
    "hostName" : "<yourhostname>",
    "startDate" : "2018-07-10",
    "startTime" : "13:44:25",
    "endDate" : "2018-07-10",
    "endTime" : "13:44:26",
    "commandLevel" : 910,
    ...
  }
}
```

Example usage

The following example shows you how to use more than one queue:

```
amqsevt -m QM1 -q SYSTEM.ADMIN.QMGR.EVENT -q SYSTEM.ADMIN.PERM.EVENT -w 1
```

Related concepts

[“Event monitoring” on page 30](#)

Event monitoring is the process of detecting occurrences of *instrumentation events* in a queue manager network. An instrumentation event is a logical combination of conditions that is detected by a queue manager or channel instance. Such an event causes the queue manager or channel instance to put a special message, called an *event message*, on an event queue.

[“Instrumentation events” on page 31](#)

An instrumentation event is a logical combination of conditions that a queue manager or channel instance detects and puts a special message, called an *event message*, on an event queue.

Related reference

[C programming](#)

[“Sample C program to monitor the logger event queue” on page 73](#)

Use this page to view a sample C program that monitors the logger event queue for new event messages, reads those messages, and puts the contents of the message to stdout.

Message monitoring

Message monitoring is the process of identifying the route a message has taken through a queue manager network. By identifying the types of activities, and the sequence of activities performed on behalf of a message, the message route can be determined.

As a message passes through a queue manager network, various processes perform activities on behalf of the message. Use one of the following techniques to determine a message route:

- The IBM MQ display route application (dspmqrte)
- Activity recording
- Trace-route messaging

These techniques all generate special messages that contain information about the activities performed on the message as it passed through a queue manager network. Use the information returned in these special messages to achieve the following objectives:

- Record message activity.

- Determine the last known location of a message.
- Detect routing problems in your queue manager network.
- Assist in determining the causes of routing problems in your queue manager network.
- Confirm that your queue manager network is running correctly.
- Familiarize yourself with the running of your queue manager network.
- Trace published messages.

Related concepts

[Types of message](#)

Activities and operations

Activities are discrete actions that an application performs on behalf of a message. Activities consist of operations, which are single pieces of work that an application performs.

The following actions are examples of activities:

- A message channel agent (MCA) sends a message from a transmission queue down a channel
- An MCA receives a message from a channel and puts it on its target queue
- An application getting a message from a queue, and putting a reply message in response.
- The IBM MQ publish/subscribe engine processes a message.

Activities consist of one or more *operations*. Operations are single pieces of work that an application performs. For example, the activity of an MCA sending a message from a transmission queue down a channel consists of the following operations:

1. Getting a message from a transmission queue (a *Get* operation).
2. Sending the message down a channel (a *Send* operation).

In a publish/subscribe network, the activity of the IBM MQ publish/subscribe engine processing a message can consist of the following multiple operations:

1. Putting a message to a topic string (a *Put* operation).
2. Zero or more operations for each of the subscribers that are considered for receipt of the message (a *Publish* operation, a *Discarded Publish* operation or an *Excluded Publish* operation).

Information from activities

You can identify the sequence of activities performed on a message by recording information as the message is routed through a queue manager network. You can determine the route of a message through the queue manager network from the sequence of activities performed on the message, and can obtain the following information:

The last known location of a message

If a message does not reach its intended destination, you can determine the last known location of the message from a complete or partial message route.

Configuration issues with a queue manager network

When studying the route of a message through a queue manager network, you might see that the message has not gone where expected. There are many reasons why this can occur, for example, if a channel is inactive, the message might take an alternative route.

For a publish/subscribe application, you can also determine the route of a message being published to a topic and any messages that flow in a queue manager network as a result of being published to subscribers.

In such situations, a system administrator can determine whether there are any problems in the queue manager network, and if appropriate, correct them.

Message routes

Depending on your reason for determining a message route, you can use the following general approaches:

Using activity information recorded for a trace-route message

Trace-route messages record activity information for a specific purpose. You can use them to determine configuration issues with a queue manager network, or to determine the last known location of a message. If a trace-route message is generated to determine the last known location of a message that did not reach its intended destination, it can mimic the original message. This gives the trace-route message the greatest chance of following the route taken by the original message.

The IBM MQ display route application can generate trace-route messages.

Using activity information recorded for the original message

You can enable any message for activity recording and have activity information recorded on its behalf. If a message does not reach its intended destination, you can use the recorded activity information to determine the last known location of the message. By using activity information from the original message, the most accurate possible message route can be determined, leading to the last known location. To use this approach, the original message must be enabled for activity recording.

Warning: Avoid enabling all messages in a queue manager network for activity recording. Messages enabled for activity recording can have many activity reports generated on their behalf. If every message in a queue manager network is enabled for activity recording, the queue manager network traffic can increase to an unacceptable level.

Related concepts

[“Message monitoring” on page 82](#)

Message monitoring is the process of identifying the route a message has taken through a queue manager network. By identifying the types of activities, and the sequence of activities performed on behalf of a message, the message route can be determined.

[“Message route techniques” on page 84](#)

Activity recording and trace-route messaging are techniques that allow you to record activity information for a message as it is routed through a queue manager network.

[“Trace-route messaging” on page 91](#)

Trace-route messaging is a technique that uses *trace-route messages* to record activity information for a message. Trace-route messaging involves sending a trace-route message into a queue manager network.

Related tasks

[Writing your own message channel agents](#)

Message route techniques

Activity recording and trace-route messaging are techniques that allow you to record activity information for a message as it is routed through a queue manager network.

Activity recording

If a message has the appropriate report option specified, it requests that applications generate *activity reports* as it is routed through a queue manager network. When an application performs an activity on behalf of a message, an activity report can be generated, and delivered to an appropriate location. An activity report contains information about the activity that was performed on the message.

The activity information collected using activity reports must be arranged in order before a message route can be determined.

Trace-route messaging

Trace-route messaging is a technique that involves sending a *trace-route message* into a queue manager network. When an application performs an activity on behalf of the trace-route message, activity information can be accumulated in the message data of the trace-route message, or activity reports can be generated. If activity information is accumulated in the message data of the trace-

route message, when it reaches its target queue a trace-route reply message containing all the information from the trace-route message can be generated and delivered to an appropriate location.

Because a trace-route message is dedicated to recording the sequence of activities performed on its behalf, there are more processing options available compared with normal messages that request activity reports.

Comparison of activity recording and trace-route messaging

Both activity recording and trace-route messaging can provide activity information to determine the route a message has taken through a queue manager network. Both methods have their own advantages.

Benefit	Activity recording	Trace-route messaging
Can determine the last known location of a message	Yes	Yes
Can determine configuration issues with a queue manager network	Yes	Yes
Can be requested by any message (is not restricted to use with trace-route messages)	Yes	No
Message data is left unmodified	Yes	No
Message processed normally	Yes	No
Activity information can be accumulated in the message data	No	Yes
Optional message delivery to target queue	No	Yes
If a message is caught in an infinite loop, it can be detected and dealt with	No	Yes
Activity information can be put in order reliably	No	Yes
Application provided to display the activity information	No	Yes

Message route completeness

In some cases it is not possible to identify the full sequence of activities performed on behalf of a message, so only a partial message route can be determined. The completeness of a message route is directly influenced by the queue manager network that the messages are routed through.

MCAs and user-written applications connected to queue managers can record information related to the activities performed on behalf of a message. The recording of activity information is controlled by the queue manager attributes ACTIVREC and ROUTEREC. A queue manager network can determine complete message routes.

How activity information is stored

IBM MQ stores activity information in activity reports, trace-route messages, or trace-route reply messages. In each case the information is stored in a structure called the *Activity* PCF group. A trace-route message or trace-route reply message can contain many Activity PCF groups, depending on the number of activities performed on the message. Activity reports contain one Activity PCF group because a separate activity report is generated for every recorded activity.

With trace-route messaging, additional information can be recorded. This additional information is stored in a structure called the *TraceRoute* PCF group. The TraceRoute PCF group contains a number of PCF structures that are used to store additional activity information, and to specify options that determine how the trace-route message is handled as it is routed through a queue manager network.

Related concepts

[“Activity recording” on page 86](#)

Activity recording is a technique for determining the routes that messages take through a queue manager network. To determine the route that a message has taken, the activities performed on behalf of the message are recorded.

[“Trace-route messaging” on page 91](#)

Trace-route messaging is a technique that uses *trace-route messages* to record activity information for a message. Trace-route messaging involves sending a trace-route message into a queue manager network.

Related reference

[“The TraceRoute PCF group” on page 96](#)

Attributes in the *TraceRoute* PCF group control the behavior of a trace-route message. The *TraceRoute* PCF group is in the message data of every trace-route message.

[“Activity report message data” on page 129](#)

Use this page to view the parameters contained by the *Activity* PCF group in an activity report message. Some parameters are returned only when specific operations have been performed.

Activity recording

Activity recording is a technique for determining the routes that messages take through a queue manager network. To determine the route that a message has taken, the activities performed on behalf of the message are recorded.

When using activity recording, each activity performed on behalf of a message can be recorded in an activity report. An activity report is a type of report message. Each activity report contains information about the application that performed the activity on behalf of the message, when the activity took place, and information about the operations that were performed as part of the activity. Activity reports are typically delivered to a reply-to queue where they are collected together. By studying the activity reports related to a message, you can determine the route that the message took through the queue manager network.

Activity report usage

When messages are routed through a queue manager network, activity reports can be generated. You can use activity report information in the following ways:

Determine the last known location of a message

If a message that is enabled for activity recording does not reach its intended destination, activity reports generated for the message as it was routed through a queue manager network can be studied to determine the last known location of the message.

Determine configuration issues with a queue manager network

A number of messages enabled for activity recording can be sent into a queue manager network. By studying the activity reports related to each message it can become apparent that they have not taken the expected route. There are many reasons why this can occur, for example, a channel could have stopped, forcing the message to take an alternative route. In these situations, a system administrator can determine whether there are any problems in the queue manager network, and if there are, correct them.

Note: You can use activity recording in conjunction with trace-route messages by using the IBM MQ display route application.

Activity report format

Activity reports are PCF messages generated by applications that have performed an activity on behalf of a message. Activity reports are standard IBM MQ report messages containing a message descriptor and message data, as follows:

The message descriptor

- An MQMD structure

Message data

- An embedded PCF header (MQEPH)
- Activity report message data

Activity report message data consists of the *Activity* PCF group, and if generated for a trace-route message, the *TraceRoute* PCF group.

Related reference

[MQMD - Message descriptor](#)

[MQEPH - Embedded PCF header](#)

Controlling activity recording

Enable activity recording at the queue manager level. To enable an entire queue manager network, individually enable every queue manager in the network for activity recording. If you enable more queue managers, more activity reports are generated.

About this task

To generate activity reports for a message as it is routed through a queue manager: define the message to request activity reports; enable the queue manager for activity recording; and ensure that applications performing activities on the message are capable of generating activity reports.

If you do not want activity reports to be generated for a message as it is routed through a queue manager, *disable* the queue manager for activity recording.

Procedure

1. Request activity reports for a message
 - a) In the message descriptor of the message, specify MQRO_ACTIVITY in the *Report* field.
 - b) In the message descriptor of the message, specify the name of a reply-to queue in the *ReplyToQ* field.

Warning: Avoid enabling all messages in a queue manager network for activity recording. Messages enabled for activity recording can have many activity reports generated on their behalf. If every message in a queue manager network is enabled for activity recording, the queue manager network traffic can increase to an unacceptable level.

2. Enable or disable the queue manager for activity recording.

Use the MQSC command ALTER QMGR, specifying the parameter ACTIVREC, to change the value of the queue manager attribute. The value can be:

MSG

The queue manager is enabled for activity recording. Any activity reports generated are delivered to the reply-to queue specified in the message descriptor of the message. This is the default value.

QUEUE

The queue manager is enabled for activity recording. Any activity reports generated are delivered to the local system queue SYSTEM.ADMIN.ACTIVITY.QUEUE. The system queue can also be used to forward activity reports to a common queue.

DISABLED

The queue manager is disabled for activity recording. No activity reports are generated while in the scope of this queue manager.

For example, to enable a queue manager for activity recording and specify that any activity reports generated are delivered to the local system queue SYSTEM.ADMIN.ACTIVITY.QUEUE, use the following MQSC command:

```
ALTER QMGR ACTIVREC(QUEUE)
```

Remember: When you modify the *ACTIVREC* queue manager attribute, a running MCA does not detect the change until the channel is restarted.

3. Ensure that your application uses the same algorithm as MCAs use to determine whether to generate an activity report for a message:
 - a) Verify that the message has requested activity reports to be generated
 - b) Verify that the queue manager where the message currently resides is enabled for activity recording
 - c) Put the activity report on the queue determined by the *ACTIVREC* queue manager attribute

Setting up a common queue for activity reports

To determine the locations of the activity reports related to a specific message when the reports are delivered to the local system queue, it is more efficient to use a common queue on a single node

Before you begin

Set the **ACTIVREC** parameter to enable the queue manager for activity recording and to specify that any activity reports generated are delivered to the local system queue `SYSTEM.ADMIN.ACTIVITY.QUEUE`.

About this task

If a number of queue managers in a queue manager network are set to deliver activity reports to the local system queue, it can be time consuming to determine the locations of the activity reports related to a specific message. Alternatively, use a single node, which is a queue manager that hosts a common queue. All the queue managers in a queue manager network can deliver activity reports to this common queue. The benefit of using a common queue is that queue managers do not have to deliver activity reports to the reply-to queue specified in a message and, when determining the locations of the activity reports related to a message, you query one queue only.

To set up a common queue, perform the following steps:

Procedure

1. Select or define a queue manager as the single node
2. On the single node, select or define a queue for use as the common queue
3. On all queue managers where activity reports are to be delivered to the common queue, redefine the local system queue `SYSTEM.ADMIN.ACTIVITY.QUEUE` as a remote queue definition:
 - a) Specify the name of the single node as the remote queue manager name
 - b) Specify the name of the common queue as the remote queue name

Determining message route information

To determine a message route, obtain the information from the activity reports collected. Determine whether enough activity reports are on the reply-to queue to enable you to determine the required information and arrange the activity reports in order.

About this task

The order that activity reports are put on the reply-to queue does not necessarily correlate to the order in which the activities were performed. You must order activity reports manually, unless they are generated for a trace-route message, in which case you can use the IBM MQ display route application to order the activity reports.

Determine whether enough activity reports are on the reply-to queue for you to obtain the necessary information:

Procedure

1. Identify all related activity reports on the reply-to queue by comparing identifiers of the activity reports and the original message. Ensure you set the report option of the original message such that the activity reports can be correlated with the original message.

2. Order the identified activity reports from the reply-to queue.
You can use the following parameters from the activity report:

OperationType

The types of operations performed might enable you to determine the activity report that was generated directly before, or after, the current activity report.

For example, an activity report details that an MCA sent a message from a transmission queue down a channel. The last operation detailed in the activity report has an *OperationType* of send and details that the message was sent using the channel, CH1, to the destination queue manager, QM1. This means that the next activity performed on the message will have occurred on queue manager, QM1, and that it will have begun with a receive operation from channel, CH1. By using this information you can identify the next activity report, providing it exists and has been acquired.

OperationDate and OperationTime

You can determine the general order of the activities from the dates and times of the operations in each activity report.

Warning: Unless every queue manager in the queue manager network has their system clocks synchronized, ordering by date and time does not guarantee that the activity reports are in the correct sequence. You must establish the order manually.

The order of the activity reports represents the route, or partial route, that the message took through the queue manager network.

3. Obtain the information you need from the activity information in the ordered activity reports.
If you have insufficient information about the message, you might be able to acquire further activity reports.

Retrieving further activity reports

To determine a message route, sufficient information must be available from the activity reports collected. If you retrieve the activity reports related to a message from the reply-to queue that the message specified, but you not have the necessary information, look for further activity reports.

About this task

To determine the locations of any further activity reports, perform the following steps:

Procedure

1. For any queue managers in the queue manager network that deliver activity reports to a common queue, retrieve activity reports from the common queue that have a *CorrelId* that matches the *MsgId* of the original message.
2. For any queue managers in the queue manager network that do not deliver activity reports to a common queue, retrieve activity reports as follows:
 - a) Examine the existing activity reports to identify queue managers through which the message was routed.
 - b) For these queue managers, identify the queue managers that are enabled for activity recording.
 - c) For these queue managers, identify any that did not return activity reports to the specified reply-to queue.
 - d) For each of the queue managers that you identify, check the system queue `SYSTEM.ADMIN.ACTIVITY.QUEUE` and retrieve any activity reports that have a *CorrelId* that matches the *MsgId* of the original message.
 - e) If you find no activity reports on the system queue, check the queue manager dead letter queue, if one exists.
An activity report can only be delivered to a dead letter queue if the report option, `MQRO_DEAD_LETTER_Q`, is set.

3. Arrange all the acquired activity reports in order.

The order of the activity reports then represents the route, or partial route, that the message took.

4. Obtain the information you need from the activity information in the ordered activity reports.

In some circumstances, recorded activity information cannot reach the specified reply-to queue, a common queue, or a system queue.

Circumstances where activity information is not acquired

To determine the complete sequence of activities performed on behalf of a message, information related to every activity must be acquired. If the information relating to any activity has not been recorded, or has not been acquired, you can determine only a partial sequence of activities.

Activity information is not recorded in the following circumstances:

- The message is processed by a queue manager earlier than IBM WebSphere® MQ 6.0.
- The message is processed by a queue manager that is not enabled for activity recording.
- The application that expected to process the message is not running.

Recorded activity information is unable to reach the specified reply-to queue in the following circumstances:

- There is no channel defined to route activity reports to the reply-to queue.
- The channel to route activity reports to the reply-to queue is not running.
- The remote queue definition to route activity reports back to the queue manager where the reply-to queue resides (the queue manager alias), is not defined.
- The user that generated the original message does not have open, or put, authority to the queue manager alias.
- The user that generated the original message does not have open, or put, authority to the reply-to queue.
- The reply-to queue is put inhibited.

Recorded activity information is unable to reach the system queue, or a common queue, in the following circumstances:

- If a common queue is to be used and there is no channel defined to route activity reports to the common queue.
- If a common queue is to be used and the channel to route activity reports to the common queue is not running.
- If a common queue is to be used and the system queue is incorrectly defined.
- The user that generated the original message does not have open, or put, authority to the system queue.
- The system queue is put inhibited.
- If a common queue is to be used and the user that generated the original message does not have open, or put, authority to the common queue.
- If a common queue is to be used and the common queue is put inhibited.

In these circumstances, providing the activity report does not have the report option MQRO_DISCARD_MSG specified, the activity report can be retrieved from a dead letter queue if one was defined on the queue manager where the activity report was rejected. An activity report will only have this report option specified if the original message, from which the activity report was generated, had both MQRO_PASS_DISCARD_AND_EXPIRY and MQRO_DISCARD_MSG specified in the Report field of the message descriptor.

Trace-route messaging

Trace-route messaging is a technique that uses *trace-route messages* to record activity information for a message. Trace-route messaging involves sending a trace-route message into a queue manager network.

As the trace-route message is routed through the queue manager network, activity information is recorded. This activity information includes information about the applications that performed the activities, when they were performed, and the operations that were performed as part of the activities. You can use the information recorded using trace-route messaging for the following purposes:

To determine the last known location of a message

If a message does not reach its intended destination, you can use the activity information recorded for a trace-route message to determine the last known location of the message. A trace-route message is sent into a queue manager network with the same target destination as the original message, intending that it follows the same route. Activity information can be accumulated in the message data of the trace-route message, or recorded using activity reports. To increase the probability that the trace-route message follows the same route as the original message, you can modify the trace-route message to mimic the original message.

To determine configuration issues with a queue manager network

Trace-route messages are sent into a queue manager network and activity information is recorded. By studying the activity information recorded for a trace-route message, it can become apparent that the trace-route message did not follow the expected route. There are many reasons why this can occur, for example, a channel might be inactive, forcing the message to take an alternative route. In these situations, a system administrator can determine whether there are any problems in the queue manager network, and if there are, correct them.

You can use the IBM MQ display route application to configure, generate, and put trace-route messages into a queue manager network.

Warning: If you put a trace-route message to a distribution list, the results are undefined.

Related concepts

[“Trace-route message reference” on page 146](#)

Use this page to obtain an overview of the trace-route message format. The trace-route message data includes parameters that describe the activities that the trace-route message has caused

How activity information is recorded

With trace-route messaging, you can record activity information in the message data of the trace-route message, or use activity reports. Alternatively, you can use both techniques.

Accumulating activity information in the message data of the trace-route message

As a trace-route message is routed through a queue manager network, information about the activities performed on behalf of the trace-route message can be accumulated in the message data of the trace-route message. The activity information is stored in *Activity* PCF groups. For every activity performed on behalf of the trace-route message, an *Activity* PCF group is written to the end of the PCF block in the message data of the trace-route message.

Additional activity information is recorded in trace-route messaging, in a PCF group called the *TraceRoute* PCF group. The additional activity information is stored in this PCF group, and can be used to help determine the sequence of recorded activities. This technique is controlled by the *Accumulate* parameter in the *TraceRoute* PCF group.

Recording activity information using activity reports

As a trace-route message is routed through a queue manager network, an activity report can be generated for every activity that was performed on behalf of the trace-route message. The activity information is stored in the *Activity* PCF group. For every activity performed on behalf of a trace-route message, an activity report is generated containing an *Activity* PCF group. Activity recording for trace-route messages works in the same way as for any other message.

Activity reports generated for trace-route messages contain additional activity information compared to the those generated for any other message. The additional information is returned in a *TraceRoute* PCF group. The information contained in the *TraceRoute* PCF group is accurate only from the time the activity report was generated. You can use the additional information to help determine the sequence of activities performed on behalf of the trace-route message.

Acquiring recorded activity information

When a trace-route message has reached its intended destination, or is discarded, the method that you use to acquire the activity information depends on how that information was recorded.

Before you begin

If you are unfamiliar with activity information, refer to [“How activity information is recorded”](#) on page 91.

About this task

Use the following methods to acquire the activity information after the trace-route message has reached its intended destination, or is discarded:

Procedure

- Retrieve the trace-route message.

The *Deliver* parameter, in the *TraceRoute* PCF group, controls whether a trace-route message is placed on the target queue on arrival, or whether it is discarded. If the trace-route message is delivered to the target queue, you can retrieve the trace-route message from this queue. Then, you can use the IBM MQ display route application to display the activity information.

To request that activity information is accumulated in the message data of a trace-route message, set the *Accumulate* parameter in the *TraceRoute* PCF group to `MQRROUTE_ACCUMULATE_IN_MSG`.

- Use a trace-route reply message.

When a trace-route message reaches its intended destination, or the trace-route message cannot be routed any further in a queue manager network, a trace-route reply message can be generated. A trace-route reply message contains a duplicate of all the activity information from the trace-route message, and is either delivered to a specified reply-to queue, or the system queue `SYSTEM.ADMIN.TRACE.ROUTE.QUEUE`. You can use the IBM MQ display route application to display the activity information.

To request a trace-route reply message, set the *Accumulate* parameter in the *TraceRoute* PCF group to `MQRROUTE_ACCUMULATE_AND_REPLY`.

- Use activity reports.

If activity reports are generated for a trace-route message, you must locate the activity reports before you can acquire the activity information. Then, to determine the sequence of activities, you must order the activity reports.

Controlling trace-route messaging

Enable trace-route messaging at the queue manager level, so that applications in the scope of that queue manager can write activity information to a trace-route message. To enable an entire queue manager network, individually enable every queue manager in the network for trace-route messaging. If you enable more queue managers, more activity reports are generated.

Before you begin

If you are using activity reports to record activity information for a trace-route message, refer to [“Controlling activity recording”](#) on page 87.

About this task

To record activity information for a trace-route message as it is routed through a queue manager, perform the following steps:

Procedure

- Define how activity information is to be recorded for the trace-route message.
Refer to [“Generating and configuring a trace-route message” on page 95](#)
- If you want to accumulate activity information in the trace-route message, ensure that the queue manager is enabled for trace-route messaging
- If you want to accumulate activity information in the trace-route message, ensure that applications performing activities on the trace-route message are capable of writing activity information to the message data of the trace-route message

Related concepts

[“Generating and configuring a trace-route message” on page 95](#)

A trace-route message comprises specific message descriptor and message data parts. To generate a trace-route message, either create the message manually or use the IBM MQ display route application.

Related tasks

[“Controlling activity recording” on page 87](#)

Enable activity recording at the queue manager level. To enable an entire queue manager network, individually enable every queue manager in the network for activity recording. If you enable more queue managers, more activity reports are generated.

Enabling queue managers for trace-route messaging

To control whether queue managers are enabled or disabled for trace-route messaging use the queue manager attribute ROUTEREC.

Use the MQSC command ALTER QMGR, specifying the parameter ROUTEREC to change the value of the queue manager attribute. The value can be any of the following values:

MSG

The queue manager is enabled for trace-route messaging. Applications within the scope of the queue manager can write activity information to the trace-route message.

If the *Accumulate* parameter in the *TraceRoute* PCF group is set as MQROUTE_ACCUMULATE_AND_REPLY, and the next activity to be performed on the trace-route message:

- is a discard
- is a put to a local queue (target queue or dead-letter queue)
- will cause the total number of activities performed on the trace-route message to exceed the value of parameter the *MaxActivities*, in the *TraceRoute* PCF group .

a trace-route reply message is generated, and delivered to the reply-to queue specified in the message descriptor of the trace-route message.

QUEUE

The queue manager is enabled for trace-route messaging. Applications within the scope of the queue manager can write activity information to the trace-route message.

If the *Accumulate* parameter in the *TraceRoute* PCF group is set as MQROUTE_ACCUMULATE_AND_REPLY, and the next activity to be performed on the trace-route message:

- is a discard
- is a put to a local queue (target queue or dead-letter queue)
- will cause the total number of activities performed on the trace-route message to exceed the value of parameter the *MaxActivities*, in the *TraceRoute* PCF group .

a trace-route reply message is generated, and delivered to the local system queue SYSTEM.ADMIN.TRACE.ROUTE.QUEUE.

DISABLED

The queue manager is disabled for trace-route messaging. Activity information is not accumulated in the trace-route message, however the *TraceRoute* PCF group can be updated while in the scope of this queue manager.

For example, to disable a queue manager for trace-route messaging, use the following MQSC command:

```
ALTER QMGR ROUTEREC(DISABLED)
```

Remember: When you modify the *ROUTEREC* queue manager attribute, a running MCA does not detect the change until the channel is restarted.

Enabling applications for trace-route messaging

To enable trace-route messaging for a user application, base your algorithm on the algorithm used by message channel agents (MCAs)

Before you begin

If you are not familiar with the format of a trace-route message, see [“Trace-route message reference”](#) on page 146.

About this task

Message channel agents (MCAs) are enabled for trace-route messaging. To enable a user application for trace-route messaging, use the following steps from the algorithm that MCAs use:

Procedure

1. Determine whether the message being processed is a trace-route message.
If the message does not conform to the format of a trace-route message, the message is not processed as a trace-route message.
2. Determine whether activity information is to be recorded.
If the detail level of the performed activity is not less than the level of detail specified by the *Detail* parameter, activity information is recorded under specific circumstances. This information is only recorded if the trace-route message requests accumulation, and the queue manager is enabled for trace-route messaging, or if the trace-route message requests an activity report and the queue manager is enabled for activity recording.
 - If activity information is to be recorded, increment the *RecordedActivities* parameter.
 - If activity information is not to be recorded, increment the *UnrecordedActivities* parameter.
3. Determine whether the total number of activities performed on the trace-route message exceeds the value of the *MaxActivities* parameter.

The total number of activities is the sum of *RecordedActivities*, *UnrecordedActivities*, and *DiscontinuityCount*.

If the total number of activities exceeds *MaxActivities*, reject the message with feedback MQFB_MAX_ACTIVITIES.
4. If value of *Accumulate* is set as MQROUTE_ACCUMULATE_IN_MSG or MQROUTE_ACCUMULATE_AND_REPLY, and the queue manager is enabled for trace-route messaging, write an Activity PCF group to the end of the PCF block in the message data of a trace-route message.
5. Deliver the trace-route message to a local queue.
 - If the parameter, *Deliver*, is specified as MQROUTE_DELIVER_NO, reject the trace-route message with feedback MQFB_NOT_DELIVERED.
 - If the parameter, *Deliver*, is specified as MQROUTE_DELIVER_YES, deliver the trace-route message to the local queue.
6. Generate a trace-route reply message if all the following conditions are true:

- The trace-route message was delivered to a local queue or rejected
- The value of the parameter, *Accumulate*, is MQROUTE_ACCUMULATE_AND_REPLY
- The queue manager is enabled for trace-route messaging

The trace-route reply message is put on the queue determined by the ROUTEREC queue manager attribute.

7. If the trace-route message requested an activity report and the queue manager is enabled for activity recording, generate an activity report.

The activity report is put on the queue determined by the ACTIVREC queue manager attribute.

Generating and configuring a trace-route message

A trace-route message comprises specific message descriptor and message data parts. To generate a trace-route message, either create the message manually or use the IBM MQ display route application.

A trace-route message consists of the following parts:

Message descriptor

An MQMD structure, with the *Format* field set to MQFMT_ADMIN or MQFMT_EMBEDDED_PCF.

Message data

One of the following combinations:

- A PCF header (MQCFH) and trace-route message data, if *Format* is set to MQFMT_ADMIN
- An embedded PCF header (MQEPH), trace-route message data, and additional user-specified message data, if *Format* is set to MQFMT_EMBEDDED_PCF

The trace-route message data consists of the *TraceRoute* PCF group and one or more *Activity* PCF groups.

Manual generation

When generating a trace-route message manually, an *Activity* PCF group is not required. *Activity* PCF groups are written to the message data of the trace-route message when an MCA or user-written application performs an activity on its behalf.

The IBM MQ display route application

Use the IBM MQ display route application, **dspmqrte**, to configure, generate and put a trace-route message into a queue manager network. Set the *Format* parameter in the message descriptor to MQFMT_ADMIN. You cannot add user data to the trace-route message generated by the IBM MQ display route application.

Restriction: dspmqrte cannot be issued on IBM MQ for z/OS queue managers. If you want the first queue manager the trace-route message is routed through to be a queue manager of this type, connect to the queue manager as a client using the optional parameter -c.

Mimicking the original message

When using a trace-route message to determine the route another message has taken through a queue manager network, the more closely a trace-route message mimics the original message, the greater the chance that the trace-route message will follow the same route as the original message.

The following message characteristics can affect where a message is forwarded to within a queue manager network:

Priority

The priority can be specified in the message descriptor of the message.

Persistence

The persistence can be specified in the message descriptor of the message.

Expiration

The expiration can be specified in the message descriptor of the message.

Report options

Report options can be specified in the message descriptor of the message.

Message size

To mimic the size of a message, additional data can be written to the message data of the message. For this purpose, additional message data can be meaningless.

Tip: The IBM MQ display route application cannot specify message size.

Message data

Some queue manager networks use content based routing to determine where messages are forwarded. In these cases the message data of the trace-route message needs to be written to mimic the message data of the original message.

Tip: The IBM MQ display route application cannot specify message data.

The TraceRoute PCF group

Attributes in the *TraceRoute* PCF group control the behavior of a trace-route message. The *TraceRoute* PCF group is in the message data of every trace-route message.

The following table lists the parameters in the *TraceRoute* group that an MCA recognizes. Further parameters can be added if user-written applications are written to recognize them, as described in “Additional activity information” on page 101.

Table 16. <i>TraceRoute</i> PCF group	
Parameter	Type
TraceRoute	MQCFGR
Detail	MQCFIN
RecordedActivities	MQCFIN
UnrecordedActivities	MQCFIN
DiscontinuityCount	MQCFIN
MaxActivities	MQCFIN
Accumulate	MQCFIN
Forward	MQCFIN
Deliver	MQCFIN

Descriptions of each parameter in the *TraceRoute* PCF group follows:

Detail

Specifies the detail level of activity information that is to be recorded. The value can be any of the following values:

MQROUTE_DETAIL_LOW

Only activities performed by user application are recorded.

MQROUTE_DETAIL_MEDIUM

Activities specified in MQROUTE_DETAIL_LOW should be recorded. Additionally, activities performed by MCAs are recorded.

MQROUTE_DETAIL_HIGH

Activities specified in MQROUTE_DETAIL_LOW, and MQROUTE_DETAIL_MEDIUM should be recorded. MCAs do not record any further activity information at this level of detail. This option is only available to user applications that are to record further activity information. For example, if a user application determines the route a message takes by considering certain message characteristics, the information about the routing logic could be included with this level of detail.

RecordedActivities

Specifies the number of recorded activities performed on behalf of the trace-route message. An activity is considered to be recorded if information about it has been written to the trace-route message, or if an activity report has been generated. For every recorded activity, *RecordedActivities* increments by one.

UnrecordedActivities

Specifies the number of unrecorded activities performed on behalf of the trace-route message. An activity is considered to be unrecorded if an application that is enabled for trace-route messaging neither accumulates, nor writes the related activity information to an activity report.

An activity performed on behalf of a trace-route message is unrecorded in the following circumstances:

- The detail level of the performed activity is less than the level of detail specified by the parameter *Detail*.
- The trace-route message requests an activity report but not accumulation, and the queue manager is not enabled for activity recording.
- The trace-route message requests accumulation but not an activity report, and the queue manager is not enabled for trace-route messaging.
- The trace-route message requests both accumulation and an activity report, and the queue manager is not enabled for activity recording and trace route messaging.
- The trace-route message requests neither accumulation nor an activity report.

For every unrecorded activity the parameter, *UnrecordedActivities*, increments by one.

DiscontinuityCount

Specifies the number of times the trace-route message has been routed through a queue manager with applications that were not enabled for trace-route messaging. This value is incremented by the queue manager. If this value is greater than 0, only a partial message route can be determined.

MaxActivities

Specifies the maximum number of activities that can be performed on behalf of the trace-route message.

The total number of activities is the sum of *RecordedActivities*, *UnrecordedActivities*, and *DiscontinuityCount*. The total number of activities must not exceed the value of *MaxActivities*.

The value of *MaxActivities* can be:

A positive integer

The maximum number of activities.

If the maximum number of activities is exceeded, the trace-route message is rejected with feedback MQFB_MAX_ACTIVITIES. This can prevent the trace-route message from being forwarded indefinitely if caught in an infinite loop.

MQROUTE_UNLIMITED_ACTIVITIES

An unlimited number of activities can be performed on behalf of the trace-route message.

Accumulate

Specifies the method used to accumulate activity information. The value can be any of the following values:

MQROUTE_ACCUMULATE_IN_MSG

If the queue manager is enabled for trace-route messaging, activity information is accumulated in the message data of the trace-route message.

If this value is specified, the trace-route message data consists of the following:

- The *TraceRoute* PCF group.
- Zero or more *Activity* PCF groups.

MQROUTE_ACCUMULATE_AND_REPLY

If the queue manager is enabled for trace-route messaging, activity information is accumulated in the message data of the trace-route message, and a trace-route reply message is generated if any of the following occur:

- The trace-route message is discarded by an IBM MQ queue manager.

- The trace-route message is put to a local queue (target queue or dead-letter queue) by an IBM MQ queue manager.
- The number of activities performed on the trace-route message exceeds the value of *MaxActivities*.

If this value is specified, the trace-route message data consists of the following:

- The *TraceRoute* PCF group.
- Zero or more *Activity* PCF groups.

MQROUTE_ACCUMULATE_NONE

Activity information is not accumulated in the message data of the trace-route message.

If this value is specified, the trace-route message data consists of the following:

- The *TraceRoute* PCF group.

Forward

Specifies where a trace-route message can be forwarded to. The value can be:

MQROUTE_FORWARD_IF_SUPPORTED

The trace-route message is only forwarded to queue managers that will honor the value of the *Deliver* parameter from the *TraceRoute* group.

MQROUTE_FORWARD_ALL

The trace-route message is forwarded to any queue manager, regardless of whether the value of the *Deliver* parameter will be honored.

Queue managers use the following algorithm when determining whether to forward a trace-route message to a remote queue manager:

1. Determine whether the remote queue manager is capable of supporting trace-route messaging.
 - If the remote queue manager is capable of supporting trace-route messaging, the algorithm continues to step [“4” on page 98](#).
 - If the remote queue manager is not capable of supporting trace-route messaging, the algorithm continues to step [“2” on page 98](#).
2. Determine whether the *Deliver* parameter from the *TraceRoute* group contains any unrecognized delivery options in the MQROUTE_DELIVER_REJ_UNSUP_MASK bit mask.
 - If any unrecognized delivery options are found, the trace-route message is rejected with feedback MQFB_UNSUPPORTED_DELIVERY.
 - If no unrecognized delivery options are found, the algorithm continues to step [“3” on page 98](#).
3. Determine the value of the parameter *Deliver* from the *TraceRoute* PCF group in the trace-route message.
 - If *Deliver* is specified as MQROUTE_DELIVER_YES, the trace-route message is forwarded to the remote queue manager.
 - If *Deliver* is specified as MQROUTE_DELIVER_NO, the algorithm continues to step [“4” on page 98](#).
4. Determine whether the *Forward* parameter from the *TraceRoute* group contains any unrecognized forwarding options in the MQROUTE_FORWARDING_REJ_UNSUP_MASK bit mask.
 - If any unrecognized forwarding options are found, the trace-route message is rejected with feedback MQFB_UNSUPPORTED_FORWARDING.
 - If no unrecognized forwarding options are found, the algorithm continues to step [“5” on page 98](#).
5. Determine the value of the parameter *Forward* from the *TraceRoute* PCF group in the trace-route message.
 - If *Forward* is specified as MQROUTE_FORWARD_IF_SUPPORTED, the trace-route message is rejected with feedback MQFB_NOT_FORWARDED.

- If *Forward* is specified as `MQROUTE_FORWARD_ALL`, trace-route message can be forwarded to the remote queue manager.

Deliver

Specifies the action to be taken if the trace-route message reaches its intended destination. User-written applications must check this attribute before placing a trace-route message on its target queue. The value can be any of the following values:

MQROUTE_DELIVER_YES

On arrival, the trace-route message is put on the target queue. Any application performing a get operation on the target queue can retrieve the trace-route message.

MQROUTE_DELIVER_NO

On arrival, the trace-route message is not delivered to the target queue. The message is processed according to its report options.

Setting up a common queue for trace-route reply messages

To determine the locations of the trace-route reply messages related to a specific message when the reports are delivered to the local system queue, it is more efficient to use a common queue on a single node

Before you begin

Set the **ROUTEREC** parameter to enable the queue manager for trace-route messaging and to specify that any trace-route reply messages generated are delivered to the local system queue `SYSTEM.ADMIN.TRACE.ROUTE.QUEUE`.

About this task

If a number of queue managers in a queue manager network are set to deliver trace-route reply messages to the local system queue, it can be time consuming to determine the locations of the trace-route reply messages related to a specific message. Alternatively, use a single node, which is a queue manager that hosts a common queue. All the queue managers in a queue manager network can deliver trace-route reply messages to this common queue. The benefit of using a common queue is that queue managers do not have to deliver trace-route reply messages to the reply-to queue specified in a message and, when determining the locations of the trace-route reply messages related to a message, you query one queue only.

To set up a common queue, perform the following steps:

Procedure

1. Select or define a queue manager as the single node
2. On the single node, select or define a queue for use as the common queue
3. On all queue managers that forward trace-route reply messages to the common queue, redefine the local system queue `SYSTEM.ADMIN.TRACE.ROUTE.QUEUE` as a remote queue definition
 - a) Specify the name of the single node as the remote queue manager name
 - b) Specify the name of the common queue as the remote queue name

Acquiring and using recorded information

Use any of the following techniques to acquire recorded activity information for a trace-route message

Note that the circumstances in which activity information is not acquired apply also to trace-route reply messages.

Activity information is not recorded when a trace-route message is processed by a queue manager that is disabled for both activity recording and trace-route messaging.

Acquiring information from trace-route reply messages

To acquire activity information you locate the trace-route reply message. Then you retrieve the message and analyze the activity information.

About this task

You can acquire activity information from a trace-route reply message only if you know the location of the trace-route reply message. Locate the message and process the activity information as follows:

Procedure

1. Check the reply-to queue that was specified in the message descriptor of the trace-route message. If the trace-route reply message is not on the reply-to queue, check the following locations:
 - The local system queue, SYSTEM.ADMIN.TRACE.ROUTE.QUEUE, on the target queue manager of the trace-route message
 - The common queue, if you have set up a common queue for trace-route reply messages
 - The local system queue, SYSTEM.ADMIN.TRACE.ROUTE.QUEUE, on any other queue manager in the queue manager network, which can occur if the trace-route message has been put to a dead-letter queue, or the maximum number of activities was exceeded
2. Retrieve the trace-route reply message
3. Use the IBM MQ display route application to display the recorded activity information
4. Study the activity information and obtain the information that you need

Acquiring information from trace-route messages

To acquire activity information you locate the trace-route message, which must have the appropriate parameters in the *TraceRoute* PCF group. Then you retrieve the message and analyze the activity information.

About this task

You can acquire activity information from a trace-route message only if you know the location of the trace-route message and it has the parameter *Accumulate* in the *TraceRoute* PCF group specified as either MQRROUTE_ACCUMULATE_IN_MSG or MQRROUTE_ACCUMULATE_AND_REPLY.

For the trace-route message to be delivered to the target queue the *Deliver* parameter in the *TraceRoute* PCF group must be specified as MQRROUTE_DELIVER_YES.

Procedure

1. Check the target queue. If the trace-route message is not on the target queue, you can try to locate the trace-route message using a trace-route message enabled for activity recording. With the generated activity reports try to determine the last known location of the trace-route message.
2. Retrieve the trace-route message
3. Use the IBM MQ display route application to display the recorded activity information
4. Study the activity information and obtain the information that you need

Acquiring information from activity reports

To acquire activity information you locate the activity report, which must have the report option specified in the message descriptor. Then you retrieve the activity report and analyze the activity information.

About this task

You can acquire activity information from an activity report only if you know the location of the activity report and the report option MQRROUTE_ACTIVITY was specified in the message descriptor of the trace-route message.

Procedure

1. Locate and order the activity reports generated for a trace-route message.
When you have located the activity reports, you can order them manually or use the IBM MQ display route application to order and display the activity information automatically.
2. Study the activity information and obtain the information that you need

Additional activity information

As a trace-route message is routed through a queue manager network, user applications can record additional information by including one or more additional PCF parameters when writing the *Activity* group to the message data of the trace-route message or activity report.

Additional activity information can help system administrators to identify the route taken by a trace-route message took, or why that route was taken.

If you use the IBM MQ display route application to display the recorded information for a trace-route message, any additional PCF parameters can only be displayed with a numeric identifier, unless the parameter identifier of each parameter is recognized by the IBM MQ display route application. To recognize a parameter identifier, additional information must be recorded using the following PCF parameters. Include these PCF parameters in an appropriate place in the *Activity* PCF group.

GroupName

Table 17. Group name	
Description	Grouped parameters specifying the additional information.
Identifier	MQGACF_VALUE_NAMING.
Data type	MQCFGR
Parameters in group	<i>ParameterName</i> <i>ParameterValue</i>

ParameterName

Table 18. Parameter name	
Description	Contains the name to be displayed by the IBM MQ display route application, which puts the value of <i>ParameterValue</i> into context.
Identifier	MQCA_VALUE_NAME.
Data type	MQCFST
Included in PCF group:	<i>GroupName</i> .
Value:	The name to be displayed.

ParameterValue

Table 19. Parameter value	
Description	Contains the value to be displayed by the IBM MQ display route application.
Identifier:	The PCF structure identifier for the additional information.
Data type:	The PCF structure data type for the additional information.
Included in PCF group:	<i>GroupName</i> .
Value:	The value to be displayed.

Examples of recording additional activity information

The following examples illustrate how a user application can record additional information when performing an activity on behalf of a trace-route message. In both examples, the IBM MQ display route application is used to generate a trace-route message, and display the activity information returned to it.

Recording additional activity information: Example 1

Additional activity information is recorded by a user application in a format where the parameter identifier *is not* recognized by the IBM MQ display route application.

1. The IBM MQ display route application is used to generate and put a trace-route message into a queue manager network. The necessary options are set to request the following:
 - Activity information is accumulated in the message data of the trace-route message.
 - On arrival at the target queue the trace-route message is discarded, and a trace-route reply message is generated and delivered to a specified reply-to queue.
 - On receipt of the trace-route reply message, the IBM MQ display route application displays the accumulated activity information.

The trace-route message is put into the queue manager network.

2. As the trace-route message is routed through the queue manager network a user application, that is enabled for trace-route messaging, performs a low detail activity on behalf of the message. In addition to writing the standard activity information to the trace-route message, the user application writes the following PCF parameter to the end of the Activity group:

ColorValue

Identifier

65536

Data type

MQCFST

Value

'Red'

This additional PCF parameter gives further information about the activity that was performed, however it is written in a format where the parameter identifier *is not* recognized by the IBM MQ display route application.

3. The trace-route messages reaches the target queue and a trace-route reply message is returned to the IBM MQ display route application. The additional activity information is displayed as follows:

```
65536: 'Red'
```

The IBM MQ display route application does not recognize the parameter identifier of the PCF parameter and displays it as a numeric value. The context of the additional information is not clear.

For an example of when the IBM MQ display route application does recognize the parameter identifier of the PCF parameter, see [“Recording additional activity information: Example 2” on page 102](#).

Recording additional activity information: Example 2

Additional activity information is recorded by a user application in a format where the parameter identifier *is* recognized by the IBM MQ display route application.

1. The IBM MQ display route application is used to generate and put a trace-route message into a queue manager network in the same fashion as in [“Recording additional activity information: Example 1” on page 102](#).
2. As the trace-route message is routed through the queue manager network a user application, that is enabled for trace-route messaging, performs a low detail activity on behalf of the message. In addition to writing the standard activity information to the trace-route message, the user application writes the following PCF parameters to the end of the Activity group:

ColorInfo

Table 20. Color information	
Description	Grouped parameters specifying information about a color.
Identifier:	MQGACF_VALUE_NAMING.
Data type:	MQCFGR.
Parameters in group:	<i>ColorName</i> <i>ColorValue</i>

ColorName

Table 21. Color name	
Description	Contains the name to be displayed by the IBM MQ display route application which puts the value of <i>ColorValue</i> into context.
Identifier:	MQCA_VALUE_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>ColorInfo</i> .
Value:	'Color'

ColorValue

Table 22. Color value	
Description	Contains the value to be displayed by the IBM MQ display route application.
Identifier:	65536.
Data type:	MQCFST.
Included in PCF group:	<i>ColorInfo</i> .
Value:	'Red'

These additional PCF parameters gives further information about the activity that was performed. These PCF parameters are written in a format where the parameter identifier is recognized by the IBM MQ display route application.

3. The trace-route messages reaches the target queue and a trace-route reply message is returned to the IBM MQ display route application. The additional activity information is displayed as follows:

```
Color: 'Red'
```

The IBM MQ display route application recognizes that the parameter identifier of the PCF structure containing the value of the additional activity information has a corresponding name. The corresponding name is displayed instead of the numeric value.

IBM MQ display route application

Use the IBM MQ display route application (**dspmqrte**) to work with trace-route messages and activity information related to a trace-route message, using a command-line interface.  The IBM MQ display route application is not shipped on IBM MQ for z/OS, but you can run it from a distributed

installation and connect it as a client to an IBM MQ for z/OS queue manager by specifying the **-c** parameter when issuing the **dspmqrte** command.

You can use the IBM MQ [dspmqrte](#) display route application for the following purposes:

- To configure, generate, and put a trace-route message into a queue manager network.

By putting a trace-route message into a queue manager network, activity information can be collected and used to determine the route that the trace-route message took. You can specify the characteristics of the trace-route messages as follows:

- The destination of the trace-route message.
- How the trace-route message mimics another message.
- How the trace-route message should be handled as it is routed through a queue manager network.
- Whether activity recording or trace-route messaging are used to record activity information.

- To order and display activity information related to a trace-route message.

If the IBM MQ display route application has put a trace-route message into a queue manager network, after the related activity information has been returned, the information can be ordered and displayed immediately. Alternatively, the IBM MQ display route application can be used to order, and display, activity information related to a trace-route message that was previously generated.

Related reference

[dspmqrte](#)

Parameters for trace-route messages

Use this page to obtain an overview of the parameters provided by the IBM MQ display route application, **dspmqrte**, to determine the characteristics of a trace-route message, including how it is treated as it is routed through a queue manager network.

Related reference

[dspmqrte](#)

Queue manager connection

Use this page to specify the queue manager that the IBM MQ display route application connects to

-c

Specifies that the IBM MQ display route application connects as a client application.

If you do not specify this parameter, the IBM MQ display route application does not connect as a client application.

-m QMgrName

The name of the queue manager to which the IBM MQ display route application connects. The name can contain up to 48 characters.

If you do not specify this parameter, the default queue manager is used.

The target destination

Use this page to specify the target destination of a trace-route message

-q TargetQName

If the IBM MQ display route application is being used to send a trace-route message into a queue manager network, *TargetQName* specifies the name of the target queue.

-ts TargetTopicString

Specifies the topic string.

-qm TargetQMGr

Qualifies the target destination; normal queue manager name resolution will then apply. The target destination is specified with *-q TargetQName* or *-ts TargetTopicString*.

If you do not specify this parameter, the queue manager to which the IBM MQ display route application is connected is used as the target queue manager.

-o

Specifies that the target destination 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 destination is opened with option MQOO_BIND_NOT_FIXED.

If you do not specify this parameter, the target destination is bound to a specific destination.

The publication topic

For publish/subscribe applications, use this page to specify the topic string of a trace-route message for the IBM MQ display route application to publish

-ts TopicName

Specifies a topic string to which the IBM MQ display route application is to publish a trace-route message, and puts this application into topic mode. In this mode, the application traces all of the messages that result from the publish request.

You can also use the IBM MQ display route application to display the results from an activity report that was generated for publish messages.

Message mimicking

Use this page to configure a trace-route message to mimic a message, for example when the original message did not reach its intended destination

One use of trace-route messaging is to help determine the last known location of a message that did not reach its intended destination. The IBM MQ display route application provides parameters that can help configure a trace-route message to mimic the original message. When mimicking a message, you can use the following parameters:

-l Persistence

Specifies the persistence of the generated trace-route message. Possible values for *Persistence* are:

yes

The generated trace-route message is persistent. (MQPER_PERSISTENT).

no

The generated trace-route message is **not** persistent. (MQPER_NOT_PERSISTENT).

q

The generated trace-route message inherits its persistence value from the destination specified by *-q TargetQName* or *-ts TargetTopicString*. (MQPER_PERSISTENCE_AS_Q_DEF).

A trace-route reply message, or any report messages, returned will share the same persistence value as the original trace-route message.

If *Persistence* is specified as **yes**, you must specify the parameter *-rq ReplyToQ*. The reply-to queue must not resolve to a temporary dynamic queue.

If you do not specify this parameter, the generated trace-route message is **not** persistent.

-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 destination specified by *-q TargetQName* or *-ts TargetTopicString*.

If you do not specify this parameter, the priority value is taken from the destination specified by *-q TargetQName* or *-ts TargetTopicString*.

-xs Expiry

Specifies the expiry time for the trace-route message, in seconds.

If you do not specify this parameter, the expiry time is specified as 60 seconds.

-ro none |ReportOption

none

Specifies no report options are set.

ReportOption

Specifies report options for the trace-route message. Multiple report options can be specified using a comma as a separator. Possible values for *ReportOption* are:

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.

If neither *-ro ReportOption* nor *-ro none* are specified, then the MQRO_ACTIVITY and MQRO_DISCARD_MSG report options are specified.

The IBM MQ display route application does not allow you to add user data to the trace-route message. If you require user data to be added to the trace-route message you must generate the trace-route message manually.

Recorded activity information

Use this page to specify the method used to return recorded activity information, which you can then use to determine the route that a trace-route message has taken

Recorded activity information can be returned as follows:

- In activity reports
- In a trace-route reply message
- In the trace-route message itself (having been put on the target queue)

When using **dspmqrte**, the method used to return recorded activity information is determined using the following parameters:

-ro activity

Specifies that activity information is returned using activity reports. By default activity recording is enabled.

-ac -ar

Specifies that activity information is accumulated in the trace-route message, and that a trace-route reply message is to be generated.

-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

Requests that a trace-route reply message containing all accumulated activity information is generated in the following circumstances:

- The trace-route message is discarded by an IBM MQ queue manager.
- The trace-route message is put to a local queue (target queue or dead-letter queue) by an IBM MQ queue manager.
- The number of activities performed on the trace-route message exceeds the value of specified in *-s Activities*.

-ac -d yes

Specifies that activity information is accumulated in the trace-route message, and that on arrival, the trace-route message will be put on the target queue.

-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.

-d yes

On arrival, the trace-route message is put to the target queue, even if the queue manager does not support trace-route messaging.

If you do not specify this parameter, the trace-route message is not put to the target queue.

The trace-route message can then be retrieved from the target queue, and the recorded activity information acquired.

You can combine these methods as required.

Additionally, the detail level of the recorded activity information can be specified using the following parameter:

-t Detail

Specifies the activities that are recorded. The possible values for *Detail* are:

low

Activities performed by user-defined application are recorded only.

medium

Activities specified in low are recorded. Additionally, publish activities and 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 could be included with this level of detail.

If you do not specify this parameter, medium level activities are recorded.

By default the IBM MQ display route application uses a temporary dynamic queue to store the returned messages. When the IBM MQ display route application ends, the temporary dynamic queue is closed, and any messages are purged. If the returned messages are required beyond the current execution of the IBM MQ display route application ends, then a permanent queue must be specified using the following parameters:

-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 then a dynamic reply-to queue is created using the system default model queue, SYSTEM.DEFAULT.MODEL.QUEUE.

-rqm ReplyToQMGr

Specifies the name of the queue manager where the reply-to queue resides. The name can contain up to 48 characters.

If you do not specify this parameter, the queue manager to which the IBM MQ display route application is connected is used as the reply-to queue manager.

How the trace-route message is handled

Use this page to control how a trace-route message is handled as it is routed through a queue manager network.

The following parameters can restrict where the trace-route message can be routed in the queue manager network:

-d Deliver

Specifies whether the trace-route message is to be delivered to the target queue on arrival. Possible values for *Deliver* are:

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.

If you do not specify this parameter, the trace-route message is not put to the target queue.

-f Forward

Specifies the type of queue manager that the trace-route message can be forwarded to. For details of the algorithm that queue managers use to determine whether to forward a message to a remote queue manager, refer to [“The TraceRoute PCF group” on page 96](#). The possible values for *Forward* are:

all

The trace-route message is forwarded to any queue manager.

Warning: If forwarded to a queue manager earlier than IBM WebSphere MQ 6.0, the trace-route message will not be recognized and can be delivered to a local queue despite the value of the *-d Deliver* parameter.

supported

The trace-route message is only forwarded to a queue manager that will honor the *Deliver* parameter from the *TraceRoute* PCF group

If you do not specify this parameter, the trace-route message will only be forwarded to a queue manager that will honor the *Deliver* parameter.

The following parameters can prevent a trace-route message from remaining in a queue manager network indefinitely:

-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 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.

-xs Expiry

Specifies the expiry time for the trace-route message, in seconds.

If you do not specify this parameter, the expiry time is specified as 60 seconds.

-xp PassExpiry

Specifies whether the expiry time from the trace-route message is passed on to a trace-route reply message. Possible values for *PassExpiry* are:

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 report option (if specified), and the remaining expiry time are passed on.

This is the default value.

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 expiry time from the trace-route message are **not** passed on.

If you do not specify this parameter, MQRO_PASS_DISCARD_AND_EXPIRY is not specified.

-ro discard

Specifies the MQRO_DISCARD_MSG report option. This can prevent the trace-route message remaining in the queue manager network indefinitely.

Display of activity information

The IBM MQ display route application can display activity information for a trace-route message that it has just put into a queue manager network, or it can display activity information for a previously generated trace-route message. It can also display additional information recorded by user-written applications.

To specify whether activity information returned for a trace-route message is displayed, specify the following parameter:

-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 from (*-ro ReportOption*), then a specific (non-model) reply-to queue must be specified using *-rq ReplyToQ* . By default, only activity report messages are requested.

After the trace-route message is put to the specified target queue, a 48 character hexadecimal string is displayed containing the message identifier of the trace-route message. The message identifier can be used by the IBM MQ display route application to display the activity information for the trace-route message at a later time, using the *-i CorrelId* parameter.

If you do not specify this parameter, activity information returned for the trace-route message is displayed in the form specified by the *-v* parameter.

When displaying activity information for a trace-route message that has just been put into a queue manager network, the following parameter can be specified:

-w WaitTime

Specifies the time, in seconds, that the IBM MQ display route application will wait for activity reports, or a trace-route reply message, to return to the specified reply-to queue.

If you do not specify this parameter, the wait time is specified as the expiry time of the trace-route message, plus 60 seconds.

When displaying previously accumulated activity information the following parameters must be set:

-q TargetQName

If the IBM MQ display route application is being used to view previously gathered activity information, *TargetQName* specifies the name of the queue where the activity information is stored.

-i CorrelId

This parameter is used when the IBM MQ display route application is used to display previously accumulated activity information only. There can be many activity reports and trace-route reply messages on the queue specified by *-q TargetQName*. *CorrelId* is used to identify the activity reports, or a trace-route reply message, related to a trace-route message. Specify the message identifier of the original trace-route message in *CorrelId*.

The format of *CorrelId* is a 48 character hexadecimal string.

The following parameters can be used when displaying previously accumulated activity information, or when displaying current activity information for a trace-route message:

-b

Specifies that the IBM MQ display route application will only browse activity reports or a trace-route reply message related to a message. This allows activity information to be displayed again at a later time.

If you do not specify this parameter, the IBM MQ display route application will destructively get activity reports or a trace-route reply message related to a message.

-v summary | all | none | outline *DisplayOption***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. Multiple display options can be specified using a comma as a separator.

If no values are supplied the following is displayed:

- The application name
- The type of each operation
- Any operation specific parameters

Possible values for *DisplayOption* are:

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 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.

tracerroute

All non-PCF group parameters in *TraceRoute* PCF groups are displayed.

If you do not specify this parameter, a summary of the message route is displayed.

Display of additional information

As a trace-route message is routed through a queue manager network, user-written applications can record additional information by writing one or more additional PCF parameters to the message data of the trace-route message or to the message data of an activity report. For the IBM MQ display route application to display additional information in a readable form it must be recorded in a specific format, as described in [“Additional activity information” on page 101](#).

IBM MQ display route application examples

The following examples show how you can use the IBM MQ display route application. In each example, two queue managers (QM1 and QM2) are inter-connected by two channels (QM2.TO.QM1 and QM1.TO.QM2).

Example 1 - Requesting activity reports

Display activity information from a trace-route message delivered to the target queue

In this example the IBM MQ display route application connects to queue manager, QM1, and is used to generate and deliver a trace-route message to the target queue, TARGET.Q, on remote queue manager, QM2. The necessary report option is specified so that activity reports are requested as the trace-route reply message is routed. On arrival at the target queue the trace-route message is discarded. Activity information returned to the IBM MQ display route application using activity reports is put in order and displayed.




Figure 13. Requesting activity reports, Diagram 1

- The ACTIVREC attribute of each queue manager (QM1 and QM2) is set to MSG.
- The following command is issued:

```
dspmqzte -m QM1 -q TARG.AT.QM2 -rq ACTIV.REPLY.Q
```

QM1 is the name of the queue manager to which the IBM MQ display route application connects, TARG.AT.QM2 is the name of the target queue, and ACTIV.REPLY.Q is the name of the queue to which it is requested that all responses to the trace-route message are sent.

Default values are assumed for all options that are not specified, but note in particular the -f option (the trace-route message is forwarded only to a queue manager that honors the Deliver parameter of the TraceRoute PCF group), the -d option (on arrival, the trace-route message is not put on the target queue), the -ro option (MQRO_ACTIVITY and MQRO_DISCARD_MSG report options are specified), and the -t option (medium detail level activity is recorded).

- DSPMQRTE generates the trace-route message and puts it on the remote queue TARG.AT.QM2.
- DSPMQRTE then looks at the value of the ACTIVREC attribute of queue manager QM1. The value is MSG, therefore DSPMQRTE generates an activity report and puts it on the reply queue ACTIV.REPLY.Q.




Figure 14. Requesting activity reports, Diagram 2

- The sending message channel agent (MCA) gets the trace-route message from the transmission queue. The message is a trace-route message, therefore the MCA begins to record the activity information.
- The ACTIVREC attribute of the queue manager (QM1) is MSG, and the MQRO_ACTIVITY option is specified in the Report field of the message descriptor, therefore the MCA will later generate an activity report. The RecordedActivities parameter value in the TraceRoute PCF group is incremented by 1.
- The MCA checks that the MaxActivities value in the TraceRoute PCF group has not been exceeded.
- Before the message is forwarded to QM2 the MCA follows the algorithm that is described in [Forwarding](#) (steps “1” on page 98, “4” on page 98, and “5” on page 98) and the MCA chooses to send the message.
- The MCA then generates an activity report and puts it on the reply queue (ACTIV.REPLY.Q).




Figure 15. Requesting activity reports, Diagram 3

- The receiving MCA receives the trace-route message from the channel. The message is a trace-route message, therefore the MCA begins to record the information about the activity.

- If the queue manager that the trace-route message has come from is IBM WebSphere MQ 5.3.1 or earlier, the MCA increments the DiscontinuityCount parameter of the TraceRoute PCF by 1. This is not the case here.
- The ACTIVREC attribute of the queue manager (QM2) is MSG, and the MQRO_ACTIVITY option is specified, therefore the MCA will generate an activity report. The RecordedActivities parameter value is incremented by 1.
- The target queue is a local queue, therefore the message is discarded with feedback MQFB_NOT_DELIVERED, in accordance with the Deliver parameter value in the TraceRoute PCF group.
- The MCA then generates the final activity report and puts it on the reply queue. This resolves to the transmission queue that is associated with queue manager QM1 and the activity report is returned to queue manager QM1 (ACTIV.REPLY.Q).




Figure 16. Requesting activity reports, Diagram 4

- Meanwhile, DSPMQRTE has been continually performing MQGETs on the reply queue (ACTIV.REPLY.Q), waiting for activity reports. It will wait for up to 120 seconds (60 seconds longer than the expiry time of the trace-route message) since -w was not specified when DSPMQRTE was started.
- DSPMQRTE gets the 3 activity reports off the reply queue.
- The activity reports are ordered using the RecordedActivities, UnrecordedActivities, and DiscontinuityCount parameters in the TraceRoute PCF group for each of the activities. The only value that is non-zero in this example is RecordedActivities, therefore this is the only parameter that is actually used.
- The program ends as soon as the discard operation is displayed. Even though the final operation was a discard, it is treated as though a put took place because the feedback is MQFB_NOT_DELIVERED.

The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM1 -q TARG.AT.QM2
-iq ACTIV.REPLY.Q'.
AMQ8659: DSPMQRTE command successfully put a message on queue 'QM2',
queue manager 'QM1'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
AMQ8666: Queue 'QM2' on queue manager 'QM1'.
AMQ8666: Queue 'TARGET.Q' on queue manager 'QM2'.
AMQ8652: DSPMQRTE command has finished.
```

Example 2 - Requesting a trace-route reply message

Generate and deliver a trace-route message to the target queue

In this example the IBM MQ display route application connects to queue manager, QM1, and is used to generate and deliver a trace-route message to the target queue, TARGET.Q, on remote queue manager,

QM2. The necessary option is specified so that activity information is accumulated in the trace-route message. On arrival at the target queue a trace-route reply message is requested, and the trace-route message is discarded.




Figure 17. Requesting a trace-route reply message, Diagram 1

- The ROUTEREC attribute of each queue manager (QM1 and QM2) is set to MSG.
- The following command is issued:

```
dspmqrte -m QM1 -q TARG.AT.QM2 -rq TR.REPLY.Q -ac -ar -ro discard
```

QM1 is the name of the queue manager to which the IBM MQ display route application connects, TARG.AT.QM2 is the name of the target queue, and ACTIV.REPLY.Q is the name of the queue to which it is requested that all responses to the trace-route message are sent. The -ac option specifies that activity information is accumulated in the trace-route message, the -ar option specifies that all accumulated activity is sent to the reply-to queue that is specified by the -rq option (that is, TR.REPLY.Q). The -ro option specifies that report option MQRO_DISCARD_MSG is set which means that activity reports are not generated in this example.

- DSPMQRTE accumulates activity information in the trace-route message before the message is put on the target route. The queue manager attribute ROUTEREC must not be DISABLED for this to happen.




Figure 18. Requesting a trace-route reply message, Diagram 2

- The message is a trace-route message, therefore the sending MCA begins to record information about the activity.
- The queue manager attribute ROUTEREC on QM1 is not DISABLED, therefore the MCA accumulates the activity information within the message, before the message is forwarded to queue manager QM2.




Figure 19. Requesting a trace-route reply message, Diagram 3

- The message is a trace-route message, therefore the receiving MCA begins to record information about the activity.
- The queue manager attribute ROUTEREC on QM2 is not DISABLED, therefore the MCA accumulates the information within the message.
- The target queue is a local queue, therefore the message is discarded with feedback MQFB_NOT_DELIVERED, in accordance with the Deliver parameter value in the TraceRoute PCF group.

- This is the last activity that will take place on the message, and because the queue manager attribute ROUTEREC on QM1 is not DISABLED, the MCA generates a trace-route reply message in accordance with the Accumulate value. The value of ROUTEREC is MSG, therefore the reply message is put on the reply queue. The reply message contains all the accumulated activity information from the trace-route message.




Figure 20. Requesting a trace-route reply message, Diagram 4

- Meanwhile DSPMQRTE is waiting for the trace-route reply message to return to the reply queue. When it returns, DSPMQRTE parses each activity that it contains and prints it out. The final operation is a discard operation. DSPMQRTE ends after it has been printed.

The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM1 -q TARG.AT.QM2 -rq
TR.REPLY.Q'.
AMQ8659: DSPMQRTE command successfully put a message on queue 'QM2', queue
manager 'QM1'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
AMQ8666: Queue 'QM2' on queue manager 'QM1'.
AMQ8666: Queue 'TARGET.Q' on queue manager 'QM2'.
AMQ8652: DSPMQRTE command has finished.
```

Example 3 - Delivering activity reports to the system queue

Detect when activity reports are delivered to queues other than the reply-to queue and use the IBM MQ display route application to read activity reports from the other queue.

This example is the same as “Example 1 - Requesting activity reports” on page 111, except that QM2 now has the value of the ACTIVREC queue manage attribute set to QUEUE. Channel QM1.TO.QM2 must have been restarted for this to take effect.

This example demonstrates how to detect when activity reports are delivered to queues other than the reply-to queue. Once detected, the IBM MQ display route application is used to read activity reports from another queue.




Figure 21. Delivering activity reports to the system queue, Diagram 1

- The message is a trace-route message, therefore the receiving MCA begins to record information about the activity.
- The value of the ACTIVREC queue manager attribute on QM2 is now QUEUE, therefore the MCA generates an activity report, but puts it on the system queue (SYSTEM.ADMIN.ACTIVITY.QUEUE) and not on the reply queue (ACTIV.REPLY.Q).




Figure 22. Delivering activity reports to the system queue, Diagram 2

- Meanwhile DSPMQRTE has been waiting for activity reports to arrive on ACTIV.REPLY.Q. Only two arrive. DSPMQRTE continues waiting for 120 seconds because it seems that the route is not yet complete.

The output that is displayed follows:

```

AMQ8653: DSPMQRTE command started with options '-m QM1 -q TARG.AT.QM2 -rq
        ACTIV.REPLY.Q -v outline identifiers'.
AMQ8659: DSPMQRTE command successfully put a message on queue 'QM2', queue
        manager 'QM1'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
-----

```

```

Activity:
  ApplName: 'cann\output\bin\dspmqrte.exe'

```

```

Operation:
  OperationType: Put

```

```

Message:

```

```

MQMD:
  MsgId: X'414D51204C41524745512020202020A3C9154220001502'
  CorrelId: X'414D51204C41524745512020202020A3C9154220001503'
  QMgrName: 'QM1'
  QName: 'TARG.AT.QM2'
  ResolvedQName: 'QM2'
  RemoteQName: 'TARGET.Q'
  RemoteQMGrName: 'QM2'
-----

```

```

Activity:
  ApplName: 'cann\output\bin\runmqchl.EXE'

```

```

Operation:
  OperationType: Get

```

```

Message:

```

```

MQMD:
  MsgId: X'414D51204C41524745512020202020A3C9154220001505'
  CorrelId: X'414D51204C41524745512020202020A3C9154220001502'

```

```

EmbeddedMQMD:
  MsgId: X'414D51204C41524745512020202020A3C9154220001502'
  CorrelId: X'414D51204C41524745512020202020A3C9154220001503'

```

```

QMGrName: 'QM1'
QName: 'QM2'
ResolvedQName: 'QM2'

```

```

Operation:
  OperationType: Send

```

```

Message:

```

```

MQMD:
  MsgId: X'414D51204C41524745512020202020A3C9154220001502'
  CorrelId: X'414D51204C41524745512020202020A3C9154220001503'
  QMgrName: 'QM1'
  RemoteQMGrName: 'QM2'
  ChannelName: 'QM1.TO.QM2'
  ChannelType: Sender
  XmitQName: 'QM2'
-----

```

```

AMQ8652: DSPMQRTE command has finished.

```

- The last operation that DSPMQRTE observed was a Send, therefore the channel is running. Now we must work out why we did not receive any more activity reports from queue manager QM2 (as identified in RemoteQMGrName).
- To check whether there is any activity information on the system queue, start DSPMQRTE on QM2 to try and collect more activity reports. Use the following command to start DSPMQRTE:

```

dspmqrte -m QM2 -q SYSTEM.ADMIN.ACTIVITY.QUEUE
        -i 414D51204C41524745512020202020A3C9154220001502 -v outline

```

where 414D51204C41524745512020202020A3C9154220001502 is the MsgId of the trace-route message that was put.

- DSPMQRTE then performs a sequence of MQGETs again, waiting for responses on the system activity queue related to the trace-route message with the specified identifier.
- DSPMQRTE gets one more activity report, which it displays. DSPMQRTE determines that the preceding activity reports are missing, and displays a message saying this. We already know about this part of the route, however.

The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM2
        -q SYSTEM.ADMIN.ACTIVITY.QUEUE
        -i 414D51204C41524745512020202020A3C915420001502 -v outline'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
```

```
-----
Activity:
Activity information unavailable.
```

```
-----
Activity:
ApplName: 'cann\output\bin\AMQRMPPA.EXE'
```

```
Operation:
OperationType: Receive
QMgrName: 'QM2
RemoteQMgrName: 'QM1
ChannelName: 'QM1.TO.QM2
ChannelType: Receiver
```

```
Operation:
OperationType: Discard
QMgrName: 'QM2
QName: 'TARGET.Q
Feedback: NotDelivered
```

```
-----
AMQ8652: DSPMQRTE command has finished.
```

- This activity report indicates that the route information is now complete. No problem occurred.
- Just because route information is unavailable, or because DSPMQRTE cannot display all of the route, this does not mean that the message was not delivered. For example, the queue manager attributes of different queue managers might be different, or a reply queue might not be defined to get the response back.

Example 4 - Diagnosing a channel problem

Diagnose a problem in which the trace-route message does not reach the target queue

In this example the IBM MQ display route application connects to queue manager, QM1, generates a trace-route message, then attempts to deliver it to the target queue, TARGET.Q, on remote queue manager, QM2. In this example the trace-route message does not reach the target queue. The available activity report is used to diagnose the problem.




Figure 23. Diagnosing a channel problem

- In this example, the channel QM1.TO.QM2 is not running.
- DSPMQRTE puts a trace-route message (as in example 1) to the target queue and generates an activity report.
- There is no MCA to get the message from the transmission queue (QM2), therefore this is the only activity report that DSPMQRTE gets back from the reply queue. This time the fact that the route is not complete does indicate a problem. The administrator can use the transmission queue found in ResolvedQName to investigate why the transmission queue is not being serviced.

The output that is displayed follows:

```
AMQ8653: DSPMQRTE command started with options '-m QM1 -q TARG.AT.QM2
        -rq ACTIV.REPLY.Q -v outline'.
AMQ8659: DSPMQRTE command successfully put a message on queue 'QM2',
        queue manager 'QM1'.
AMQ8674: DSPMQRTE command is now waiting for information to display.
-----
Activity:
  ApplName: 'cann\output\bin\dspmqrte.exe'

  Operation:
    OperationType: Put
    QMgrName: 'QM1'
    QName: 'TARG.AT.QM2'
    ResolvedQName: 'QM2'
    RemoteQName: 'TARGET.Q'
    RemoteQMgrName: 'QM2'
-----
AMQ8652: DSPMQRTE command has finished.
```

Activity report reference

Use this page to obtain an overview of the activity report message format. The activity report message data contains the parameters that describe the activity.

Activity report format

Activity reports are standard IBM MQ report messages containing a message descriptor and message data. Activity reports are PCF messages generated by applications that have performed an activity on behalf of a message as it has been routed through a queue manager network.

Activity reports contain the following information:

A message descriptor

An MQMD structure

Message data

Consists of the following:

- An embedded PCF header (MQEPH).
- Activity report message data.

Activity report message data consists of the *Activity* PCF group and, if generated for a trace-route message, the *TraceRoute* PCF group.

[Table 23 on page 121](#) shows the structure of these reports, including parameters that are returned only under certain conditions.

Table 23. Activity report format

MQMD structure	Embedded PCF header MQEPH structure	Activity report message data
Structure identifier Structure version Report options Message type Expiration time Feedback Encoding Coded character set ID Message format Priority Persistence Message identifier Correlation identifier Backout count Reply-to queue Reply-to queue manager User identifier Accounting token Application identity data Application type Application name Put date Put time Application origin data Group identifier Message sequence number Offset Message flags Original length	Structure identifier Structure version Structure length Encoding Coded character set ID Message format Flags PCF header (MQCFH) Structure type Structure length Structure version Command identifier Message sequence number Control options Completion code Reason code Parameter count	Activity Activity application name Activity application type Activity description Operation Operation type Operation date Operation time Message Message length MQMD ⁸ EmbeddedMQMD Queue manager name Queue sharing group name Queue name ¹ ² ³ ⁷ Resolved queue name ¹ ³ ⁷ Remote queue name ³ ⁷ Remote queue manager name ² ³ ⁴ ⁵ ⁷ Subscription level ⁹ Subscription identifier ⁹ Feedback ² ¹⁰ Channel name ⁴ ⁵ Channel type ⁴ ⁵ Transmission queue name ⁵ TraceRoute ⁶ Detail Recorded activities Unrecorded activities Discontinuity count Max activities Accumulate Deliver

Notes:

1. Returned for Get and Browse operations.
2. Returned for Discard operations.
3. Returned for Put, Put Reply, and Put Report operations.
4. Returned for Receive operations.

5. Returned for Send operations.
6. Returned for trace-route messages.
7. Not returned for Put operations to a topic, contained within Publish activities.
8. Not returned for Excluded Publish operations. For Publish and Discarded Publish operations, returned containing a subset of parameters.
9. Returned for Publish, Discarded Publish, and Excluded Publish operations.
10. Returned for Discarded Publish and Excluded Publish operations.

Activity report MQMD (message descriptor)

Use this page to view the values contained by the MQMD structure for an activity report

StrucId

Structure identifier:

Data type

MQCHAR4

Value

MQMD_STRUC_ID.

Version

Structure version number

Data type

MLONG

Values

Copied from the original message descriptor. Possible values are:

MQMD_VERSION_1

Version-1 message descriptor structure, supported in all environments.

MQMD_VERSION_2

Version-2 message descriptor structure, supported in the following environments:

-  AIX
-  IBM i
-  Linux
-  Windows
-  z/OS

and all IBM MQ MQI clients connected to these systems.

Report

Options for further report messages

Data type

MLONG

Value

If MQRO_PASS_DISCARD_AND_EXPIRY or MQRO_DISCARD_MSG were specified in the *Report* field of the original message descriptor:

MQRO_DISCARD

The report is discarded if it cannot be delivered to the destination queue.

Otherwise:

MQRO_NONE

No reports required.

MsgType

Indicates type of message

Data type

MLONG

Value

MQMT_REPORT

Expiry

Report message lifetime

Data type

MLONG

Value

If the *Report* field in the original message descriptor is specified as MQRO_PASS_DISCARD_AND_EXPIRY, the remaining expiry time from the original message is used.

Otherwise:

MQEI_UNLIMITED

The report does not have an expiry time.

Feedback

Description: Feedback or reason code.

Data type: MLONG.

Value: **MQFB_ACTIVITY**
Activity report.

Encoding

Description: Numeric encoding of report message data.

Data type: MLONG.

Value: MQENC_NATIVE.

CodedCharSetId

Description: Character set identifier of report message data.

Data type: MLONG.

Value: Set as appropriate.

Format

Description: Format name of report message data

Data type: MQCHAR8.

Value: **MQFMT_EMBEDDED_PCF**
Embedded PCF message.

Priority

Description: Report message priority.

Data type: MLONG.

Value: Copied from the original message descriptor.

Persistence

Description: Report message persistence.
Data type: MQLONG.
Value: Copied from the original message descriptor.

MsgId

Description: Message identifier.
Data type: MQBYTE24.
Values: If the *Report* field in the original message descriptor is specified as MQRO_PASS_MSG_ID, the message identifier from the original message is used. Otherwise, a unique value will be generated by the queue manager.

CorrelId

Description: Correlation identifier.
Data type: MQBYTE24.
Value: If the *Report* field in the original message descriptor is specified as MQRO_PASS_CORREL_ID, the correlation identifier from the original message is used. Otherwise, the message identifier is copied from the original message.

BackoutCount

Description: Backout counter.
Data type: MQLONG.
Value: 0.

ReplyToQ

Description: Name of reply queue.
Data type: MQCHAR48.
Values: Blank.

ReplyToQMgr

Description: Name of reply queue manager.
Data type: MQCHAR48.
Value: The queue manager name that generated the report message.

UserIdentifier

Description: The user identifier of the application that generated the report message.
Data type: MQCHAR12.
Value: Copied from the original message descriptor.

AccountingToken

Description: Accounting token that allows an application to charge for work done as a result of the message.

Data type: MQBYTE32.
Value: Copied from the original message descriptor.

ApplIdentityData

Description: Application data relating to identity.
Data type: MQCHAR32.
Values: Copied from the original message descriptor.

PutApplType

Description: Type of application that put the report message.
Data type: MQLONG.
Value: **MQAT_QMGR**
Queue manager generated message.

PutApplName

Description: Name of application that put the report message.
Data type: MQCHAR28.
Value: Either the first 28 bytes of the queue manager name, or the name of the MCA that generated the report message.

PutDate

Description: Date when message was put.
Data type: MQCHAR8.
Value: As generated by the queue manager.

PutTime

Description: Time when message was put.
Data type: MQCHAR8.
Value: As generated by the queue manager.

ApplOriginData

Description: Application data relating to origin.
Data type: MQCHAR4.
Value: Blank.

If *Version* is MQMD_VERSION_2, the following additional fields are present:

GroupId

Description: Identifies to which message group or logical message the physical message belongs.
Data type: MQBYTE24.
Value: Copied from the original message descriptor.

MsgSeqNumber

Description:	Sequence number of logical message within group.
Data type:	MLONG.
Value:	Copied from the original message descriptor.

Offset

Description:	Offset of data in physical message from start of logical message.
Data type:	MLONG.
Value:	Copied from the original message descriptor.

MsgFlags

Description:	Message flags that specify attributes of the message or control its processing.
Data type:	MLONG.
Value:	Copied from the original message descriptor.

OriginalLength

Description:	Length of original message.
Data type:	MLONG.
Value:	Copied from the original message descriptor.

Activity report MQEPH (Embedded PCF header)

Use this page to view the values contained by the MQEPH structure for an activity report

The MQEPH structure contains a description of both the PCF information that accompanies the message data of an activity report, and the application message data that follows it.

For an activity report, the MQEPH structure contains the following values:

StrucId

Description:	Structure identifier.
Data type:	MQCHAR4.
Value:	MQEPH_STRUC_ID.

Version

Description:	Structure version number.
Data type:	MLONG.
Values:	MQEPH_VERSION_1.

StrucLength

Description:	Structure length.
Data type:	MLONG.
Value:	Total length of the structure including the PCF parameter structures that follow it.

Encoding

Description:	Numeric encoding of the message data that follows the last PCF parameter structure.
Data type:	MLONG.
Value:	If any data from the original application message data is included in the report message, the value will be copied from the <i>Encoding</i> field of the original message descriptor. Otherwise, 0.

CodedCharSetId

Description:	Character set identifier of the message data that follows the last PCF parameter structure.
Data type:	MLONG.
Value:	If any data from the original application message data is included in the report message, the value will be copied from the <i>CodedCharSetId</i> field of the original message descriptor. Otherwise, MQCCSI_UNDEFINED.

Format

Description:	Format name of message data that follows the last PCF parameter structure.
Data type:	MQCHAR8.
Value:	If any data from the original application message data is included in the report message, the value will be copied from the <i>Format</i> field of the original message descriptor. Otherwise, MQFMT_NONE.

Flags

Description:	Flags that specify attributes of the structure or control its processing.
Data type:	MLONG.
Value:	MQEPH_CCSID_EMBEDDED Specifies that the character set of the parameters containing character data is specified individually within the <i>CodedCharSetId</i> field in each structure.

PCFHeader

Description:	Programmable Command Format Header
Data type:	MQCFH.
Value:	See “Activity report MQCFH (PCF header)” on page 127 .

Activity report MQCFH (PCF header)

Use this page to view the PCF values contained by the MQCFH structure for an activity report
For an activity report, the MQCFH structure contains the following values:

Type

Description:	Structure type that identifies the content of the report message.
Data type:	MLONG.

Value: **MQCFT_REPORT**
Message is a report.

StrucLength

Description: Structure length.

Data type: MQLONG.

Value: **MQCFH_STRUC_LENGTH**
Length in bytes of MQCFH structure.

Version

Description: Structure version number.

Data type: MQLONG.

Values: MQCFH_VERSION_3

Command

Description: Command identifier. This identifies the category of the message.

Data type: MQLONG.

Values: **MQCMD_ACTIVITY_MSG**
Message activity.

MsgSeqNumber

Description: Message sequence number. This is the sequence number of the message within a group of related messages.

Data type: MQLONG.

Values: 1.

Control

Description: Control options.

Data type: MQLONG.

Values: MQCFC_LAST.

CompCode

Description: Completion code.

Data type: MQLONG.

Values: MQCC_OK.

Reason

Description: Reason code qualifying completion code.

Data type: MQLONG.

Values: MQRC_NONE.

ParameterCount

Description:	Count of parameter structures. This is the number of parameter structures that follow the MQCFH structure. A group structure (MQCFGR), and its included parameter structures, are counted as one structure only.
Data type:	MLONG.
Values:	1 or greater.

Activity report message data

Use this page to view the parameters contained by the *Activity* PCF group in an activity report message. Some parameters are returned only when specific operations have been performed.

Activity report message data consists of the *Activity* PCF group and, if generated for a trace-route message, the *TraceRoute* PCF group. The *Activity* PCF group is detailed in this topic.

Some parameters, which are described as Operation-specific activity report message data, are returned only when specific operations have been performed.

For an activity report, the activity report message data contains the following parameters:

Activity

Description:	Grouped parameters describing the activity.
Identifier:	MQGACF_ACTIVITY.
Data type:	MQCFGR.
Included in PCF group:	None.
Parameters in PCF group:	<i>ActivityApplName</i> <i>ActivityApplType</i> <i>ActivityDescription</i> <i>Operation</i> <i>TraceRoute</i>
Returned:	Always.

ActivityApplName

Description:	Name of application that performed the activity.
Identifier:	MQCACF_APPL_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Activity</i> .
Maximum length:	MQ_APPL_NAME_LENGTH.
Returned:	Always.

ActivityApplType

Description:	Type of application that performed the activity.
Identifier:	MQIA_APPL_TYPE.
Data type:	MQCFIN.
Included in PCF group:	<i>Activity</i> .

Returned: Always.

ActivityDescription

Description: Description of activity performed by the application.

Identifier: MQCACF_ACTIVITY_DESCRIPTION.

Data type: MQCFST.

Included in PCF group: *Activity*.

Maximum length: 64

Returned: Always.

Operation

Description: Grouped parameters describing an operation of the activity.

Identifier: MQGACF_OPERATION.

Data type: MQCFGR.

Included in PCF group: *Activity*.

Parameters in PCF group: *OperationType*
OperationDate
OperationTime
Message
QMgrName
QSGName

Note: Additional parameters are returned in this group depending on the operation type. These additional parameters are described as [Operation-specific activity report message data](#).

Returned: One *Operation* PCF group per operation in the activity.

OperationType

Description: Type of operation performed.

Identifier: MQIACF_OPERATION_TYPE.

Data type: MQCFIN.

Included in PCF group: *Operation*.

Values: MQOPER_*

Returned: Always.

OperationDate

Description: Date when the operation was performed.

Identifier: MQCACF_OPERATION_DATE.

Data type: MQCFST.

Included in PCF group: *Operation*.

Maximum length: MQ_DATE_LENGTH.
Returned: Always.

OperationTime

Description: Time when the operation was performed.
Identifier: MQCACF_OPERATION_TIME.
Data type: MQCFST.
Included in PCF group: *Operation*.
Maximum length: MQ_TIME_LENGTH.
Returned: Always.

Message

Description: Grouped parameters describing the message that caused the activity.
Identifier: MQGACF_MESSAGE.
Data type: MQCFGR.
Included in PCF group: *Operation*.
Parameters in group: *MsgLength*
MQMD
EmbeddedMQMD
Returned: Always, except for Excluded Publish operations.

MsgLength

Description: Length of the message that caused the activity, before the activity occurred.
Identifier: MQIACF_MSG_LENGTH.
Data type: MQCFIN.
Included in PCF group: *Message*.
Returned: Always.

MQMD

Description: Grouped parameters related to the message descriptor of the message that caused the activity.
Identifier: MQGACF_MQMD.
Data type: MQCFGR.
Included in PCF group: *Message*.

Parameters in group:	<i>StrucId</i> <i>Version</i> <i>Report</i> <i>MsgType</i> <i>Expiry</i> <i>Feedback</i> <i>Encoding</i> <i>CodedCharSetId</i> <i>Format</i> <i>Priority</i> <i>Persistence</i> <i>MsgId</i> <i>CorrelId</i> <i>BackoutCount</i> <i>ReplyToQ</i> <i>ReplyToQMgr</i> <i>UserIdentifier</i> <i>AccountingToken</i> <i>ApplIdentityData</i> <i>PutApplType</i> <i>PutApplName</i> <i>PutDate</i> <i>PutTime</i> <i>ApplOriginData</i> <i>GroupId</i> <i>MsgSeqNumber</i> <i>Offset</i> <i>MsgFlags</i> <i>OriginalLength</i>
----------------------	---

Returned:	Always, except for Excluded Publish operations.
-----------	---

EmbeddedMQMD

Description:	Grouped parameters describing the message descriptor embedded within a message on a transmission queue.
Identifier:	MQGACF_EMBEDDED_MQMD.
Data type:	MQCFGR.
Included in PCF group:	<i>Message</i> .

Parameters in group:	<i>StrucId</i> <i>Version</i> <i>Report</i> <i>MsgType</i> <i>Expiry</i> <i>Feedback</i> <i>Encoding</i> <i>CodedCharSetId</i> <i>Format</i> <i>Priority</i> <i>Persistence</i> <i>MsgId</i> <i>CorrelId</i> <i>BackoutCount</i> <i>ReplyToQ</i> <i>ReplyToQMGr</i> <i>UserIdentifier</i> <i>AccountingToken</i> <i>ApplIdentityData</i> <i>PutApplType</i> <i>PutApplName</i> <i>PutDate</i> <i>PutTime</i> <i>ApplOriginData</i> <i>GroupId</i> <i>MsgSeqNumber</i> <i>Offset</i> <i>MsgFlags</i> <i>OriginalLength</i>
----------------------	---

Returned:	For Get operations where the queue resolves to a transmission queue.
-----------	--

StrucId

Description:	Structure identifier
Identifier:	MQCACF_STRUC_ID.
Data type:	MQCFST.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Maximum length:	4.
Returned:	Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

Version

Description:	Structure version number.
Identifier:	MQIACF_VERSION.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .

Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

Report

Description: Options for report messages.

Identifier: MQIACF_REPORT.

Data type: MQCFIN.

Included in PCF group: *MQMD* or *EmbeddedMQMD*.

Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

MsgType

Description: Indicates type of message.

Identifier: MQIACF_MSG_TYPE.

Data type: MQCFIN.

Included in PCF group: *MQMD* or *EmbeddedMQMD*.

Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

Expiry

Description: Message lifetime.

Identifier: MQIACF_EXPIRY.

Data type: MQCFIN.

Included in PCF group: *MQMD* or *EmbeddedMQMD*.

Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

Feedback

Description: Feedback or reason code.

Identifier: MQIACF_FEEDBACK.

Data type: MQCFIN.

Included in PCF group: *MQMD* or *EmbeddedMQMD*.

Returned: Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

Encoding

Description: Numeric encoding of message data.

Identifier: MQIACF_ENCODING.

Data type: MQCFIN.

Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Returned:	Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

CodedCharSetId

Description:	Character set identifier of message data.
Identifier:	MQIA_CODED_CHAR_SET_ID.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Returned:	Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

Format

Description:	Format name of message data
Identifier:	MQCACH_FORMAT_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Maximum length:	MQ_FORMAT_LENGTH.
Returned:	Always, except for Excluded Publish operations.

Priority

Description:	Message priority.
Identifier:	MQIACF_PRIORITY.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Returned:	Always, except for Excluded Publish operations.

Persistence

Description:	Message persistence.
Identifier:	MQIACF_PERSISTENCE.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Returned:	Always, except for Excluded Publish operations.

MsgId

Description:	Message identifier.
Identifier:	MQBACF_MSG_ID.
Data type:	MQCFBS.

Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Maximum length:	MQ_MSG_ID_LENGTH.
Returned:	Always, except for Excluded Publish operations.

CorrelId

Description:	Correlation identifier.
Identifier:	MQBACF_CORREL_ID.
Data type:	MQCFBS.
Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Maximum length:	MQ_CORREL_ID_LENGTH.
Returned:	Always, except for Excluded Publish operations.

BackoutCount

Description:	Backout counter.
Identifier:	MQIACF_BACKOUT_COUNT.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Returned:	Always, except for Excluded Publish operations and in MQMD for Publish and Discarded Publish operations.

ReplyToQ

Description:	Name of reply queue.
Identifier:	MQCACF_REPLY_TO_Q.
Data type:	MQCFST.
Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Maximum length:	MQ_Q_NAME_LENGTH.
Returned:	Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish operations.

ReplyToQMgr

Description:	Name of reply queue manager.
Identifier:	MQCACF_REPLY_TO_Q_MGR.
Data type:	MQCFST.
Included in PCF group:	<i>MQMD or EmbeddedMQMD.</i>
Maximum length:	MQ_Q_MGR_NAME_LENGTH.
Returned:	Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

UserIdentifier

Description:	The user identifier of the application that originated the message.
Identifier:	MQCACF_USER_IDENTIFIER.
Data type:	MQCFST.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Maximum length:	MQ_USER_ID_LENGTH.
Returned:	Always, except for Excluded Publish Operations.

AccountingToken

Description:	Accounting token that allows an application to charge for work done as a result of the message.
Identifier:	MQBACF_ACCOUNTING_TOKEN.
Data type:	MQCFBS.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Maximum length:	MQ_ACCOUNTING_TOKEN_LENGTH.
Returned:	Always, except for Excluded Publish Operations.

ApplIdentityData

Description:	Application data relating to identity.
Identifier:	MQCACF_APPL_IDENTITY_DATA.
Data type:	MQCFST.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Maximum length:	MQ_APPL_IDENTITY_DATA_LENGTH.
Returned:	Always, except for Excluded Publish Operations.

PutApplType

Description:	Type of application that put the message.
Identifier:	MQIA_APPL_TYPE.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Returned:	Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

PutApplName

Description:	Name of application that put the message.
Identifier:	MQCACF_APPL_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .

Maximum length: MQ_APPL_NAME_LENGTH.
 Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

PutDate

Description: Date when message was put.
 Identifier: MQCACF_PUT_DATE.
 Data type: MQCFST.
 Included in PCF group: *MQMD or EmbeddedMQMD*.
 Maximum length: MQ_PUT_DATE_LENGTH.
 Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

PutTime

Description: Time when message was put.
 Identifier: MQCACF_PUT_TIME.
 Data type: MQCFST.
 Included in PCF group: *MQMD or EmbeddedMQMD*.
 Maximum length: MQ_PUT_TIME_LENGTH.
 Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

ApplOriginData

Description: Application data relating to origin.
 Identifier: MQCACF_APPL_ORIGIN_DATA.
 Data type: MQCFST.
 Included in PCF group: *MQMD or EmbeddedMQMD*.
 Maximum length: MQ_APPL_ORIGIN_DATA_LENGTH.
 Returned: Always, except for Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

GroupId

Description: Identifies to which message group or logical message the physical message belongs.
 Identifier: MQBACF_GROUP_ID.
 Data type: MQCFBS.
 Included in PCF group: *MQMD or EmbeddedMQMD*.
 Maximum length: MQ_GROUP_ID_LENGTH.
 Returned: If the *Version* is specified as MQMD_VERSION_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

MsgSeqNumber

Description:	Sequence number of logical message within group.
Identifier:	MQIACH_MSG_SEQUENCE_NUMBER.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Returned:	If <i>Version</i> is specified as MQMD_VERSION_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

Offset

Description:	Offset of data in physical message from start of logical message.
Identifier:	MQIACF_OFFSET.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Returned:	If <i>Version</i> is specified as MQMD_VERSION_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

MsgFlags

Description:	Message flags that specify attributes of the message or control its processing.
Identifier:	MQIACF_MSG_FLAGS.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Returned:	If <i>Version</i> is specified as MQMD_VERSION_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

OriginalLength

Description:	Length of original message.
Identifier:	MQIACF_ORIGINAL_LENGTH.
Data type:	MQCFIN.
Included in PCF group:	<i>MQMD</i> or <i>EmbeddedMQMD</i> .
Returned:	If <i>Version</i> is specified as MQMD_VERSION_2. Not returned in Excluded Publish Operations and in MQMD for Publish and Discarded Publish Operations.

QMGrName

Description:	Name of the queue manager where the activity was performed.
Identifier:	MQCA_Q_MGR_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation</i> .
Maximum length:	MQ_Q_MGR_NAME_LENGTH

Returned: Always.

QSGName

Description: Name of the queue sharing group to which the queue manager where the activity was performed belongs.

Identifier: MQCA_QSG_NAME.

Data type: MQCFST.

Included in PCF group: *Operation.*

Maximum length: MQ_QSG_NAME_LENGTH

Returned: If the activity was performed on an IBM MQ for z/OS queue manager.

TraceRoute

Description: Grouped parameters specifying attributes of the trace-route message.

Identifier: MQGACF_TRACE_ROUTE.

Data type: MQCFGR.

Contained in PCF group: *Activity.*

Parameters in group: *Detail*
RecordedActivities
UnrecordedActivities
DiscontinuityCount
MaxActivities
Accumulate
Forward
Deliver

Returned: If the activity was performed on behalf of the trace-route message.

The values of the parameters in the *TraceRoute* PCF group are those from the trace-route message at the time the activity report was generated.

Operation-specific activity report message data

Use this page to view the additional PCF parameters that might be returned in the PCF group *Operation* in an activity report, depending on the value of the *OperationType* parameter

The additional parameters vary depending on the following operation types:

Get/Browse (MQOPER_GET/MQOPER_BROWSE)

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Get/Browse (MQOPER_GET/MQOPER_BROWSE) operation type (a message on a queue was got, or browsed).

QName

Description: The name of the queue that was opened.

Identifier: MQCA_Q_NAME.

Data type: MQCFST.

Included in PCF group: *Operation.*

Maximum length: MQ_Q_NAME_LENGTH
Returned: Always.

ResolvedQName

Description: The name that the opened queue resolves to.
Identifier: MQCACF_RESOLVED_Q_NAME.
Data type: MQCFST.
Included in PCF group: *Operation*.
Maximum length: MQ_Q_NAME_LENGTH
Returned: Always.

Discard (MQOPER_DISCARD)

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Discard (MQOPER_DISCARD) operation type (a message was discarded).

Feedback

Description: The reason for the message being discarded.
Identifier: MQIACF_FEEDBACK.
Data type: MQCFIN.
Included in PCF group: *Operation*.
Returned: Always.

QName

Description: The name of the queue that was opened.
Identifier: MQCA_Q_NAME.
Data type: MQCFST.
Maximum length: MQ_Q_NAME_LENGTH
Included in PCF group: *Operation*.
Returned: If the message was discarded because it was unsuccessfully put to a queue.

RemoteQMgrName

Description: The name of the queue manager to which the message was destined.
Identifier: MQCA_REMOTE_Q_MGR_NAME.
Data type: MQCFST.
Maximum length: MQ_Q_MGR_NAME_LENGTH
Included in PCF group: *Operation*.
Returned: If the value of *Feedback* is MQFB_NOT_FORWARDED.

Publish/Discarded Publish/Excluded Publish (MQOPER_PUBLISH/MQOPER_DISCARDED_PUBLISH/MQOPER_EXCLUDED_PUBLISH)

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Publish/Discarded Publish/Excluded Publish (MQOPER_PUBLISH/MQOPER_DISCARDED_PUBLISH/MQOPER_EXCLUDED_PUBLISH) operation type (a publish/subscribe message was delivered, discarded, or excluded).

SubId

Description:	The subscription identifier.
Identifier:	MQBACF_SUB_ID.
Data type:	MQCFBS.
Included in PCF group:	<i>Operation.</i>
Returned:	Always.

SubLevel

Description:	The subscription level.
Identifier:	MQIACF_SUB_LEVEL.
Data type:	MQCFIN.
Included in PCF group:	<i>Operation.</i>
Returned:	Always.

Feedback

Description:	The reason for discarding the message.
Identifier:	MQIACF_FEEDBACK.
Data type:	MQCFIN.
Included in PCF group:	<i>Operation.</i>
Returned:	If the message was discarded because it was not delivered to a subscriber, or the message was not delivered because the subscriber was excluded.

The Publish operation MQOPER_PUBLISH provides information about a message delivered to a particular subscriber. This operation describes the elements of the onward message that might have changed from the message described in the associated Put operation. Similarly to a Put operation, it contains a message group MQGACF_MESSAGE and, inside that, an MQMD group MQGACF_MQMD. However, this MQMD group contains only the following fields, which can be overridden by a subscriber: *Format, Priority, Persistence, MsgId, CorrelId, UserIdentifier, AccountingToken, ApplIdentityData*.

The *SubId* and *SubLevel* of the subscriber are included in the operation information. You can use the *SubID* with the MQCMD_INQUIRE_SUBSCRIBER PCF command to retrieve all other attributes for a subscriber.

The Discarded Publish operation MQOPER_DISCARDED_PUBLISH is analogous to the Discard operation that is used when a message is not delivered in point-to-point messaging. A message is not delivered to a subscriber if the message was explicitly requested not to be delivered to a local destination and this subscriber specifies a local destination. A message is also considered not delivered if there is a problem getting the message to the destination queue, for example, because the queue is full.

The information in a Discarded Publish operation is the same as for a Publish operation, with the addition of a *Feedback* field that gives the reasons why the message was not delivered. This feedback field contains MQFB_* or MQRC_* values that are common with the MQOPER_DISCARD operation. The reason for discarding a publish, as opposed to excluding it, are the same as the reasons for discarding a put.

The Excluded Publish operation MQOPER_EXCLUDED_PUBLISH provides information about a subscriber that was considered for delivery of the message, because the topic on which the subscriber is subscribing matches that of the associated Put operation, but the message was not delivered to the subscriber because other selection criteria do not match with the message that is being put to the topic. As with a Discarded Publish operation, the *Feedback* field provides information about the reason why this subscription was excluded. However, unlike the Discarded Publish operation, no message-related information is provided because no message was generated for this subscriber.

Put/Put Reply/Put Report (MQOPER_PUT/MQOPER_PUT_REPLY/MQOPER_PUT_REPORT)

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Put/Put Reply/Put Report (MQOPER_PUT/MQOPER_PUT_REPLY/MQOPER_PUT_REPORT) operation type (a message, reply message, or report message was put to a queue).

QName

Description:	The name of the queue that was opened.
Identifier:	MQCA_Q_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Maximum length:	MQ_Q_NAME_LENGTH
Returned:	Always, apart from one exception: not returned if the Put operation is to a topic, contained within a publish activity.

ResolvedQName

Description:	The name that the opened queue resolves to.
Identifier:	MQCACF_RESOLVED_Q_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Maximum length:	MQ_Q_NAME_LENGTH
Returned:	When the opened queue could be resolved. Not returned if the Put operation is to a topic, contained within a publish activity.

RemoteQName

Description:	The name of the opened queue, as it is known on the remote queue manager.
Identifier:	MQCA_REMOTE_Q_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Maximum length:	MQ_Q_NAME_LENGTH
Returned:	If the opened queue is a remote queue. Not returned if the Put operation is to a topic, contained within a publish activity.

RemoteQMGrName

Description:	The name of the remote queue manager on which the remote queue is defined.
Identifier:	MQCA_REMOTE_Q_MGR_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Maximum length:	MQ_Q_MGR_NAME_LENGTH
Returned:	If the opened queue is a remote queue. Not returned if the Put operation is to a topic, contained within a publish activity.

TopicString

Description:	The full topic string to which the message is being put.
Identifier:	MQCA_TOPIC_STRING.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Returned:	If the Put operation is to a topic, contained within a publish activity.

Feedback

Description:	The reason for the message being put on the dead-letter queue.
Identifier:	MQIACF_FEEDBACK.
Data type:	MQCFIN.
Included in PCF group:	<i>Operation.</i>
Returned:	If the message was put on the dead-letter queue.

Receive (MQOPER_RECEIVE)

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Receive (MQOPER_RECEIVE) operation type (a message was received on a channel).

ChannelName

Description:	The name of the channel on which the message was received.
Identifier:	MQCACH_CHANNEL_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Maximum length:	MQ_CHANNEL_NAME_LENGTH
Returned:	Always.

ChannelType

Description:	The type of channel on which the message was received.
Identifier:	MQIACH_CHANNEL_TYPE.
Data type:	MQCFIN.

Included in PCF group:	<i>Operation.</i>
Returned:	Always.

RemoteQMgrName

Description:	The name of the queue manager from which the message was received.
Identifier:	MQCA_REMOTE_Q_MGR_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Maximum length:	MQ_Q_MGR_NAME_LENGTH
Returned:	Always.

Send (MQOPER_SEND)

The additional activity report message data parameters that are returned in the PCF group *Operation* for the Send (MQOPER_SEND) operation type (a message was sent on a channel).

ChannelName

Description:	The name of the channel where the message was sent.
Identifier:	MQCACH_CHANNEL_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Maximum length:	MQ_CHANNEL_NAME_LENGTH.
Returned:	Always.

ChannelType

Description:	The type of channel where the message was sent.
Identifier:	MQIACH_CHANNEL_TYPE.
Data type:	MQCFIN.
Included in PCF group:	<i>Operation.</i>
Returned:	Always.

XmitQName

Description:	The transmission queue from which the message was retrieved.
Identifier:	MQCACH_XMIT_Q_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation.</i>
Maximum length:	MQ_Q_NAME_LENGTH.
Returned:	Always.

RemoteQMgrName

Description:	The name of the remote queue manager to which the message was sent.
Identifier:	MQCA_REMOTE_Q_MGR_NAME.
Data type:	MQCFST.
Included in PCF group:	<i>Operation</i> .
Maximum length:	MQ_Q_MGR_NAME_LENGTH
Returned:	Always.

Trace-route message reference

Use this page to obtain an overview of the trace-route message format. The trace-route message data includes parameters that describe the activities that the trace-route message has caused

Trace-route message format

Trace-route messages are standard IBM MQ messages containing a message descriptor and message data. The message data contains information about the activities performed on a trace-route message as it has been routed through a queue manager network.

Trace-route messages contain the following information:

A message descriptor

An MQMD structure, with the *Format* field set to MQFMT_ADMIN or MQFMT_EMBEDDED_PCF.

Message data

Consists of either:

- A PCF header (MQCFH) and trace-route message data, if *Format* is set to MQFMT_ADMIN, or
- An embedded PCF header (MQEPH), trace-route message data, and additional user-specified message data, if *Format* is set to MQFMT_EMBEDDED_PCF.

When using the IBM MQ display route application to generate a trace-route message, *Format* is set to MQFMT_ADMIN.

The content of the trace-route message data is determined by the *Accumulate* parameter from the *TraceRoute* PCF group, as follows:

- If *Accumulate* is set to MQROUTE_ACCUMULATE_NONE, the trace-route message data contains the *TraceRoute* PCF group.
- If *Accumulate* is set to either MQROUTE_ACCUMULATE_IN_MSG or MQROUTE_ACCUMULATE_AND_REPLY, the trace-route message data contains the *TraceRoute* PCF group and zero or more *Activity* PCF groups.

[Table 24 on page 147](#) shows the structure of a trace-route message.

Table 24. Trace-route message format

MQMD structure	Embedded PCF header MQEPH structure	Trace-route message data
Structure identifier Structure version Report options Message type Expiration time Feedback Encoding Coded character set ID Message format Priority Persistence Message identifier Correlation identifier Backout count Reply-to queue Reply-to queue manager User identifier Accounting token Application identity data Application type Application name Put date Put time Application origin data Group identifier Message sequence number Offset Message flags Original length	Structure identifier Structure version Structure length Encoding Coded character set ID Message format Flags PCF header (MQCFH) Structure type Structure length Structure version Command identifier Message sequence number Control options Completion code Reason code Parameter count	TraceRoute Detail Recorded activities Unrecorded activities Discontinuity count Max activities Accumulate Deliver

Trace-route message MQMD (message descriptor)

Use this page to view the values contained by the MQMD structure for a trace-route message

StrucId

Description: Structure identifier.
 Data type: MQCHAR4.
 Value: MQMD_STRUC_ID.

Version

Description: Structure version number.
 Data type: MQLONG.
 Values: **MQMD_VERSION_1.**

Report

Description: Options for report messages.
 Data type: MQLONG.

Value: Set according to requirements. Common report options follow:

MQRO_DISCARD_MSG

The message is discarded on arrival to a local queue.

MQRO_PASS_DISCARD_AND_EXPIRY

Every response (activity reports or trace-route reply message) will have the report option MQRO_DISCARD_MSG set, and the remaining expiry passed on. This ensures that responses do not remain in the queue manager network indefinitely.

MsgType

Description: Type of message.

Data type: MQLONG.

Value: If the *Accumulate* parameter in the TraceRoute group is specified as MQROUTE_ACCUMULATE_AND_REPLY, then message type is MQMT_REQUEST
Otherwise:
MQMT_DATAGRAM.

Expiry

Description: Message lifetime.

Data type: MQLONG.

Value: Set according to requirements. This parameter can be used to ensure trace-route messages are not left in a queue manager network indefinitely.

Feedback

Description: Feedback or reason code.

Data type: MQLONG.

Value: **MQFB_NONE.**

Encoding

Description: Numeric encoding of message data.

Data type: MQLONG.

Value: Set as appropriate.

CodedCharSetId

Description: Character set identifier of message data.

Data type: MQLONG.

Value: Set as appropriate.

Format

Description: Format name of message data

Data type: MQCHAR8.

Value: **MQFMT_ADMIN**
Admin message. No user data follows the *TraceRoute* PCF group.

MQFMT_EMBEDDED_PCF
Embedded PCF message. User data follows the *TraceRoute* PCF group.

Priority

Description: Message priority.

Data type: MQLONG.

Value: Set according to requirements.

Persistence

Description: Message persistence.

Data type: MQLONG.

Value: Set according to requirements.

MsgId

Description: Message identifier.

Data type: MQBYTE24.

Value: Set according to requirements.

CorrelId

Description: Correlation identifier.

Data type: MQBYTE24.

Value: Set according to requirements.

BackoutCount

Description: Backout counter.

Data type: MQLONG.

Value: 0.

ReplyToQ

Description: Name of reply queue.

Data type: MQCHAR48.

Values: Set according to requirements.

If *MsgType* is set to MQMT_REQUEST or if *Report* has any report generating options set, then this parameter must be non-blank.

ReplyToQMgr

Description: Name of reply queue manager.

Data type: MQCHAR48.

Value: Set according to requirements.

UserIdentifier

Description: The user identifier of the application that originated the message.
Data type: MQCHAR12.
Value: Set as normal.

AccountingToken

Description: Accounting token that allows an application to charge for work done as a result of the message.
Data type: MQBYTE32.
Value: Set as normal.

ApplIdentityData

Description: Application data relating to identity.
Data type: MQCHAR32.
Values: Set as normal.

PutApplType

Description: Type of application that put the message.
Data type: MQLONG.
Value: Set as normal.

PutApplName

Description: Name of application that put the message.
Data type: MQCHAR28.
Value: Set as normal.

PutDate

Description: Date when message was put.
Data type: MQCHAR8.
Value: Set as normal.

PutTime

Description: Time when message was put.
Data type: MQCHAR8.
Value: Set as normal.

ApplOriginData

Description: Application data relating to origin.
Data type: MQCHAR4.
Value: Set as normal..

Trace-route message MQEPH (Embedded PCF header)

Use this page to view the values contained by the MQEPH structure for a trace-route message

The MQEPH structure contains a description of both the PCF information that accompanies the message data of a trace-route message, and the application message data that follows it. An MQEPH structure is used only if additional user message data follows the TraceRoute PCF group.

For a trace-route message, the MQEPH structure contains the following values:

StrucId

Description:	Structure identifier.
Data type:	MQCHAR4.
Value:	MQEPH_STRUC_ID.

Version

Description:	Structure version number.
Data type:	MQLONG.
Values:	MQEPH_VERSION_1.

StrucLength

Description:	Structure length.
Data type:	MQLONG.
Value:	Total length of the structure including the PCF parameter structures that follow it.

Encoding

Description:	Numeric encoding of the message data that follows the last PCF parameter structure.
Data type:	MQLONG.
Value:	The encoding of the message data.

CodedCharSetId

Description:	Character set identifier of the message data that follows the last PCF parameter structure.
Data type:	MQLONG.
Value:	The character set of the message data.

Format

Description:	Format name of the message data that follows the last PCF parameter structure.
Data type:	MQCHAR8.
Value:	The format name of the message data.

Flags

Description:	Flags that specify attributes of the structure or control its processing.
Data type:	MQLONG.

Value:	MQEPH_NONE No flags specified.
	MQEPH_CCSID_EMBEDDED Specifies that the character set of the parameters containing character data is specified individually within the <i>CodedCharSetId</i> field in each structure.

PCFHeader

Description:	Programmable Command Format Header
Data type:	MQCFH.
Value:	See “Trace-route message MQCFH (PCF header)” on page 152.

Trace-route message MQCFH (PCF header)

Use this page to view the PCF values contained by the MQCFH structure for a trace-route message

For a trace-route message, the MQCFH structure contains the following values:

Type

Description:	Structure type that identifies the content of the message.
Data type:	MQLONG.
Value:	MQCFT_TRACE_ROUTE Message is a trace-route message.

StrucLength

Description:	Structure length.
Data type:	MQLONG.
Value:	MQCFH_STRUC_LENGTH Length in bytes of MQCFH structure.

Version

Description:	Structure version number.
Data type:	MQLONG.
Values:	MQCFH_VERSION_3

Command

Description:	Command identifier. This identifies the category of the message.
Data type:	MQLONG.
Values:	MQCMD_TRACE_ROUTE Trace-route message.

MsgSeqNumber

Description:	Message sequence number. This is the sequence number of the message within a group of related messages.
Data type:	MQLONG.
Values:	1.

Control

Description:	Control options.
Data type:	MLONG.
Values:	MQCFC_LAST.

CompCode

Description:	Completion code.
Data type:	MLONG.
Values:	MQCC_OK.

Reason

Description:	Reason code qualifying completion code.
Data type:	MLONG.
Values:	MQRC_NONE.

ParameterCount

Description:	Count of parameter structures. This is the number of parameter structures that follow the MQCFH structure. A group structure (MQCFGR), and its included parameter structures, are counted as one structure only.
Data type:	MLONG.
Values:	1 or greater.

Trace-route message data

Use this page to view the parameters that make up the *TraceRoute* PCF group part of trace-route message data

The content of trace-route message data depends on the *Accumulate* parameter from the *TraceRoute* PCF group. Trace-route message data consists of the *TraceRoute* PCF group, and zero or more *Activity* PCF groups. The *TraceRoute* PCF group is detailed in this topic. Refer to the related information for details of the *Activity* PCF group.

Trace-route message data contains the following parameters:

TraceRoute

Description:	Grouped parameters specifying attributes of the trace-route message. For a trace-route message, some of these parameters can be altered to control how it is processed.
Identifier:	MQGACF_TRACE_ROUTE.
Data type:	MQCFGR.
Contained in PCF group:	None.

Parameters in group:	<i>Detail</i>
	<i>RecordedActivities</i>
	<i>UnrecordedActivities</i>
	<i>DiscontinuityCount</i>
	<i>MaxActivities</i>
	<i>Accumulate</i>
	<i>Forward</i>
	<i>Deliver</i>

Detail

Description:	The detail level that will be recorded for the activity.
Identifier:	MQIACF_ROUTE_DETAIL.
Data type:	MQCFIN.
Contained in PCF group:	<i>TraceRoute.</i>
Values:	MQROUTE_DETAIL_LOW Activities performed by user-written application are recorded. MQROUTE_DETAIL_MEDIUM Activities specified in MQROUTE_DETAIL_LOW are recorded. Additionally, activities performed by MCAs are recorded. MQROUTE_DETAIL_HIGH Activities specified in MQROUTE_DETAIL_LOW, and MQROUTE_DETAIL_MEDIUM are recorded. MCAs do not record any further activity information at this level of detail. This option is only available to user-written applications that are to record further activity information.

RecordedActivities

Description:	The number of activities that the trace-route message has caused, where information was recorded.
Identifier:	MQIACF_RECORDED_ACTIVITIES.
Data type:	MQCFIN.
Contained in PCF group:	<i>TraceRoute.</i>

UnrecordedActivities

Description:	The number of activities that the trace-route message has caused, where information was not recorded.
Identifier:	MQIACF_UNRECORDED_ACTIVITIES.
Data type:	MQCFIN.
Contained in PCF group:	<i>TraceRoute.</i>

DiscontinuityCount

Description:	The number of times a trace-route message has been received from a queue manager that does not support trace-route messaging.
Identifier:	MQIACF_DISCONTINUITY_COUNT.

Data type: MQCFIN.
 Contained in PCF group: *TraceRoute*.

MaxActivities

Description: The maximum number of activities the trace-route message can be involved in before it stops being processed.
 Identifier: MQIACF_MAX_ACTIVITIES.
 Data type: MQCFIN.
 Contained in PCF group: *TraceRoute*.
 Value: **A positive integer**
 The maximum number of activities.
MQROUTE_UNLIMITED_ACTIVITIES
 An unlimited number of activities.

Accumulate

Description: Specifies whether activity information is accumulated within the trace-route message, and whether a reply message containing the accumulated activity information is generated before the trace-route message is discarded or is put on a non-transmission queue.
 Identifier: MQIACF_ROUTE_ACCUMULATION.
 Data type: MQCFIN.
 Contained in PCF group: *TraceRoute*.
 Value: **MQROUTE_ACCUMULATE_NONE**
 Activity information is not accumulated in the message data of the trace-route message.
MQROUTE_ACCUMULATE_IN_MSG
 Activity information is accumulated in the message data of the trace-route message.
MQROUTE_ACCUMULATE_AND_REPLY
 Activity information is accumulated in the message data of the trace-route message, and a trace-route reply message will be generated.

Forward

Description: Specifies queue managers that the trace-route message can be forwarded to. When determining whether to forward a message to a remote queue manager, queue managers use the algorithm that is described in [Forwarding](#).
 Identifier: MQIACF_ROUTE_FORWARDING.
 Data type: MQCFIN.
 Contained in PCF group: *TraceRoute*.

Value:	MQROUTE_FORWARD_IF_SUPPORTED The trace-route message is only forwarded to queue managers that will honor the value of the <i>Deliver</i> parameter from the <i>TraceRoute</i> group.
	MQROUTE_FORWARD_ALL The trace-route message is forwarded to any queue manager, regardless of whether the value of the <i>Deliver</i> parameter will be honored.

Deliver

Description:	Specifies the action to be taken if the trace-route message arrives at the destination queue successfully.
Identifier:	MQIACF_ROUTE_DELIVERY.
Data type:	MQCFIN.
Contained in PCF group:	<i>TraceRoute</i> .
Value:	MQROUTE_DELIVER_YES On arrival, the trace-route message is put on the target queue. Any application performing a destructive get on the target queue can receive the trace-route message.
	MQROUTE_DELIVER_NO On arrival, the trace-route message is discarded.

Trace-route reply message reference

Use this page to obtain an overview of the trace-route reply message format. The trace-route reply message data is a duplicate of the trace-route message data from the trace-route message for which it was generated

Trace-route reply message format

Trace-route reply messages are standard IBM MQ messages containing a message descriptor and message data. The message data contains information about the activities performed on a trace-route message as it has been routed through a queue manager network.

Trace-route reply messages contain the following information:

A message descriptor

An MQMD structure

Message data

A PCF header (MQCFH) and trace-route reply message data

Trace-route reply message data consists of one or more *Activity* PCF groups.

When a trace-route message reaches its target queue, a trace-route reply message can be generated that contains a copy of the activity information from the trace-route message. The trace-route reply message will be delivered to a reply-to queue or to a system queue.

Table 25 on page 157 shows the structure of a trace-route reply message, including parameters that are only returned under certain conditions.

Table 25. Trace-route reply message format

MQMD structure	PCF header MQCFH structure	Trace-route reply message data
Structure identifier Structure version Report options Message type Expiration time Feedback Encoding Coded character set ID Message format Priority Persistence Message identifier Correlation identifier Backout count Reply-to queue Reply-to queue manager User identifier Accounting token Application identity data Application type Application name Put date Put time Application origin data Group identifier Message sequence number Offset Message flags Original length	PCF header (MQCFH) Structure type Structure length Structure version Command identifier Message sequence number Control options Completion code Reason code Parameter count	Activity Activity application name Activity application type Activity description Operation Operation type Operation date Operation time Message Message length MQMD EmbeddedMQMD Queue manager name Queue sharing group name Queue name ^{1 2 3} Resolved queue name ^{1 3} Remote queue name ³ Remote queue manager-name ^{2 3 4 5} Feedback ² Channel name ^{4 5} Channel type ^{4 5} Transmission queue name ⁵ TraceRoute Detail Recorded activities Unrecorded activities Discontinuity count Max activities Accumulate Deliver

Note:

1. Returned for Get and Browse operations.
2. Returned for Discard operations.
3. Returned for Put, Put Reply, and Put Report operations.
4. Returned for Receive operations.
5. Returned for Send operations.

Trace-route reply message MQMD (message descriptor)

Use this page to view the values contained by the MQMD structure for a trace-route reply message

For a trace-route reply message, the MQMD structure contains the parameters described in [Activity report message descriptor](#). Some of the parameter values in a trace-route reply message descriptor are different from those in an activity report message descriptor, as follows:

MsgType

Description: Type of message.
 Data type: MQLONG.

Value: **MQMT_REPLY**

Feedback

Description: Feedback or reason code.

Data type: MQLONG.

Value: **MQFB_NONE**

Encoding

Description: Numeric encoding of message data.

Data type: MQLONG.

Value: Copied from trace-route message descriptor.

CodedCharSetId

Description: Character set identifier of message data.

Data type: MQLONG.

Value: Copied from trace-route message descriptor.

Format

Description: Format name of message data

Data type: MQCHAR8.

Value: **MQFMT_ADMIN**
Admin message.

Trace-route reply message MQCFH (PCF header)

Use this page to view the PCF values contained by the MQCFH structure for a trace-route reply message

The PCF header (MQCFH) for a trace-route reply message is the same as for a trace-route message.

Trace-route reply message data

The trace-route reply message data is a duplicate of the trace-route message data from the trace-route message for which it was generated

The trace-route reply message data contains one or more *Activity* groups. The parameters are described in [“Activity report message data” on page 129](#).

Accounting and statistics messages

Queue managers generate accounting and statistics messages to record information about the MQI operations performed by IBM MQ applications, or to record information about the activities occurring in an IBM MQ system.

Accounting messages

Accounting messages are used to record information about the MQI operations performed by IBM MQ applications, see [“Accounting messages” on page 159](#).

Statistics messages

Statistics messages are used to record information about the activities occurring in an IBM MQ system, see [“Statistics messages” on page 162](#). Some activity recorded in statistics messages relates to internal queue manager operations.



Accounting messages and statistics messages as described here are not available on IBM MQ for z/OS, but equivalent functionality is available through the System Management Facility (SMF).

Accounting and statistics messages are delivered to one of two system queues. User applications can retrieve the messages from these system queues and use the recorded information for various purposes:

- Account for application resource use.
- Record application activity.
- Capacity planning.
- Detect problems in your queue manager network.
- Assist in determining the causes of problems in your queue manager network.
- Improve the efficiency of your queue manager network.
- Familiarize yourself with the running of your queue manager network.
- Confirm that your queue manager network is running correctly.

Related concepts

[“Using System Management Facility” on page 343](#)

You can use SMF to collect statistics and accounting information. To use SMF, certain parameters must be set in z/OS and in IBM MQ.

Accounting messages

Accounting messages record information about the MQI operations performed by IBM MQ applications. An accounting message is a PCF message that contains a number of PCF structures.

When an application disconnects from a queue manager, an accounting message is generated and delivered to the system accounting queue (SYSTEM.ADMIN.ACCOUNTING.QUEUE). For long running IBM MQ applications, intermediate accounting messages are generated as follows:

- When the time since the connection was established exceeds the configured interval.
- When the time since the last intermediate accounting message exceeds the configured interval.

Accounting messages are in the following categories:

MQI accounting messages

MQI accounting messages contain information relating to the number of MQI calls made using a connection to a queue manager.

Queue accounting messages

Queue accounting messages contain information relating to the number of MQI calls made using connections to a queue manager, grouped by queue.

Each queue accounting message can contain up to 100 records, with every record relating to an activity performed by the application with respect to a specific queue.

Accounting messages are recorded only for local queues. If an application makes an MQI call against an alias queue, the accounting data is recorded against the base queue, and, for a remote queue, the accounting data is recorded against the transmission queue.

Note: As the information relates to MQI operations performed by IBM MQ applications, the information does not include operations relating to a streaming queue, unless the operation is directly on that queue by the application.

Related reference

[“MQI accounting message data” on page 177](#)

Use this page to view the structure of an MQI accounting message

[“Queue accounting message data” on page 188](#)

Use this page to view the structure of a queue accounting message

Accounting message format

Accounting messages comprise a set of PCF fields that consist of a message descriptor and message data.

Message descriptor

- An accounting message MQMD (message descriptor)

Accounting message data

- An accounting message MQCFH (PCF header)
- Accounting message data that is always returned
- Accounting message data that is returned if available

The accounting message MQCFH (PCF header) contains information about the application, and the interval for which the accounting data was recorded.

Accounting message data comprises PCF parameters that store the accounting information. The content of accounting messages depends on the message category as follows:

MQI accounting message

MQI accounting message data consists of a number of PCF parameters, but no PCF groups.

Queue accounting message

Queue accounting message data consists of a number of PCF parameters, and in the range 1 through 100 *QAccountingData* PCF groups.

There is one *QAccountingData* PCF group for every queue that had accounting data collected. If an application accesses more than 100 queues, multiple accounting messages are generated. Each message has the *SeqNumber* in the MQCFH (PCF header) updated accordingly, and the last message in the sequence has the *Control* parameter in the MQCFH specified as MQCFC_LAST.

Accounting information collection

Use queue and queue manager attributes to control the collection of accounting information. You can also use MQCONN options to control collection at the connection level.

Controlling the collection of MQI accounting information

Use the queue manager attribute ACCTMQI to control the collection of MQI accounting information.

To change the value of this attribute, use the MQSC command, ALTER QMGR, and specify the parameter ACCTMQI. Accounting messages are generated only for connections that begin after accounting is enabled. The **ACCTMQI** parameter can have the following values:

ON

MQI accounting information is collected for every connection to the queue manager.

OFF

MQI accounting information is not collected. This is the default value.

For example, to enable MQI accounting information collection use the following MQSC command:

```
ALTER QMGR ACCTMQI(ON)
```

Queue accounting information

Use the queue attribute ACCTQ and the queue manager attribute ACCTQ to control the collection of queue accounting information.

To change the value of the queue attribute, use the MQSC command, ALTER QLOCAL, and specify the parameter ACCTQ. Accounting messages are generated only for connections that begin after accounting is enabled. Note that changes to this value are only effective for connections to the queue manager that occur after the change to the attribute.

The queue attribute ACCTQ can have the following values:

ON

Queue accounting information for this queue is collected for every connection to the queue manager that opens the queue.

OFF

Queue accounting information for this queue is not collected.

QMGR

The collection of queue accounting information for this queue is controlled according to the value of the queue manager attribute ACCTQ. This is the default value.

To change the value of the queue manager attribute, use the MQSC command, ALTER QMGR and specify the parameter ACCTQ. The queue manager attribute ACCTQ can have the following values:

ON

Queue accounting information is collected for queues that have the queue attribute ACCTQ set as QMGR.

OFF

Queue accounting information is not collected for queues that have the queue attribute ACCTQ set as QMGR. This is the default value.

NONE

The collection of queue accounting information is disabled for all queues, regardless of the queue attribute ACCTQ.

If the queue manager attribute, ACCTQ, is set to NONE, the collection of queue accounting information is disabled for all queues, regardless of the queue attribute ACCTQ.

For example, to enable accounting information collection for the queue, Q1, use the following MQSC command:

```
ALTER QLOCAL(Q1) ACCTQ(ON)
```

To enable accounting information collection for all queues that specify the queue attribute ACCTQ as QMGR, use the following MQSC command:

```
ALTER QMGR ACCTQ(ON)
```

MQCONN options for controlling accounting information collection

Use the **ConnectOpts** parameter on the MQCONN call to modify the collection of both MQI and queue accounting information at the connection level by overriding the effective values of the queue manager attributes ACCTMQI and ACCTQ

The **ConnectOpts** parameter can have the following values:

MQCNO_ACCOUNTING_MQI_ENABLED

If the value of the queue manager attribute ACCTMQI is specified as OFF, MQI accounting is enabled for this connection. This is equivalent of the queue manager attribute ACCTMQI being specified as ON.

If the value of the queue manager attribute ACCTMQI is not specified as OFF, this attribute has no effect.

MQCNO_ACCOUNTING_MQI_DISABLED

If the value of the queue manager attribute ACCTMQI is specified as ON, MQI accounting is disabled for this connection. This is equivalent of the queue manager attribute ACCTMQI being specified as OFF.

If the value of the queue manager attribute ACCTMQI is not specified as ON, this attribute has no effect.

MQCNO_ACCOUNTING_Q_ENABLED

If the value of the queue manager attribute ACCTQ is specified as OFF, queue accounting is enabled for this connection. All queues with ACCTQ specified as QMGR, are enabled for queue accounting. This is equivalent of the queue manager attribute ACCTQ being specified as ON.

If the value of the queue manager attribute ACCTQ is not specified as OFF, this attribute has no effect.

MQCNO_ACCOUNTING_Q_DISABLED

If the value of the queue manager attribute ACCTQ is specified as ON, queue accounting is disabled for this connection. This is equivalent of the queue manager attribute ACCTQ being specified as OFF.

If the value of the queue manager attribute ACCTQ is not specified as ON, this attribute has no effect.

These overrides are by disabled by default. To enable them, set the queue manager attribute ACCTCONO to ENABLED. To enable accounting overrides for individual connections use the following MQSC command:

```
ALTER QMGR ACCTCONO(ENABLED)
```

Accounting message generation when applications disconnect

Accounting messages are generated when an application disconnects from the queue manager. Intermediate accounting messages are also written for long running IBM MQ applications.

Accounting messages are generated in either of the following ways when an application disconnects:

- The application issues an MQDISC call
- The queue manager recognises that the application has terminated

Intermediate accounting messages are written for long running IBM MQ applications when the interval since the connection was established or since the last intermediate accounting message that was written exceeds the configured interval. The queue manager attribute, ACCTINT, specifies the time, in seconds, after which intermediate accounting messages can be automatically written. Accounting messages are generated only when the application interacts with the queue manager, so applications that remain connected to the queue manager for long periods without executing MQI requests do not generate accounting messages until the execution of the first MQI request following the completion of the accounting interval.

The default accounting interval is 1800 seconds (30 minutes). For example, to change the accounting interval to 900 seconds (15 minutes) use the following MQSC command:

```
ALTER QMGR ACCTINT(900)
```

Statistics messages

Statistics messages record information about the activities occurring in an IBM MQ system. A statistics message is a PCF message that contains a number of PCF structures.

Statistics messages are delivered to the system queue (SYSTEM.ADMIN.STATISTICS.QUEUE) at configured intervals, whenever there is some activity.

Statistics messages are in the following categories:

MQI statistics messages

MQI statistics messages contain information relating to the number of MQI calls made during a configured interval. For example, the information can include the number of MQI calls issued by a queue manager.

Queue statistics messages

Queue statistics messages contain information relating to the activity of a queue during a configured interval. The information includes the number of messages put on, and retrieved from, the queue, and the total number of bytes processed by a queue.

Each queue statistics message can contain up to 100 records, with each record relating to the activity per queue for which statistics were collected.

Statistics messages are recorded only for local queues. If an application makes an MQI call against an alias queue, the statistics data is recorded against the base queue, and, for a remote queue, the statistics data is recorded against the transmission queue.

Channel statistics messages

Channel statistics messages contain information relating to the activity of a channel during a configured interval. For example the information might be the number of messages transferred by the channel, or the number of bytes transferred by the channel.

Each channel statistics message contains up to 100 records, with each record relating to the activity per channel for which statistics were collected.

Note: As the information relates to activities occurring in an IBM MQ system, the information does include operations relating to streaming queues.

Related reference

[“Controlling MQI statistics collection” on page 164](#)

Use the queue manager attribute `STATMQI` to control the collection of MQI statistics information.

[“Controlling queue statistics collection” on page 164](#)

Use the queue attribute `STATQ` and the queue manager attribute `STATQ` to control the collection of queue statistics information

[“Controlling channel statistics collection” on page 165](#)

Use the channel attribute `STATCHL` to control the collection of channel statistics information. You can also set queue manager attributes to control information collection.

Statistics messages format

Statistics messages comprise a set of PCF fields that consist of a message descriptor and message data.

Message descriptor

- A statistics message `MQMD` (message descriptor)

Accounting message data

- A statistics message `MQCFH` (PCF header)
- Statistics message data that is always returned
- Statistics message data that is returned if available

The statistics message `MQCFH` (PCF header) contains information about the interval for which the statistics data was recorded.

Statistics message data comprises PCF parameters that store the statistics information. The content of statistics messages depends on the message category as follows:

MQI statistics message

MQI statistics message data consists of a number of PCF parameters, but no PCF groups.

Queue statistics message

Queue statistics message data consists of a number of PCF parameters, and in the range 1 through 100 *QStatisticsData* PCF groups.

There is one *QStatisticsData* PCF group for every queue was active in the interval. If more than 100 queues were active in the interval, multiple statistics messages are generated. Each message has the *SeqNumber* in the `MQCFH` (PCF header) updated accordingly, and the last message in the sequence has the *Control* parameter in the `MQCFH` specified as `MQCFC_LAST`.

Channel statistics message

Channel statistics message data consists of a number of PCF parameters, and in the range 1 through 100 *ChlStatisticsData* PCF groups.

There is one *ChlStatisticsData* PCF group for every channel that was active in the interval. If more than 100 channels were active in the interval, multiple statistics messages are generated. Each message has the *SeqNumber* in the `MQCFH` (PCF header) updated accordingly, and the last message in the sequence has the *Control* parameter in the `MQCFH` specified as `MQCFC_LAST`.

Statistics information collection

Use queue, queue manager, and channel attributes to control the collection of statistics information

Controlling MQI statistics collection

Use the queue manager attribute **STATMQI** to control the collection of MQI statistics information.

To change the value of this attribute, use the MQSC command, **ALTER QMGR** and specify the parameter **STATMQI**. Statistics messages are generated only for queues that are opened after statistics collection has been enabled. The **STATMQI** parameter can have the following values:

ON

MQI statistics information is collected for every connection to the queue manager.

OFF

MQI statistics information is not collected. This is the default value.

For example, to enable MQI statistics information collection use the following MQSC command:

```
ALTER QMGR STATMQI(ON)
```

Controlling queue statistics collection

Use the queue attribute **STATQ** and the queue manager attribute **STATQ** to control the collection of queue statistics information

You can enable or disable queue statistics information collection for individual queues or for multiple queues. To control individual queues, set the queue attribute **STATQ**. You enable or disable queue statistics information collection at the queue manager level by using the queue manager attribute **STATQ**. For all queues that have the queue attribute **STATQ** specified with the value **QMGR**, queue statistics information collection is controlled at the queue manager level.

Queue statistics are incremented only for operations using IBM MQ MQI Object Handles that were opened after statistics collection has been enabled.

Queue Statistics messages are generated only for queues for which statistics data has been collected in the previous time period.

The same queue can have several put operations and get operations through several Object Handles. Some Object Handles might have been opened before statistics collection was enabled, but others were opened afterwards. Therefore, it is possible for the queue statistics to record the activity of some put operations and get operations, and not all.

To ensure that the Queue Statistics are recording the activity of all applications, you must close and reopen new Object Handles on the queue, or queues, that you are monitoring. The best way to achieve this, is to end and restart all applications after enabling statistics collection.

To change the value of the queue attribute **STATQ**, use the MQSC command, **ALTER QLOCAL** and specify the parameter **STATQ**. The queue attribute **STATQ** can have the following values:

ON

Queue statistics information is collected for every connection to the queue manager that opens the queue.

OFF

Queue statistics information for this queue is not collected.

QMGR

The collection of queue statistics information for this queue is controlled according to the value of the queue manager attribute, **STATQ**. This is the default value.

To change the value of the queue manager attribute **STATQ**, use the MQSC command, **ALTER QMGR** and specify the parameter **STATQ**. The queue manager attribute **STATQ** can have the following values:

ON

Queue statistics information is collected for queues that have the queue attribute **STATQ** set as **QMGR**

OFF

Queue statistics information is not collected for queues that have the queue attribute STATQ set as QMGR. This is the default value.

NONE

The collection of queue statistics information is disabled for all queues, regardless of the queue attribute STATQ.

If the queue manager attribute STATQ is set to NONE, the collection of queue statistics information is disabled for all queues, regardless of the queue attribute STATQ.

For example, to enable statistics information collection for the queue, Q1, use the following MQSC command:

```
ALTER QLOCAL(Q1) STATQ(ON)
```

To enable statistics information collection for all queues that specify the queue attribute STATQ as QMGR, use the following MQSC command:

```
ALTER QMGR STATQ(ON)
```

Controlling channel statistics collection

Use the channel attribute STATCHL to control the collection of channel statistics information. You can also set queue manager attributes to control information collection.

You can enable or disable channel statistics information collection for individual channels, or for multiple channels. To control individual channels, you must set the channel attribute STATCHL to enable or disable channel statistic information collection. To control many channels together, you enable or disable channel statistics information collection at the queue manager level by using the queue manager attribute STATCHL. For all channels that have the channel attribute STATCHL specified with the value QMGR, channel statistics information collection is controlled at the queue manager level.

Automatically defined cluster-sender channels are not IBM MQ objects, so do not have attributes in the same way as channel objects. To control automatically defined cluster-sender channels, use the queue manager attribute STATACLS. This attribute determines whether automatically defined cluster-sender channels within a queue manager are enabled or disabled for channel statistics information collection.

You can set channel statistics information collection to one of the three monitoring levels: low, medium or high. You can set the monitoring level at either object level or at the queue manager level. The choice of which level to use is dependent on your system. Collecting statistics information data might require some instructions that are relatively expensive computationally, so to reduce the impact of channel statistics information collection, the medium and low monitoring options measure a sample of the data at regular intervals rather than collecting data all the time. [Table 26 on page 165](#) summarizes the levels available with channel statistics information collection:

Table 26. Detail level of channel statistics information collection		
Level	Description	Usage
Low	Measure a small sample of the data, at regular intervals.	For objects that process a high volume of messages.
Medium	Measure a sample of the data, at regular intervals.	For most objects.
High	Measure all data, at regular intervals.	For objects that process only a few messages per second, on which the most current information is important.

To change the value of the channel attribute STATCHL, use the MQSC command, ALTER CHANNEL and specify the parameter STATCHL.

To change the value of the queue manager attribute STATCHL, use the MQSC command, ALTER QMGR and specify the parameter STATCHL.

To change the value of the queue manager attribute STATACLS, use the MQSC command, ALTER QMGR and specify the parameter STATACLS.

The channel attribute, STATCHL, can have the following values:

LOW

Channel statistics information is collected with a low level of detail.

MEDIUM

Channel statistics information is collected with a medium level of detail.

HIGH

Channel statistics information is collected with a high level of detail.

OFF

Channel statistics information is not collected for this channel.

QMGR

The channel attribute is set as QMGR. The collection of statistics information for this channel is controlled by the value of the queue manager attribute, STATCHL.

This is the default value.

► **z/OS** On z/OS systems, enabling this parameter simply turns on statistics data collection, regardless of the value you select. Specifying LOW, MEDIUM, or HIGH makes no difference to your results. This parameter must be enabled in order to collect channel accounting records.

The queue manager attribute, STATCHL, can have the following values:

LOW

Channel statistics information is collected with a low level of detail, for all channels that have the channel attribute STATCHL set as QMGR.

MEDIUM

Channel statistics information is collected with a medium level of detail, for all channels that have the channel attribute STATCHL set as QMGR.

HIGH

Channel statistics information is collected with a high level of detail, for all channels that have the channel attribute STATCHL set as QMGR.

OFF

Channel statistics information is not collected for all channels that have the channel attribute STATCHL set as QMGR.

This is the default value.

NONE

The collection of channel statistics information is disabled for all channel, regardless of the channel attribute STATCHL.

► **z/OS** On z/OS systems, enabling this parameter simply turns on statistics data collection, regardless of the value you select. Specifying LOW, MEDIUM, or HIGH makes no difference to your results. This parameter must be enabled in order to collect channel accounting records.

The queue manager attribute, STATACLS, can have the following values:

LOW

Statistics information is collected with a low level of detail for automatically defined cluster-sender channels.

MEDIUM

Statistics information is collected with a medium level of detail for automatically defined cluster-sender channels.

HIGH

Statistics information is collected with a high level of detail for automatically defined cluster-sender channels.

OFF

Statistics information is not collected for automatically defined cluster-sender channels.

QMGR

The collection of statistics information for automatically defined cluster-sender channels is controlled by the value of the queue manager attribute, STATCHL.

This is the default value.

 On z/OS systems, enabling this parameter simply turns on statistics data collection, regardless of the value you select. Specifying LOW, MEDIUM, or HIGH makes no difference to your results. This parameter must be enabled in order to collect channel accounting records.

For example, to enable statistics information collection, with a medium level of detail, for the sender channel QM1.T0.QM2, use the following MQSC command:

```
ALTER CHANNEL(QM1.T0.QM2) CHLTYPE(SDR) STATCHL(MEDIUM)
```

To enable statistics information collection, at a medium level of detail, for all channels that specify the channel attribute STATCHL as QMGR, use the following MQSC command:

```
ALTER QMGR STATCHL(MEDIUM)
```

To enable statistics information collection, at a medium level of detail, for all automatically defined cluster-sender channels, use the following MQSC command:

```
ALTER QMGR STATACLS(MEDIUM)
```

Statistics message generation

Statistics messages are generated at configured intervals, and when a queue manager shuts down in a controlled fashion.

The configured interval is controlled by the STATINT queue manager attribute, which specifies the interval, in seconds, between the generation of statistics messages. The default statistics interval is 1800 seconds (30 minutes). To change the statistics interval, use the MQSC command ALTER QMGR and specify the **STATINT** parameter. For example, to change the statistics interval to 900 seconds (15 minutes) use the following MQSC command:

```
ALTER QMGR STATINT(900)
```

To write the currently collected statistics data to the statistics queue before the statistics collection interval is due to expire, use the MQSC command RESET QMGR TYPE(STATISTICS). Issuing this command causes the collected statistics data to be written to the statistics queue and a new statistics data collection interval to begin.

Displaying accounting and statistics information

To use the information recorded in accounting and statistics messages, run an application such as the **amqsmon** sample program to transform the recorded information into a suitable format


Accounting and statistics messages are written to the system accounting and statistics queues. **amqsmon** is a sample program supplied with IBM MQ that processes messages from the accounting and statistics queues and displays the information to the screen in a readable form.

Because **amqsmon** is a sample program, you can use the supplied source code as template for writing your own application to process accounting or statistics messages, or modify the **amqsmon** source code to meet your own particular requirements.

amqsmon (Display formatted monitoring information)

Use the **amqsmon** sample program to display in a readable format the information contained within accounting and statistics messages. The **amqsmon** program reads accounting messages from the accounting queue, SYSTEM.ADMIN.ACCOUNTING.QUEUE. and reads statistics messages from the statistics queue, SYSTEM.ADMIN.STATISTICS.QUEUE.

Syntax



Required parameters

-t *Type*

The type of messages to process. Specify *Type* as one of the following:

accounting

Accounting records are processed. Messages are read from the system queue SYSTEM.ADMIN.ACCOUNTING.QUEUE unless you have used the -n parameter to select a particular queue to read from.

statistics

Statistics records are processed. Messages are read from the system queue SYSTEM.ADMIN.STATISTICS.QUEUE unless you have used the -n parameter to select a particular queue to read from.

Optional Parameters

-m *QMGrName*

The name of the queue manager from which accounting or statistics messages are to be processed.

If you do not specify this parameter, the default queue manager is used.

-a

Process messages containing MQI records only.

Only display MQI records. Messages not containing MQI records will always be left on the queue they were read from.

-q *QueueName*

QueueName is an optional parameter.

If *QueueName* is not supplied:

Displays queue accounting and queue statistics records only.

If *QueueName* is supplied:

Displays queue accounting and queue statistics records for the queue specified by *QueueName* only.

If *-b* is not specified then the accounting and statistics messages from which the records came are discarded. Since accounting and statistics messages can also contain records from other queues, if *-b* is not specified then unseen records can be discarded.

-c *ChannelName*

ChannelName is an optional parameter.

If *ChannelName* is not supplied:

Displays channel statistics records only.

If *ChannelName* is supplied:

Displays channel statistics records for the channel specified by *ChannelName* only.

If *-b* is not specified then the statistics messages from which the records came are discarded. Since statistics messages can also contain records from other channels, if *-b* is not specified then unseen records can be discarded.

This parameter is available when displaying statistics messages only, (*-t statistics*).

-qm *qm name*

This parameter is optional but if you specify it you must provide a queue manager name

Specifies that only accounting or statistics messages produced by the given queue manager should be read from the queue. Typically, only used if messages from several queue managers have been forwarded to a central queue manager.

-i *ConnectionId*

Displays records related to the connection identifier specified by *ConnectionId* only.

This parameter is available when displaying accounting messages only, (*-t accounting*).

If *-b* is not specified then the statistics messages from which the records came are discarded. Since statistics messages can also contain records from other channels, if *-b* is not specified then unseen records can be discarded.

-n *input queue*

The queue to read accounting or statistics messages from. If this is not specified the default queue for the chosen *-t <type>* is used.

-b

Browse messages.

Messages are retrieved non-destructively.

-d *Depth*

The maximum number of messages that can be processed.

If you do not specify this parameter, then an unlimited number of messages can be processed.

-w *TimeOut*

Time maximum number of seconds to wait for a message to become available.

If you do not specify this parameter, amqsmon will end once there are no more messages to process.

-s StartTime

Process messages put after the specified *StartTime* only.

StartTime is specified in the format yyyy-mm-dd hh.mm.ss. If a date is specified without a time, then the time will default to 00.00.00 on the date specified. Times are in GMT.

For the effect of not specifying this parameter, see [Note 1](#).

-e EndTime

Process messages put before the specified *EndTime* only.

The *EndTime* is specified in the format yyyy-mm-dd hh.mm.ss. If a date is specified without a time, then the time will default to 23.59.59 on the date specified. Times are in GMT.

For the effect of not specifying this parameter, see [Note 1](#).

-l Parameter

Only display the selected fields from the records processed. *Parameter* is a comma-separated list of integer values, with each integer value mapping to the numeric constant of a field, see [amqsmon example 5](#).

If you do not specify this parameter, then all available fields are displayed.

Note:

1. If you do not specify *-s StartTime* or *-e EndTime*, the messages that can be processed are not restricted by put time.

amqsmon (Display formatted monitoring information) examples

Use this page to view examples of running the amqsmon (Display formatted monitoring information) sample program

1. See [“Queue statistics message data”](#) on page 211 for an explanation of the attributes.

The following command displays all MQI statistics messages from queue manager saturn.queue.manager:

```
amqsmon -m saturn.queue.manager -t statistics -a
```

The output from this command follows:

```
RecordType: MQIStatistics
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.02'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.39.02'
CommandLevel: 600
ConnCount: 23
ConnFailCount: 0
ConnsMax: 8
DiscCount: [17, 0, 0]
OpenCount: [0, 80, 1, 0, 0, 3, 0, 0, 0, 0, 0, 0, 0]
OpenFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
CloseCount: [0, 73, 1, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0]
CloseFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
InqCount: [4, 2102, 0, 0, 0, 46, 0, 0, 0, 0, 0, 0, 0]
InqFailCount: [0, 31, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
SetCount: [0, 1, 0, 0, 0, 1, 0, 0, 0, 0, 0, 0, 0]
SetFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
PutCount: [26, 1]
PutFailCount: 0
Put1Count: [40, 0]
Put1FailCount: 0
PutBytes: [57064, 12320]
GetCount: [18, 1]
GetBytes: [52, 12320]
GetFailCount: 2254
BrowseCount: [18, 60]
BrowseBytes: [23784, 30760]
BrowseFailCount: 9
```

```
CommitCount: 0
CommitFailCount: 0
BackCount: 0
ExpiredMsgCount: 0
PurgeCount: 0
```

2. The following command displays all queue statistics messages for queue LOCALQ on queue manager saturn.queue.manager:

```
amqsmon -m saturn.queue.manager -t statistics -q LOCALQ
```

The output from this command follows:

```
RecordType: QueueStatistics
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.02'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.39.02'
CommandLevel: 600
ObjectCount: 3
QueueStatistics:
  QueueName: 'LOCALQ'
  CreateDate: '2005-03-08'
  CreateTime: '17.07.02'
  QueueType: Predefined
  QueueDefinitionType: Local
  QMinDepth: 0
  QMaxDepth: 18
  AverageQueueTime: [29827281, 0]
  PutCount: [26, 0]
  PutFailCount: 0
  Put1Count: [0, 0]
  Put1FailCount: 0
  PutBytes: [88, 0]
  GetCount: [18, 0]
  GetBytes: [52, 0]
  GetFailCount: 0
  BrowseCount: [0, 0]
  BrowseBytes: [0, 0]
  BrowseFailCount: 1
  NonQueuedMsgCount: 0
  ExpiredMsgCount: 0
  PurgedMsgCount: 0
```

3. The following command displays all of the statistics messages recorded since 15:30 on 30 April 2005 from queue manager saturn.queue.manager.

```
amqsmon -m saturn.queue.manager -t statistics -s "2005-04-30 15.30.00"
```

The output from this command follows:

```
RecordType: MQIStatistics
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.02'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.39.02'
CommandLevel: 600
ConnCount: 23
ConnFailCount: 0
ConnsMax: 8
DiscCount: [17, 0, 0]
OpenCount: [0, 80, 1, 0, 0, 3, 0, 0, 0, 0, 0, 0]
...
RecordType: QueueStatistics
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.02'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.39.02'
CommandLevel: 600
ObjectCount: 3
QueueStatistics: 0
  QueueName: 'LOCALQ'
  CreateDate: '2005-03-08'
  CreateTime: '17.07.02'
```

```

QueueType: Predefined
...
QueueStatistics: 1
QueueName: 'SAMPLEQ'
CreateDate: '2005-03-08'
CreateTime: '17.07.02'
QueueType: Predefined
...

```

4. See “Queue accounting message data” on page 188 for an explanation of the attributes.

The following command displays all accounting messages recorded on 30 April 2005 from queue manager saturn.queue.manager:

```
amqsmmon -m saturn.queue.manager -t accounting -s "2005-04-30" -e "2005-04-30"
```

The output from this command follows:

```

RecordType: MQIAccounting
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-04-30'
IntervalStartTime: '15.09.29'
IntervalEndDate: '2005-04-30'
IntervalEndTime: '15.09.30'
CommandLevel: 600
ConnectionId: x'414d514354524556312020202020208d0b3742010a0020'
SeqNumber: 0
ApplicationName: 'amqsput'
ApplicationPid: 8572
ApplicationTid: 1
UserId: 'admin'
ConnDate: '2005-03-16'
ConnTime: '15.09.29'
DiscDate: '2005-03-16'
DiscTime: '15.09.30'
DiscType: Normal
OpenCount: [0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
OpenFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
CloseCount: [0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
CloseFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
PutCount: [1, 0]
PutFailCount: 0
PutBytes: [4, 0]
GetCount: [0, 0]
GetFailCount: 0
GetBytes: [0, 0]
BrowseCount: [0, 0]
BrowseFailCount: 0
BrowseBytes: [0, 0]
CommitCount: 0
CommitFailCount: 0
BackCount: 0
InqCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
InqFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
SetCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]
SetFailCount: [0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0]

RecordType: MQIAccounting
QueueManager: 'saturn.queue.manager'
IntervalStartDate: '2005-03-16'
IntervalStartTime: '15.16.22'
IntervalEndDate: '2005-03-16'
IntervalEndTime: '15.16.22'
CommandLevel: 600
ConnectionId: x'414d514354524556312020202020208d0b3742010c0020'
SeqNumber: 0
ApplicationName: 'runmqsc'
ApplicationPid: 8615
ApplicationTid: 1
...

```

5. The following command browses the accounting queue and displays the application name and connection identifier of every application for which MQI accounting information is available:

```
amqsmmon -m saturn.queue.manager -t accounting -b -l 7006,3024
```

```
MonitoringType: QueueAccounting
ConnectionId: x'414d5143514d39303520202020202020fcf1855e01e80322'
ApplicationName: 'WebSphere MQ\bin\amqsput.exe'
QueueAccounting: 0

MonitoringType: QueueAccounting
ConnectionId: x'414d5143514d39303520202020202020fcf1855e01ea0322'
ApplicationName: 'BM\MQ_4\bin64\MQExplorer.exe'
QueueAccounting: 0
QueueAccounting: 1
QueueAccounting: 2
QueueAccounting: 3
QueueAccounting: 4
QueueAccounting: 5
QueueAccounting: 6
QueueAccounting: 7
QueueAccounting: 8
QueueAccounting: 9

MonitoringType: QueueAccounting
ConnectionId: x'414d5143514d39303520202020202020fcf1855e01e90322'
ApplicationName: 's\IBM\MQ_4\bin64\amqsput.exe'
QueueAccounting: 0

MonitoringType: QueueAccounting
ConnectionId: x'414d5143514d39303520202020202020fcf1855e01ef0322'
ApplicationName: 'BM\MQ_4\bin64\MQExplorer.exe'
QueueAccounting: 0
QueueAccounting: 1
QueueAccounting: 2
QueueAccounting: 3
QueueAccounting: 4
QueueAccounting: 5
QueueAccounting: 6
QueueAccounting: 7
QueueAccounting: 8
QueueAccounting: 9

MonitoringType: QueueAccounting
ConnectionId: x'414d5143514d39303520202020202020fcf1855e01e60322'
ApplicationName: 's\IBM\MQ_4\bin64\runmqsc.exe'
QueueAccounting: 0

5 Records Processed
```

See “Finding the mapping for AvgTimeOnQ (703) and QmaxDepth (739)” on page 174 for details on how you locate variables if you selected *statistics* for the **Type** parameter in the **amqsmmon** command.

Use this topic to explain how the variables used in the example are found, for monitoring accounting data.

- 3024 means MQCACF_APPL_NAME which is the **ApplicationName**
- 7006 means MQBACF_CONNECTION_ID which is the **ConnectionId**

1. Visit [“Queue accounting message data”](#) on page 188 for an explanation of the attributes [“ApplicationName”](#) on page 190 and [“ConnectionId”](#) on page 189 as you selected *accounting* for the **Type** parameter in the **amqsmon** command.

The **Identifier** for **ApplicationName** is MQCACF_APPL_NAME, and for **ConnectionId** is MOBACF CONNECTION ID

IBM MQ Monitoring and performance **173**

Go to the [Constants](#) section and scroll down until you find the [MQCACF_* \(Command format Character Parameter Types\)](#) list. Locate MQCACF_APPL_NAME and you see the value 3024.

Similarly, find the [MQBACF_* \(Command format Byte Parameter Types\)](#) list. Locate MQBACF_CONNECTION_ID and you see the value 7006.

Finding the mapping for AvgTimeOnQ (703) and QmaxDepth (739)

Use this topic to explain how the variables used in the example are found for monitoring statistics data.

Summary:

- 703 means MQIAMO64_AVG_Q_TIME which is the **AvgTimeOnQ**
- 739 means MQIAMO_Q_MAX_DEPTH which is the **QMaxDepth**

To find out the mapping you need to carry out two procedures:

1. Visit “[Queue statistics message data](#)” on [page 211](#) for an explanation of the attributes “[AvgTimeOnQ](#)” on [page 214](#) and “[QMaxDepth](#)” on [page 214](#) as you selected *statistics* for the **Type** parameter in the **amqsmn** command.

In each case, look for the attribute **Identifier**.

The **Identifier** for **AvgTimeOnQ** is MQIAMO64_AVG_Q_TIME and for **QMaxDepth** is MQIAMO_Q_MAX_DEPTH.

2. Search for the identifiers you found in Step “1” on [page 174](#).

Go to the [Constants](#) section and scroll down until you find the [MQIAMO_* \(Command format Integer Monitoring Parameter Types\)](#) list. Locate MQIAMO_Q_MAX_DEPTH and you see the value 739.

Similarly, find the [MQIAMO64_* \(Command format 64-bit Integer Monitoring Parameter Types\)](#) list. Locate MQIAMO64_AVG_Q_TIME and you see the value 703.

Accounting and statistics message reference

Use this page to obtain an overview of the format of accounting and statistics messages and the information returned in these messages

Accounting and statistics message messages are standard IBM MQ messages containing a message descriptor and message data. The message data contains information about the MQI operations performed by IBM MQ applications, or information about the activities occurring in an IBM MQ system.

Message descriptor

- An MQMD structure

Message data

- A PCF header (MQCFH)
- Accounting or statistics message data that is always returned
- Accounting or statistics message data that is returned if available

Accounting and statistics message format

Use this page as an example of the structure of an MQI accounting message

Table 27. MQI accounting message structure		
MQMD structure	Accounting message header MQCFH structure	MQI accounting message data ¹
Structure identifier Structure version Report options Message type Expiration time Feedback code Encoding Coded character set ID Message format Message priority Persistence Message identifier Correlation identifier Backout count Reply-to queue Reply-to queue manager User identifier Accounting token Application identity data Application type Application name Put date Put time Application origin data Group identifier Message sequence number Offset Message flags Original length	Structure type Structure length Structure version Command identifier Message sequence number Control options Completion code Reason code Parameter count	Queue manager Interval start date Interval start time Interval end date Interval end time Command level Connection identifier Sequence number Application name Application process identifier Application thread identifier User identifier Connection date Connection time Connection name Channel name Disconnect date Disconnect time Disconnect type Open count Open fail count Close count Close fail count Put count Put fail count Put1 count Put1 fail count Put bytes Get count Get fail count Get bytes Browse count Browse fail count Browse bytes Commit count Commit fail count Backout count Inquire count Inquire fail count Set count Set fail count
Note: 1. The parameters shown are those returned for an MQI accounting message. The actual accounting or statistics message data depends on the message category.		

Accounting and statistics message MQMD (message descriptor)

Use this page to understand the differences between the message descriptor of accounting and statistics messages and the message descriptor of event messages

The parameters and values in the message descriptor of accounting and statistics message are the same as in the message descriptor of event messages, with the following exception:

Format

Description:	Format name of message data.
Data type:	MQCHAR8.
Value:	MQFMT_ADMIN Admin message.

Some of the parameters contained in the message descriptor of accounting and statistics message contain fixed data supplied by the queue manager that generated the message.

The MQMD also specifies the name of the queue manager (truncated to 28 characters) that put the message, and the date and time when the message was put on the accounting, or statistics, queue.

Message data in accounting and statistics messages

The message data in accounting and statistics messages is based on the programmable command format (PCF), which is used in PCF command inquiries and responses. The message data in accounting and statistics messages consists of a PCF header (MQCFH) and an accounting or statistics report.

Accounting and statistics message MQCFH (PCF header)

The message header of accounting and statistics messages is an MQCFH structure. The parameters and values in the message header of accounting and statistics message are the same as in the message header of event messages, with the following exceptions:

Command

Description:	Command identifier. This identifies the accounting or statistics message category.
Data type:	MQLONG.
Values:	MQCMD_ACCOUNTING_MQI MQI accounting message. MQCMD_ACCOUNTING_Q Queue accounting message. MQCMD_STATISTICS_MQI MQI statistics message. MQCMD_STATISTICS_Q Queue statistics message. MQCMD_STATISTICS_CHANNEL Channel statistics message.

Version

Description:	Structure version number.
Data type:	MQLONG.
Value:	MQCFH_VERSION_3 Version-3 for accounting and statistics messages.

Accounting and statistics message data

The content of accounting and statistics message data is dependent on the category of the accounting or statistics message, as follows:

MQI accounting message

MQI accounting message data consists of a number of PCF parameters, but no PCF groups.

Queue accounting message

Queue accounting message data consists of a number of PCF parameters, and in the range 1 through 100 *QAccountingData* PCF groups.

MQI statistics message

MQI statistics message data consists of a number of PCF parameters, but no PCF groups.

Queue statistics message

Queue statistics message data consists of a number of PCF parameters, and in the range 1 through 100 *QStatisticsData* PCF groups.

Channel statistics message

Channel statistics message data consists of a number of PCF parameters, and in the range 1 through 100 *ChlStatisticsData* PCF groups.

MQI accounting message data

Use this page to view the structure of an MQI accounting message

Message name:	MQI accounting message.
Platforms:	All, except IBM MQ for z/OS.
System queue:	SYSTEM.ADMIN.ACCOUNTING.QUEUE.

QueueManager

Description:	The name of the queue manager
Identifier:	MQCA_Q_MGR_NAME
Data type:	MQCFST
Maximum length:	MQ_Q_MGR_NAME_LENGTH
Returned:	Always

IntervalStartDate

Description:	The date of the start of the monitoring period
Identifier:	MQCAMO_START_DATE
Data type:	MQCFST
Maximum length:	MQ_DATE_LENGTH
Returned:	Always

IntervalStartTime

Description:	The time of the start of the monitoring period
Identifier:	MQCAMO_START_TIME
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	Always

IntervalEndDate

Description:	The date of the end of the monitoring period
Identifier:	MQCAMO_END_DATE
Data type:	MQCFST
Maximum length:	MQ_DATE_LENGTH
Returned:	Always

IntervalEndTime

Description:	The time of the end of the monitoring period
Identifier:	MQCAMO_END_TIME
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	Always

CommandLevel

Description:	The queue manager command level
Identifier:	MQIA_COMMAND_LEVEL
Data type:	MQCFIN
Returned:	Always

ConnectionId

Description:	The connection identifier for the IBM MQ connection
Identifier:	MQBACF_CONNECTION_ID
Data type:	MQCFBS
Maximum length:	MQ_CONNECTION_ID_LENGTH
Returned:	Always

SeqNumber

Description:	The sequence number. This value is incremented for each subsequent record for long running connections.
Identifier:	MQIACF_SEQUENCE_NUMBER
Data type:	MQCFIN
Returned:	Always

ApplicationName

Description:	The name of the application. The contents of this field are equivalent to the contents of the <i>PutApplName</i> field in the message descriptor.
Identifier:	MQCACF_APPL_NAME
Data type:	MQCFST
Maximum length:	MQ_APPL_NAME_LENGTH
Returned:	Always

ApplicationPid

Description:	The operating system process identifier of the application
Identifier:	MQIACF_PROCESS_ID
Data type:	MQCFIN
Returned:	Always

ApplicationTid

Description:	The IBM MQ thread identifier of the connection in the application
Identifier:	MQIACF_THREAD_ID
Data type:	MQCFIN
Returned:	Always

UserId

Description:	The user identifier context of the application
Identifier:	MQCACF_USER_IDENTIFIER
Data type:	MQCFST
Maximum length:	MQ_USER_ID_LENGTH
Returned:	Always

ConnDate

Description:	Date of MQCONN operation
Identifier:	MQCAMO_CONN_DATE
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	When available

ConnTime

Description:	Time of MQCONN operation
Identifier:	MQCAMO_CONN_TIME
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	When available

ConnName

Description:	Connection name for client connection
Identifier:	MQCACH_CONNECTION_NAME
Data type:	MQCFST
Maximum length:	MQ_CONN_NAME_LENGTH
Returned:	When available

ChannelName

Description: Channel name for client connection
Identifier: MQCACH_CHANNEL_NAME
Data type: MQCFST
Maximum length: MQ_CHANNEL_NAME_LENGTH
Returned: When available

RemoteProduct

Description: Remote product identifier for client connection as displayed in the **RPRODUCT** field of [DISPLAY CHSATUS](#)
Identifier: MQCACH_REMOTE_PRODUCT
Data type: MQCFST
Maximum length: MQ_REMOTE_PRODUCT_LENGTH
Returned: When available

RemoteVersion

Description: Remote product version for client connection as displayed in the **RVERSION** field of [DISPLAY CHSTATUS](#)
Identifier: MQCACH_REMOTE_VERSION
Data type: MQCFST
Maximum length: MQ_REMOTE_VERSION_LENGTH
Returned: When available

DiscDate

Description: Date of MQDISC operation
Identifier: MQCAMO_DISC_DATE
Data type: MQCFST
Maximum length: MQ_DATE_LENGTH
Returned: When available

DiscTime

Description: Time of MQDISC operation
Identifier: MQCAMO_DISC_TIME
Data type: MQCFST
Maximum length: MQ_TIME_LENGTH
Returned: When available

DiscType

Description: Type of disconnect
Identifier: MQIAMO_DISC_TYPE
Data type: MQCFIN

Values: The possible values are:

- MQDISCONNECT_NORMAL**
Requested by application
- MQDISCONNECT_IMPLICIT**
Abnormal application termination
- MQDISCONNECT_Q_MGR**
Connection broken by queue manager

Returned: When available

OpenCount

Description: The number of objects successfully opened, either by directly issuing a call to MQOPEN or by using the MQPUT1 verb. This parameter is an integer list indexed by object type, see [Reference note 1](#).

Identifier: MQIAMO_OPENS

Data type: MQCFIL

Returned: When available

OpenFailCount

Description: The number of unsuccessful attempts to open an object. This parameter is an integer list indexed by object type, see [Reference note 1](#).

Identifier: MQIAMO_OPENS_FAILED

Data type: MQCFIL

Returned: When available

CloseCount

Description: The number of objects closed. This parameter is an integer list indexed by object type, see [Reference note 1](#).

Identifier: MQIAMO_CLOSES

Data type: MQCFIL

Returned: When available

CloseFailCount

Description: The number of unsuccessful attempts to close an object. This parameter is an integer list indexed by object type, see [Reference note 1](#).

Identifier: MQIAMO_CLOSES_FAILED

Data type: MQCFIL

Returned: When available

PutCount

Description: The number persistent and nonpersistent messages successfully put to a queue, with the exception of messages put using the MQPUT1 call. This parameter is an integer list indexed by persistence value, see [Reference note 2](#).

Identifier: MQIAMO_PUTS

Data type: MQCFIL

Returned: When available

PutFailCount

Description: The number of unsuccessful attempts to put a message

Identifier: MQIAMO_PUTS_FAILED

Data type: MQCFIN

Returned: When available

Put1Count

Description: The number of persistent and nonpersistent messages successfully put to the queue using MQPUT1 calls. This parameter is an integer list indexed by persistence value, see [Reference note 2](#).

Identifier: MQIAMO_PUT1S

Data type: MQCFIL

Included in PCF group: *QAccountingData*

Returned: When available

Put1FailCount

Description: The number of unsuccessful attempts to put a message using MQPUT1 calls

Identifier: MQIAMO_PUT1S_FAILED

Data type: MQCFIN

Included in PCF group: *QAccountingData*

Returned: When available

PutBytes

Description: The number bytes written using put calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see [Reference note 2](#).

Identifier: MQIAMO64_PUT_BYTES

Data type: MQCFIL64

Returned: When available

GetCount

Description: The number of successful destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see [Reference note 2](#).

Identifier: MQIAMO_GETS

Data type: MQCFIL

Returned: When available

GetFailCount

Description: The number of failed destructive MQGET calls

Identifier:	MQIAMO_GETS_FAILED
Data type:	MQCFIN
Returned:	When available

GetBytes

Description:	Total number of bytes retrieved for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_GET_BYTES
Data type:	MQCFIL64
Returned:	When available

BrowseCount

Description:	The number of successful non-destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_BROWSES
Data type:	MQCFIL
Returned:	When available

BrowseFailCount

Description:	The number of unsuccessful non-destructive MQGET calls
Identifier:	MQIAMO_BROWSES_FAILED
Data type:	MQCFIN
Returned:	When available

BrowseBytes

Description:	Total number of bytes browsed for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_BROWSE_BYTES
Data type:	MQCFIL64
Returned:	When available

CommitCount

Description:	The number of successful transactions. This number includes those transactions committed implicitly by the connected application. Commit requests where there is no outstanding work are included in this count.
Identifier:	MQIAMO_COMMITS
Data type:	MQCFIN
Returned:	When available

CommitFailCount

Description:	The number of unsuccessful attempts to complete a transaction
Identifier:	MQIAMO_COMMITS_FAILED

Data type:	MQCFIN
Returned:	When available

BackCount

Description:	The number of backouts processed, including implicit backouts due to abnormal disconnection
Identifier:	MQIAMO_BACKOUTS
Data type:	MQCFIN
Returned:	When available

InqCount

Description:	The number of successful objects inquired upon. This parameter is an integer list indexed by object type, see Reference note 1 .
Identifier:	MQIAMO_INQS
Data type:	MQCFIL
Returned:	When available

InqFailCount

Description:	The number of unsuccessful object inquire attempts. This parameter is an integer list indexed by object type, see Reference note 1 .
Identifier:	MQIAMO_INQS_FAILED
Data type:	MQCFIL
Returned:	When available

SetCount

Description:	The number of successful MQSET calls. This parameter is an integer list indexed by object type, see Reference note 1 .
Identifier:	MQIAMO_SETS
Data type:	MQCFIL
Returned:	When available

SetFailCount

Description:	The number of unsuccessful MQSET calls. This parameter is an integer list indexed by object type, see Reference note 1 .
Identifier:	MQIAMO_SETS_FAILED
Data type:	MQCFIL
Returned:	When available

SubCountDur

Description: The number of successful subscribe requests which created, altered or resumed durable subscriptions. This is an array of values indexed by the type of operation

- 0 = The number of subscriptions created
- 1 = The number of subscriptions altered
- 2 = The number of subscriptions resumed

Identifier: MQIAMO_SUBS_DUR

Data type: MQCFIL

Returned: When available.

SubCountNDur

Description: The number of successful subscribe requests which created, altered or resumed non-durable subscriptions. This is an array of values indexed by the type of operation

- 0 = The number of subscriptions created
- 1 = The number of subscriptions altered
- 2 = The number of subscriptions resumed

Identifier: MQIAMO_SUBS_NDUR

Data type: MQCFIL

Returned: When available.

SubFailCount

Description: The number of unsuccessful Subscribe requests.

Identifier: MQIAMO_SUBS_FAILED

Data type: MQCFIN

Returned: When available.

UnsubCountDur

Description: The number of successful unsubscribe requests for durable subscriptions. This is an array of values indexed by the type of operation

- 0 - The subscription was closed but not removed
- 1 - The subscription was closed and removed

Identifier: MQIAMO_UNSUBS_DUR

Data type: MQCFIL

Returned: When available.

UnsubCountNDur

Description: The number of successful unsubscribe requests for durable subscriptions. This is an array of values indexed by the type of operation

- 0 - The subscription was closed but not removed
- 1 - The subscription was closed and removed

Identifier:	MQIAMO_UNSUBS_NDUR
Data type:	MQCFIL
Returned:	When available.

UnsubFailCount

Description:	The number of unsuccessful unsubscribe requests.
Identifier:	MQIAMO_UNSUBS_FAILED
Data type:	MQCFIN
Returned:	When available.

SubRqCount

Description:	The number of successful MQSUBRQ requests.
Identifier:	MQIAMO_SUBRQS
Data type:	MQCFIN
Returned:	When available.

SubRqFailCount

Description:	The number of unsuccessful MQSUB requests.
Identifier:	MQIAMO_SUBRQS_FAILED
Data type:	MQCFIN
Returned:	When available.

CBCount

Description:	The number of successful MQCB requests. This is an array of values indexed by the type of operation 0 - A callback was created or altered 1 - A callback was removed 2 - A callback was resumed 3 - A callback was suspended
Identifier:	MQIAMO_CBS
Data type:	MQCFIN
Returned:	When available.

CBFailCount

Description:	The number of unsuccessful MQCB requests.
Identifier:	MQIAMO_CBS_FAILED
Data type:	MQCFIN
Returned:	When available.

CtlCount

Description:	The number of successful MQCTL requests. This is an array of values indexed by the type of operation 0 - The connection was started 1 - The connection was stopped 2 - The connection was resumed 3 - The connection was suspended
Identifier:	MQIAMO_CTL5
Data type:	MQCFIL
Returned:	When available.

CtlFailCount

Description:	The number of unsuccessful MQCTL requests.
Identifier:	MQIAMO_CTL5_FAILED
Data type:	MQCFIN
Returned:	When available.

StatCount

Description:	The number of successful MQSTAT requests.
Identifier:	MQIAMO_STATS.
Data type:	MQCFIN
Returned:	When available.

StatFailCount

Description:	The number of unsuccessful MQSTAT requests.
Identifier:	MQIAMO_STATS_FAILED
Data type:	MQCFIN
Returned:	When available.

PutTopicCount

Description:	The number persistent and nonpersistent messages successfully put to a topic, with the exception of messages put using the MQPUT1 call. This parameter is an integer list indexed by persistence value, see Reference note 2 . Note: Messages put using a queue alias which resolve to a topic are included in this value.
Identifier:	MQIAMO_TOPIC_PUTS
Data type:	MQCFIL
Returned:	When available.

PutTopicFailCount

Description:	The number of unsuccessful attempts to put a message to a topic.
Identifier:	MQIAMO_TOPIC_PUTS_FAILED

Data type:	MQCFIN
Returned:	When available.

Put1TopicCount

Description:	The number of persistent and nonpersistent messages successfully put to a topic using MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2 . Note: Messages put using a queue alias which resolve to a topic are included in this value.
Identifier:	MQIAMO_TOPIC_PUT1S
Data type:	MQCFIL
Returned:	When available.

Put1TopicFailCount

Description:	The number of unsuccessful attempts to put a message to a topic using MQPUT1 calls.
Identifier:	MQIAMO_TOPIC_PUT1S_FAILED
Data type:	MQCFIN
Returned:	When available.

PutTopicBytes

Description:	The number bytes written using put calls for persistent and nonpersistent messages which resolve to a publish operation. This is number of bytes put by the application and not the resultant number of bytes delivered to subscribers. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_TOPIC_PUT_BYTES
Data type:	MQCFIL64
Returned:	When available.

Queue accounting message data

Use this page to view the structure of a queue accounting message

Message name:	Queue accounting message.
Platforms:	All, except IBM MQ for z/OS.
System queue:	SYSTEM.ADMIN.ACCOUNTING.QUEUE.

QueueManager

Description:	The name of the queue manager
Identifier:	MQCA_Q_MGR_NAME
Data type:	MQCFST
Maximum length:	MQ_Q_MGR_NAME_LENGTH
Returned:	Always

IntervalStartDate

Description:	The date of the start of the monitoring period
Identifier:	MQCAMO_START_DATE
Data type:	MQCFST
Maximum length:	MQ_DATE_LENGTH
Returned:	Always

IntervalStartTime

Description:	The time of the start of the monitoring period
Identifier:	MQCAMO_START_TIME
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	Always

IntervalEndDate

Description:	The date of the end of the monitoring period
Identifier:	MQCAMO_END_DATE
Data type:	MQCFST
Maximum length:	MQ_DATE_LENGTH
Returned:	Always

IntervalEndTime

Description:	The time of the end of the monitoring period
Identifier:	MQCAMO_END_TIME
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	Always

CommandLevel

Description:	The queue manager command level
Identifier:	MQIA_COMMAND_LEVEL
Data type:	MQCFIN
Returned:	Always

ConnectionId

Description:	The connection identifier for the IBM MQ connection
Identifier:	MQBACF_CONNECTION_ID
Data type:	MQCFBS
Maximum length:	MQ_CONNECTION_ID_LENGTH
Returned:	Always

SeqNumber

Description:	The sequence number. This value is incremented for each subsequent record for long running connections.
Identifier:	MQIACF_SEQUENCE_NUMBER
Data type:	MQCFIN
Returned:	Always

ApplicationName

Description:	The name of the application. The contents of this field are equivalent to the contents of the PutApplName field in the message descriptor.
Identifier:	MQCACF_APPL_NAME
Data type:	MQCFST
Maximum length:	MQ_APPL_NAME_LENGTH
Returned:	Always

ApplicationPid

Description:	The operating system process identifier of the application
Identifier:	MQIACF_PROCESS_ID
Data type:	MQCFIN
Returned:	Always

ApplicationTid

Description:	The IBM MQ thread identifier of the connection in the application
Identifier:	MQIACF_THREAD_ID
Data type:	MQCFIN
Returned:	Always

UserId

Description:	The user identifier context of the application
Identifier:	MQCACF_USER_IDENTIFIER
Data type:	MQCFST
Maximum length:	MQ_USER_ID_LENGTH
Returned:	Always

ChannelName

Description:	Channel name for client connection
Identifier:	MQCACH_CHANNEL_NAME
Data type:	MQCFST
Maximum length:	MQ_CHANNEL_NAME_LENGTH
Returned:	When available

ConnName

Description:	Connection name for client connection
Identifier:	MQCACH_CONNECTION_NAME
Data type:	MQCFST
Maximum length:	MQ_CONN_NAME_LENGTH
Returned:	When available

ObjectCount

Description:	The number of queues accessed in the interval for which accounting data has been recorded. This value is set to the number of <i>QAccountingData</i> PCF groups contained in the message.
Identifier:	MQIAMO_OBJECT_COUNT
Data type:	MQCFIN
Returned:	Always

QAccountingData

Description:	Grouped parameters specifying accounting details for a queue
Identifier:	MQGACF_Q_ACCOUNTING_DATA
Data type:	MQCFGR

Parameters in group:	<i>QName</i> <i>CreateDate</i> <i>CreateTime</i> <i>QType</i> <i>QDefinitionType</i> <i>OpenCount</i> <i>OpenDate</i> <i>OpenTime</i> <i>CloseDate</i> <i>CloseTime</i> <i>PutCount</i> <i>PutFailCount</i> <i>Put1Count</i> <i>Put1FailCount</i> <i>PutBytes</i> <i>PutMinBytes</i> <i>PutMaxBytes</i> <i>GetCount</i> <i>GetFailCount</i> <i>GetBytes</i> <i>GetMinBytes</i> <i>GetMaxBytes</i> <i>BrowseCount</i> <i>BrowseFailCount</i> <i>BrowseBytes</i> <i>BrowseMinBytes</i> <i>BrowseMaxBytes</i> <i>TimeOnQMin</i> <i>TimeOnQAvg</i> <i>TimeOnQMax</i>
----------------------	---

Returned:	Always
-----------	--------

QName

Description:	The name of the queue
Identifier:	MQCA_Q_NAME
Data type:	MQCFST
Included in PCF group:	<i>QAccountingData</i>
Maximum length:	MQ_Q_NAME_LENGTH
Returned:	When available

CreateDate

Description:	The date the queue was created
Identifier:	MQCA_CREATION_DATE
Data type:	MQCFST
Included in PCF group:	<i>QAccountingData</i>

Maximum length: MQ_DATE_LENGTH
Returned: When available

CreateTime

Description: The time the queue was created
Identifier: MQCA_CREATION_TIME
Data type: MQCFST
Included in PCF group: *QAccountingData*
Maximum length: MQ_TIME_LENGTH
Returned: When available

QType

Description: The type of the queue
Identifier: MQIA_Q_TYPE
Data type: MQCFIN
Included in PCF group: *QAccountingData*
Value: MQQT_LOCAL
Returned: When available

QDefinitionType

Description: The queue definition type
Identifier: MQIA_DEFINITION_TYPE
Data type: MQCFIN
Included in PCF group: *QAccountingData*
Values: Possible values are:
MQQDT_PREDEFINED
MQQDT_PERMANENT_DYNAMIC
MQQDT_TEMPORARY_DYNAMIC
Returned: When available

OpenCount

Description: The number of times this queue was opened by the application in this interval, either by directly issuing a call to MQOPEN or by using the MQPUT1 verb.
Identifier: MQIAMO_OPENS
Data type: MQCFIL
Included in PCF group: *QAccountingData*
Returned: When available

OpenDate

Description:	The date the queue was first opened in this recording interval. If the queue was already open at the start of this interval, this value reflects the date the queue was originally opened.
Identifier:	MQCAMO_OPEN_DATE
Data type:	MQCFST
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

OpenTime

Description:	The time the queue was first opened in this recording interval. If the queue was already open at the start of this interval, this value reflects the time the queue was originally opened.
Identifier:	MQCAMO_OPEN_TIME
Data type:	MQCFST
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

CloseDate

Description:	The date of the final close of the queue in this recording interval. If the queue is still open then the value is not returned.
Identifier:	MQCAMO_CLOSE_DATE
Data type:	MQCFST
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

CloseTime

Description:	The time of final close of the queue in this recording interval. If the queue is still open then the value is not returned.
Identifier:	MQCAMO_CLOSE_TIME
Data type:	MQCFST
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

PutCount

Description:	The number of persistent and nonpersistent messages successfully put to the queue, with the exception of MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_PUTS
Data type:	MQCFIL

Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

PutFailCount

Description:	The number of unsuccessful attempts to put a message, with the exception of MQPUT1 calls
Identifier:	MQIAMO_PUTS_FAILED
Data type:	MQCFIN
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

Put1Count

Description:	The number of persistent and nonpersistent messages successfully put to the queue using MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_PUT1S
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

Put1FailCount

Description:	The number of unsuccessful attempts to put a message using MQPUT1 calls
Identifier:	MQIAMO_PUT1S_FAILED
Data type:	MQCFIN
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

PutBytes

Description:	The total number of bytes put for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_PUT_BYTES
Data type:	MQCFIL64
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

PutMinBytes

Description:	The smallest persistent and nonpersistent message size placed on the queue. This parameter is an integer list indexed by persistence value, see Reference note 2 .
--------------	--

Identifier:	MQIAMO_PUT_MIN_BYTES
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

PutMaxBytes

Description:	The largest persistent and nonpersistent message size placed on the queue. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_PUT_MAX_BYTES
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

GeneratedMsgCount

Description:	The number of generated messages. Generated messages are • Queue Depth Hi Events • Queue Depth Low Events
Identifier:	MQIAMO_GENERATED_MSGS
Data type:	MQCFIN
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

GetCount

Description:	The number of successful destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_GETS
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

GetFailCount

Description:	The number of failed destructive MQGET calls
Identifier:	MQIAMO_GETS_FAILED
Data type:	MQCFIN
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

GetBytes

Description:	The number of bytes read in destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_GET_BYTES
Data type:	MQCFIL64
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

GetMinBytes

Description:	The size of the smallest persistent and nonpersistent message retrieved from the queue. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_GET_MIN_BYTES
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

GetMaxBytes

Description:	The size of the largest persistent and nonpersistent message retrieved from the queue. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_GET_MAX_BYTES
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

BrowseCount

Description:	The number of successful non-destructive MQGET calls for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_BROWSES
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

BrowseFailCount

Description:	The number of unsuccessful non-destructive MQGET calls
Identifier:	MQIAMO_BROWSES_FAILED
Data type:	MQCFIN

Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

BrowseBytes

Description:	The number of bytes read in non-destructive MQGET calls that returned persistent messages
Identifier:	MQIAMO64_BROWSE_BYTES
Data type:	MQCFIL64
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

BrowseMinBytes

Description:	The size of the smallest persistent and nonpersistent message browsed from the queue. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_BROWSE_MIN_BYTES
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

BrowseMaxBytes

Description:	The size of the largest persistent and nonpersistent message browsed from the queue. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_BROWSE_MAX_BYTES
Data type:	MQCFIL
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

TimeOnQMin

Description:	The shortest time a persistent and nonpersistent message remained on the queue before being destructively retrieved, in microseconds. For messages retrieved under syncpoint this value does not include the time before the get operation is committed. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_Q_TIME_MIN
Data type:	MQCFIL64
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

TimeOnQAvg

Description:	The average time a persistent and nonpersistent message remained on the queue before being destructively retrieved, in microseconds. For messages retrieved under syncpoint this value does not include the time before the get operation is committed. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_Q_TIME_AVG
Data type:	MQCFIL64
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

TimeOnQMax

Description:	The longest time a persistent and nonpersistent message remained on the queue before being destructively retrieved, in microseconds. For messages retrieved under syncpoint this value does not include the time before the get operation is committed. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_Q_TIME_MAX
Data type:	MQCFIL64
Included in PCF group:	<i>QAccountingData</i>
Returned:	When available

CBCount

Description:	The number of successful MQCB requests. This is an array of values indexed by the type of operation 0 - A callback was created or altered 1 - A callback was removed 2 - A callback was resumed 3 - A callback was suspended
Identifier:	MQIAMO_CBS
Data type:	MQCFIN
Returned:	When available.

CBFailCount

Description:	The number of unsuccessful MQCB requests.
Identifier:	MQIAMO_CBS_FAILED
Data type:	MQCFIN
Returned:	When available.

MQI statistics message data

Use this page to view the structure of an MQI statistics message

Message name: MQI statistics message.

Platforms: All, except IBM MQ for z/OS.

System queue: SYSTEM.ADMIN.STATISTICS.QUEUE.

QueueManager

Description: Name of the queue manager.
Identifier: MQCA_Q_MGR_NAME.
Data type: MQCFST.
Maximum length: MQ_Q_MGR_NAME_LENGTH.
Returned: Always.

IntervalStartDate

Description: The date at the start of the monitoring period.
Identifier: MQCAMO_START_DATE.
Data type: MQCFST.
Maximum length: MQ_DATE_LENGTH
Returned: Always.

IntervalStartTime

Description: The time at the start of the monitoring period.
Identifier: MQCAMO_START_TIME.
Data type: MQCFST.
Maximum length: MQ_TIME_LENGTH
Returned: Always.

IntervalEndDate

Description: The date at the end of the monitoring period.
Identifier: MQCAMO_END_DATE.
Data type: MQCFST.
Maximum length: MQ_DATE_LENGTH
Returned: Always.

IntervalEndTime

Description: The time at the end of the monitoring period.
Identifier: MQCAMO_END_TIME.
Data type: MQCFST.
Maximum length: MQ_TIME_LENGTH
Returned: Always.

CommandLevel

Description: The queue manager command level.
Identifier: MQIA_COMMAND_LEVEL.

Data type: MQCFIN.
Returned: Always.

ConnCount

Description: The number of successful connections to the queue manager.
Identifier: MQIAMO_CONNS.
Data type: MQCFIN.
Returned: When available.

ConnFailCount

Description: The number of unsuccessful connection attempts.
Identifier: MQIAMO_CONNS_FAILED.
Data type: MQCFIN.
Returned: When available.

ConnsMax

Description: The maximum number of concurrent connections in the recording interval.
Identifier: MQIAMO_CONNS_MAX.
Data type: MQCFIN.
Returned: When available.

DiscCount

Description: The number of disconnects from the queue manager. This is an integer array, indexed by the following constants:

- MQDISCONNECT_NORMAL
- MQDISCONNECT_IMPLICIT
- MQDISCONNECT_Q_MGR

Identifier: MQIAMO_DISCS.
Data type: MQCFIL.
Returned: When available.

OpenCount

Description: The number of objects successfully opened, either by directly issuing a call to MQOPEN or by using the MQPUT1 verb. This parameter is an integer list indexed by object type, see [Reference note 1](#).
Identifier: MQIAMO_OPENS.
Data type: MQCFIL.
Returned: When available.

OpenFailCount

Description: The number of unsuccessful open object attempts. This parameter is an integer list indexed by object type, see [Reference note 1](#).
Identifier: MQIAMO_OPENS_FAILED.

Data type: MQCFIL.
Returned: When available.

CloseCount

Description: The number of objects successfully closed. This parameter is an integer list indexed by object type, see [Reference note 1](#).
Identifier: MQIAMO_CLOSES.
Data type: MQCFIL.
Returned: When available.

CloseFailCount

Description: The number of unsuccessful close object attempts. This parameter is an integer list indexed by object type, see [Reference note 1](#).
Identifier: MQIAMO_CLOSES_FAILED.
Data type: MQCFIL.
Returned: When available.

InqCount

Description: The number of objects successfully inquired upon. This parameter is an integer list indexed by object type, see [Reference note 1](#).
Identifier: MQIAMO_INQS.
Data type: MQCFIL.
Returned: When available.

InqFailCount

Description: The number of unsuccessful object inquire attempts. This parameter is an integer list indexed by object type, see [Reference note 1](#).
Identifier: MQIAMO_INQS_FAILED.
Data type: MQCFIL.
Returned: When available.

SetCount

Description: The number of objects successfully updated (SET). This parameter is an integer list indexed by object type, see [Reference note 1](#).
Identifier: MQIAMO_SETS.
Data type: MQCFIL.
Returned: When available.

SetFailCount

Description: The number of unsuccessful SET attempts. This parameter is an integer list indexed by object type, see [Reference note 1](#).
Identifier: MQIAMO_SETS_FAILED.
Data type: MQCFIL.

Returned: When available.

PutCount

Description: The number of persistent and nonpersistent messages successfully put to a queue, with the exception of MQPUT1 requests. This parameter is an integer list indexed by persistence value, see [Reference note 2](#).

Identifier: MQIAMO_PUTS.

Data type: MQCFIL.

Returned: When available.

PutFailCount

Description: The number of unsuccessful put message attempts.

Identifier: MQIAMO_PUTS_FAILED.

Data type: MQCFIN.

Returned: When available.

Put1Count

Description: The number of persistent and nonpersistent messages successfully put to a queue using MQPUT1 requests. This parameter is an integer list indexed by persistence value, see [Reference note 2](#)

Identifier: MQIAMO_PUT1S.

Data type: MQCFIL.

Returned: When available.

Put1FailCount

Description: The number of unsuccessful attempts to put a persistent and nonpersistent message to a queue using MQPUT1 requests. This parameter is an integer list indexed by persistence value, see [Reference note 2](#)

Identifier: MQIAMO_PUT1S_FAILED.

Data type: MQCFIL.

Returned: When available.

PutBytes

Description: The number bytes for persistent and nonpersistent messages written in using put requests. This parameter is an integer list indexed by persistence value, see [Reference note 2](#)

Identifier: MQIAMO64_PUT_BYTES.

Data type: MQCFIL64.

Returned: When available.

GetCount

Description: The number of successful destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see [Reference note 2](#)

Identifier: MQIAMO_GETS.

Data type: MQCFIL.
Returned: When available.

GetFailCount

Description: The number of unsuccessful destructive get requests.
Identifier: MQIAMO_GETS_FAILED.
Data type: MQCFIN.
Returned: When available.

GetBytes

Description: The number of bytes read in destructive gets requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see [Reference note 2](#)
Identifier: MQIAMO64_GET_BYTES.
Data type: MQCFIL64.
Returned: When available.

BrowseCount

Description: The number of successful non-destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see [Reference note 2](#)
Identifier: MQIAMO_BROWSES.
Data type: MQCFIL.
Returned: When available.

BrowseFailCount

Description: The number of unsuccessful non-destructive get requests.
Identifier: MQIAMO_BROWSES_FAILED.
Data type: MQCFIN.
Returned: When available.

BrowseBytes

Description: The number of bytes read in non-destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value, see [Reference note 2](#)
Identifier: MQIAMO64_BROWSE_BYTES.
Data type: MQCFIL64.
Returned: When available.

CommitCount

Description: The number of transactions successfully completed. This number includes transactions committed implicitly by the application disconnecting, and commit requests where there is no outstanding work.
Identifier: MQIAMO_COMMITS.

Data type: MQCFIN.
Returned: When available.

CommitFailCount

Description: The number of unsuccessful attempts to complete a transaction.
Identifier: MQIAMO_COMMITS_FAILED.
Data type: MQCFIN.
Returned: When available.

BackCount

Description: The number of backouts processed, including implicit backout upon abnormal disconnect.
Identifier: MQIAMO_BACKOUTS.
Data type: MQCFIN.
Returned: When available.

ExpiredMsgCount

Description: The number of persistent and nonpersistent messages that were discarded because they had expired, before they could be retrieved.
Identifier: MQIAMO_MSGS_EXPIRED.
Data type: MQCFIN.
Returned: When available.

PurgeCount

Description: The number of times the queue has been cleared.
Identifier: MQIAMO_MSGS_PURGED.
Data type: MQCFIN.
Returned: When available.

SubCountDur

Description: The number of successful Subscribe requests which created, altered or resumed durable subscriptions. This is an array of values indexed by the type of operation
0 = The number of subscriptions created
1 = The number of subscriptions altered
2 = The number of subscriptions resumed
Identifier: MQIAMO_SUBS_DUR.
Data type: MQCFIL
Returned: When available.

SubCountNDur

Description:	The number of successful Subscribe requests which created, altered or resumed non-durable subscriptions. This is an array of values indexed by the type of operation 0 = The number of subscriptions created 1 = The number of subscriptions altered 2 = The number of subscriptions resumed
Identifier:	MQIAMO_SUBS_NDUR.
Data type:	MQCFIL.
Returned:	When available.

SubFailCount

Description:	The number of unsuccessful Subscribe requests.
Identifier:	MQIAMO_SUBS_FAILED.
Data type:	MQCFIN.
Returned:	When available.

UnsubCountDur

Description:	The number of successful unsubscribe requests for durable subscriptions. This is an array of values indexed by the type of operation 0 - The subscription was closed but not removed 1 - The subscription was closed and removed
Identifier:	MQIAMO_UNSUBS_DUR.
Data type:	MQCFIL.
Returned:	When available.

UnsubCountNDur

Description:	The number of successful unsubscribe requests for non-durable subscriptions. This is an array of values indexed by the type of operation 0 - The subscription was closed but not removed 1 - The subscription was closed and removed
Identifier:	MQIAMO_UNSUBS_NDUR.
Data type:	MQCFIL.
Returned:	When available.

UnsubFailCount

Description:	The number of failed unsubscribe requests.
Identifier:	MQIAMO_UNSUBS_FAILED.
Data type:	MQCFIN.
Returned:	When available.

SubRqCount

Description: The number of successful MQSUBRQ requests.
Identifier: MQIAMO_SUBRQS
Data type: MQCFIN
Returned: When available.

SubRqFailCount

Description: The number of unsuccessful MQSUBRQ requests.
Identifier: MQIAMO_SUBRQS_FAILED.
Data type: MQCFIN.
Returned: When available.

CBCount

Description: The number of successful MQCB requests. This is an array of values indexed by the type of operation
0 - A callback was created or altered
1 - A callback was removed
2 - A callback was resumed
3 - A callback was suspended
Identifier: MQIAMO_CBS.
Data type: MQCFIL.
Returned: When available.

CBFailCount

Description: The number of unsuccessful MQCB requests.
Identifier: MQIAMO_CBS_FAILED.
Data type: MQCFIN.
Returned: When available.

CtlCount

Description: The number of successful MQCTL requests. This is an array of values indexed by the type of operation:
0 - The connection was started
1 - The connection was stopped
2 - The connection was resumed
3 - The connection was suspended
Identifier: MQIAMO_CTLs.
Data type: MQCFIL.
Returned: When available.

CtlFailCount

Description: The number of unsuccessful MQCTL requests.
Identifier: MQIAMO_CTLS_FAILED.
Data type: MQCFIN.
Returned: When available.

StatCount

Description: The number of successful MQSTAT requests.
Identifier: MQIAMO_STATS.
Data type: MQCFIN.
Returned: When available.

StatFailCount

Description: The number of unsuccessful MQSTAT requests.
Identifier: MQIAMO_STATS_FAILED.
Data type: MQCFIN.
Returned: When available.

SubCountDurHighWater

Description: The high-water mark on the number of durable subscriptions during the time interval. This is an array of values indexed by SUBTYPE
0 - The high-water mark for all durable subscriptions in the system
1 - The high-water mark for durable application subscriptions (MQSUBTYPE_API)
2 - The high-water mark for durable admin subscription (MQSUBTYPE_ADMIN)
3 - The high-water mark for durable proxy subscriptions (MQSUBTYPE_PROXY)
Identifier: MQIAMO_SUB_DUR_HIGHWATER
Data type: MQCFIL.
Returned: When available.

SubCountDurLowWater

Description: The low-water mark on the number of durable subscriptions during the time interval. This is an array of values indexed by SUBTYPE.
0 - The low-water mark for all durable subscriptions in the system
1 - The low-water mark for durable application subscriptions (MQSUBTYPE_API)
2 - The low-water mark for durable admin subscriptions (MQSUBTYPE_ADMIN)
3 - The low-water mark for durable proxy subscriptions (MQSUBTYPE_PROXY)
Identifier: MQIAMO_SUB_DUR_LOWWATER
Data type: MQCFIL.
Returned: When available.

SubCountNDurHighWater

Description: The high-water mark on the number of non-durable subscriptions during the time interval. This is an array of values indexed by SUBTYPE

- 0 - The high-water mark for all non-durable subscriptions in the system
- 1 - The high-water mark for non-durable application subscriptions (MQSUBTYPE_API)
- 2 - The high-water mark for non-durable admin subscription (MQSUBTYPE_ADMIN)
- 3 - The high-water mark for non-durable proxy subscriptions (MQSUBTYPE_PROXY)

Identifier: MQIAMO_SUB_NDUR_HIGHWATER

Data type: MQCFIL.

Returned: When available.

SubCountNDurLowWater

Description: The low-water mark on the number of non-durable subscriptions during the time interval. This is an array of values indexed by SUBTYPE.

- 0 - The low-water mark for all non-durable subscriptions in the system
- 1 - The low-water mark for non-durable application subscriptions (MQSUBTYPE_API)
- 2 - The low-water mark for non-durable admin subscriptions (MQSUBTYPE_ADMIN)
- 3 - The low-water mark for non-durable proxy subscriptions (MQSUBTYPE_PROXY)

Identifier: MQIAMO_SUB_NDUR_LOWWATER

Data type: MQCFIL.

Returned: When available.

PutTopicCount

Description: The number persistent and nonpersistent messages successfully put to a topic, with the exception of messages put using the MQPUT1 call. This parameter is an integer list indexed by persistence value, see [Reference note 2](#).

Note: Messages put using a queue alias which resolve to a topic are included in this value.

Identifier: MQIAMO_TOPIC_PUTS.

Data type: MQCFIL.

Returned: When available.

PutTopicFailCount

Description: The number of unsuccessful attempts to put a message to a topic.

Identifier: MQIAMO_TOPIC_PUTS_FAILED.

Data type: MQCFIN.

Returned: When available.

Put1TopicCount

Description:	The number of persistent and nonpersistent messages successfully put to a topic using MQPUT1 calls. This parameter is an integer list indexed by persistence value, see Reference note 2 . Note: Messages put using a queue alias which resolve to a topic are included in this value.
Identifier:	MQIAMO_TOPIC_PUT1S.
Data type:	MQCFIL.
Returned:	When available.

Put1TopicFailCount

Description:	The number of unsuccessful attempts to put a message to a topic using MQPUT1 calls.
Identifier:	MQIAMO_TOPIC_PUT1S_FAILED.
Data type:	MQCFIN.
Returned:	When available.

PutTopicBytes

Description:	The number bytes written using put calls for persistent and nonpersistent messages which resolve to a publish operation. This is number of bytes put by the application and not the resultant number of bytes delivered to subscribers, see PublishMsgBytes for this value. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_TOPIC_PUT_BYTES.
Data type:	MQCFIL64.
Returned:	When available.

PublishMsgCount

Description:	The number of messages delivered to subscriptions in the time interval. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO_PUBLISH_MSG_COUNT
Data type:	MQCFIL.
Returned:	When available.

PublishMsgBytes

Description:	The number of bytes delivered to subscriptions in the time interval. This parameter is an integer list indexed by persistence value, see Reference note 2 .
Identifier:	MQIAMO64_PUBLISH_MSG_BYTES
Data type:	MQCFIL64.
Returned:	When available.

Queue statistics message data

Use this page to view the structure of a queue statistics message

Message name:	Queue statistics message.
Platforms:	All, except IBM MQ for z/OS.
System queue:	SYSTEM.ADMIN.STATISTICS.QUEUE.

QueueManager

Description:	Name of the queue manager
Identifier:	MQCA_Q_MGR_NAME
Data type:	MQCFST
Maximum length:	MQ_Q_MGR_NAME_LENGTH
Returned:	Always

IntervalStartDate

Description:	The date at the start of the monitoring period
Identifier:	MQCAMO_START_DATE
Data type:	MQCFST
Maximum length:	MQ_DATE_LENGTH
Returned:	Always

IntervalStartTime

Description:	The time at the start of the monitoring period
Identifier:	MQCAMO_START_TIME
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	Always

IntervalEndDate

Description:	The date at the end of the monitoring period
Identifier:	MQCAMO_END_DATE
Data type:	MQCFST
Maximum length:	MQ_DATE_LENGTH
Returned:	Always

IntervalEndTime

Description:	The time at the end of the monitoring period
Identifier:	MQCAMO_END_TIME
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	Always

CommandLevel

Description:	The queue manager command level
Identifier:	MQIA_COMMAND_LEVEL
Data type:	MQCFIN
Returned:	Always

ObjectCount

Description:	The number of queue objects accessed in the interval for which statistics data has been recorded. This value is set to the number of QStatisticsData PCF groups contained in the message.
Identifier:	MQIAMO_OBJECT_COUNT
Data type:	MQCFIN
Returned:	Always

QStatisticsData

Description:	Grouped parameters specifying statistics details for a queue
Identifier:	MQGACF_Q_STATISTICS_DATA
Data type:	MQCFGR
Parameters in group:	<i>QName</i> <i>CreateDate</i> <i>CreateTime</i> <i>QType</i> <i>QDefinitionType</i> <i>QMinDepth</i> <i>QMaxDepth</i> <i>AvgTimeOnQ</i> <i>PutCount</i> <i>PutFailCount</i> <i>Put1Count</i> <i>Put1FailCount</i> <i>PutBytes</i> <i>GetCount</i> <i>GetFailCount</i> <i>GetBytes</i> <i>BrowseCount</i> <i>BrowseFailCount</i> <i>BrowseBytes</i> <i>NonQueuedMsgCount</i> <i>ExpiredMsgCount</i> <i>PurgeCount</i>
Returned:	Always

QName

Description:	The name of the queue
Identifier:	MQCA_Q_NAME

Data type:	MQCFST
Maximum length:	MQ_Q_NAME_LENGTH
Returned:	Always

CreateDate

Description:	The date when the queue was created
Identifier:	MQCA_CREATION_DATE
Data type:	MQCFST
Maximum length:	MQ_DATE_LENGTH
Returned:	Always

CreateTime

Description:	The time when the queue was created
Identifier:	MQCA_CREATION_TIME
Data type:	MQCFST
Maximum length:	MQ_TIME_LENGTH
Returned:	Always

QType

Description:	The type of the queue
Identifier:	MQIA_Q_TYPE
Data type:	MQCFIN
Value:	MQOT_LOCAL
Returned:	Always

QDefinitionType

Description:	The queue definition type
Identifier:	MQIA_DEFINITION_TYPE
Data type:	MQCFIN
Values:	Possible values are • MQQDT_PREDEFINED • MQQDT_PERMANENT_DYNAMIC • MQQDT_TEMPORARY_DYNAMIC
Returned:	When available

QMinDepth

Description:	The minimum queue depth during the monitoring period
Identifier:	MQIAMO_Q_MIN_DEPTH
Data type:	MQCFIN
Included in PCF group:	<i>QStatisticsData</i>

Returned: When available

QMaxDepth

Description: The maximum queue depth during the monitoring period

Identifier: MQIAMO_Q_MAX_DEPTH

Data type: MQCFIN

Included in PCF group: *QStatisticsData*

Returned: When available

AvgTimeOnQ

Description: The average latency, in microseconds, of messages destructively retrieved from the queue during the monitoring period. This parameter is an integer list indexed by persistence value, see [Reference note 2](#).

Identifier: MQIAMO64_AVG_Q_TIME

Data type: MQCFIL64

Included in PCF group: *QStatisticsData*

Returned: When available

PutCount

Description: The number of persistent and nonpersistent messages successfully put to the queue, with exception of MQPUT1 requests. This parameter is an integer list indexed by persistence value. See [Reference note 2](#).

Identifier: MQIAMO_PUTS

Data type: MQCFIL

Included in PCF group: *QStatisticsData*

Returned: When available

PutFailCount

Description: The number of unsuccessful attempts to put a message to the queue

Identifier: MQIAMO_PUTS_FAILED

Data type: MQCFIN

Included in PCF group: *QStatisticsData*

Returned: When available

Put1Count

Description: The number of persistent and nonpersistent messages successfully put to the queue using MQPUT1 calls. This parameter is an integer list indexed by persistence value. See [Reference note 2](#).

Identifier: MQIAMO_PUT1S

Data type: MQCFIL

Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

Put1FailCount

Description:	The number of unsuccessful attempts to put a message using MQPUT1 calls
Identifier:	MQIAMO_PUT1S_FAILED
Data type:	MQCFIN
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

PutBytes

Description:	The number of bytes written in put requests to the queue
Identifier:	MQIAMO64_PUT_BYTES
Data type:	MQCFIL64
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

GetCount

Description:	The number of successful destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value. See Reference note 2 .
Identifier:	MQIAMO_GETS
Data type:	MQCFIL
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

GetFailCount

Description:	The number of unsuccessful destructive get requests
Identifier:	MQIAMO_GETS_FAILED
Data type:	MQCFIN
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

GetBytes

Description:	The number of bytes read in destructive put requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value. See Reference note 2 .
Identifier:	MQIAMO64_GET_BYTES
Data type:	MQCFIL64

Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

BrowseCount

Description:	The number of successful non-destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value. See Reference note 2 .
Identifier:	MQIAMO_BROWSES
Data type:	MQCFIL
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

BrowseFailCount

Description:	The number of unsuccessful non-destructive get requests
Identifier:	MQIAMO_BROWSES_FAILED
Data type:	MQCFIN
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

BrowseBytes

Description:	The number of bytes read in non-destructive get requests for persistent and nonpersistent messages. This parameter is an integer list indexed by persistence value. See Reference note 2 .
Identifier:	MQIAMO64_BROWSE_BYTES
Data type:	MQCFIL64
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

NonQueuedMsgCount

Description:	The number of messages that bypassed the queue and were transferred directly to a waiting application. Bypassing a queue can only occur in certain circumstances. This number represents how many times IBM MQ was able to bypass the queue, and not the number of times an application was waiting.
Identifier:	MQIAMO_MSGS_NOT_QUEUED
Data type:	MQCFIN
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

ExpiredMsgCount

Description:	The number of persistent and nonpersistent messages that were discarded because they had expired before they could be retrieved.
Identifier:	MQIAMO_MSGS_EXPIRED
Data type:	MQCFIN
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

PurgeCount

Description:	The number of messages purged.
Identifier:	MQIAMO_MSGS_PURGED
Data type:	MQCFIN
Included in PCF group:	<i>QStatisticsData</i>
Returned:	When available

Channel statistics message data

Use this page to view the structure of a channel statistics message

Message name:	Channel statistics message.
Platforms:	All, except IBM MQ for z/OS.
System queue:	SYSTEM.ADMIN.STATISTICS.QUEUE.

QueueManager

Description:	The name of the queue manager.
Identifier:	MQCA_Q_MGR_NAME.
Data type:	MQCFST.
Maximum length:	MQ_Q_MGR_NAME_LENGTH.
Returned:	Always.

IntervalStartDate

Description:	The date at the start of the monitoring period.
Identifier:	MQCAMO_START_DATE.
Data type:	MQCFST.
Maximum length:	MQ_DATE_LENGTH.
Returned:	Always.

IntervalStartTime

Description:	The time at the start of the monitoring period.
Identifier:	MQCAMO_START_TIME.
Data type:	MQCFST.
Maximum length:	MQ_TIME_LENGTH.

Returned: Always.

IntervalEndDate

Description: The date at the end of the monitoring period

Identifier: MQCAMO_END_DATE.

Data type: MQCFST.

Maximum length: MQ_DATE_LENGTH.

Returned: Always.

IntervalEndTime

Description: The time at the end of the monitoring period

Identifier: MQCAMO_END_TIME.

Data type: MQCFST.

Maximum length: MQ_TIME_LENGTH

Returned: Always.

CommandLevel

Description: The queue manager command level.

Identifier: MQIA_COMMAND_LEVEL.

Data type: MQCFIN.

Returned: Always.

ObjectCount

Description: The number of Channel objects accessed in the interval for which statistics data has been recorded. This value is set to the number of ChlStatisticsData PCF groups contained in the message.

Identifier: MQIAMO_OBJECT_COUNT

Data type: MQCFIN.

Returned: Always.

ChlStatisticsData

Description: Grouped parameters specifying statistics details for a channel.

Identifier: MQGACF_CHL_STATISTICS_DATA.

Data type: MQCFGR.

Parameters in group:	<i>ChannelName</i> <i>ChannelType</i> <i>RemoteQmgr</i> <i>ConnectionName</i> <i>MsgCount</i> <i>TotalBytes</i> <i>NetTimeMin</i> <i>NetTimeAvg</i> <i>NetTimeMax</i> <i>ExitTimeMin</i> <i>ExitTimeAvg</i> <i>ExitTimeMax</i> <i>FullBatchCount</i> <i>IncplBatchCount</i> <i>AverageBatchSize</i> <i>PutRetryCount</i>
----------------------	---

Returned:	Always.
-----------	---------

ChannelName

Description:	The name of the channel.
Identifier:	MQCACH_CHANNEL_NAME.
Data type:	MQCFST.
Maximum length:	MQ_CHANNEL_NAME_LENGTH.
Returned:	Always.

ChannelType

Description:	The channel type.
Identifier:	MQIACH_CHANNEL_TYPE.
Data type:	MQCFIN.
Values:	Possible values are: MQCHT_SENDER Sender channel. MQCHT_SERVER Server channel. MQCHT_RECEIVER Receiver channel. MQCHT_REQUESTER Requester channel. MQCHT_CLUSRCVR Cluster receiver channel. MQCHT_CLUSSDR Cluster sender channel.
Returned:	Always.

RemoteQmgr

Description:	The name of the remote queue manager.
--------------	---------------------------------------

Identifier: MQCA_REMOTE_Q_MGR_NAME.
Data type: MQCFST.
Maximum length: MQ_Q_MGR_NAME_LENGTH
Returned: When available.

ConnectionName

Description: Connection name of remote queue manager.
Identifier: MQCACH_CONNECTION_NAME.
Data type: MQCFST
Maximum length: MQ_CONN_NAME_LENGTH
Returned: When available.

MsgCount

Description: The number of persistent and nonpersistent messages sent or received.
Identifier: MQIAMO_MSGS.
Data type: MQCFIN
Returned: When available.

TotalBytes

Description: The number of bytes sent or received for persistent and nonpersistent messages.
Identifier: MQIAMO64_BYTES.
Data type: MQCFIN64.
Returned: When available.

NetTimeMin

Description: The shortest recorded channel round trip measured in the recording interval, in microseconds.
Identifier: MQIAMO_NET_TIME_MIN.
Data type: MQCFIN.
Returned: When available.

NetTimeAvg

Description: The average recorded channel round trip measured in the recording interval, in microseconds.
Identifier: MQIAMO_NET_TIME_AVG.
Data type: MQCFIN.
Returned: When available.

NetTimeMax

Description: The longest recorded channel round trip measured in the recording interval, in microseconds.
Identifier: MQIAMO_NET_TIME_MAX.

Data type: MQCFIN.
Returned: When available.

ExitTimeMin

Description: The shortest recorded time, in microseconds, spent executing a user exit in the recording interval,
Identifier: MQIAMO_EXIT_TIME_MIN.
Data type: MQCFIN.
Returned: When available.

ExitTimeAvg

Description: The average recorded time, in microseconds, spent executing a user exit in the recording interval. Measured in microseconds.
Identifier: MQIAMO_EXIT_TIME_AVG.
Data type: MQCFIN.
Returned: When available.

ExitTimeMax

Description: The longest recorded time, in microseconds, spent executing a user exit in the recording interval. Measured in microseconds.
Identifier: MQIAMO_EXIT_TIME_MAX.
Data type: MQCFIN.
Returned: When available.

FullBatchCount

Description: The number of batches processed by the channel that were sent because the value of the channel attributes BATCHSZ or BATCHLIM was reached.
Identifier: MQIAMO_FULL_BATCHES.
Data type: MQCFIN.
Returned: When available.

IncplBatchCount

Description: The number of batches processed by the channel that were sent without the value of the channel attributes BATCHSZ or BATCHLIM being reached.
Identifier: MQIAMO_INCOMPLETE_BATCHES.
Data type: MQCFIN.
Returned: When available.

AverageBatchSize

Description: The average batch size of batches processed by the channel.
Identifier: MQIAMO_AVG_BATCH_SIZE.
Data type: MQCFIN.
Returned: When available.

PutRetryCount

Description:	The number of times in the time interval that a message failed to be put, and entered a retry loop.
Identifier:	MQIAMO_PUT_RETRIES.
Data type:	MQCFIN.
Returned:	When available.

Reference notes

Use this page to view the notes to which descriptions of the structure of accounting and statistics messages refer

The following message data descriptions refer to these notes:

- [“MQI accounting message data” on page 177](#)
- [“Queue accounting message data” on page 188](#)
- [“MQI statistics message data” on page 199](#)
- [“Queue statistics message data” on page 211](#)
- [“Channel statistics message data” on page 217](#)

1. This parameter relates to IBM MQ objects. This parameter is an array of values (MQCFIL or MQCFIL64) indexed by the following constants:

Table 28. Array indexed by object type	
Object type	Value context
MQOT_Q (1)	Contains the value relating to queue objects.
MQOT_NAMELIST (2)	Contains the value relating to namelist objects.
MQOT_PROCESS (3)	Contains the value relating to process objects.
MQOT_Q_MGR (5)	Contains the value relating to queue manager objects.
MQOT_CHANNEL (6)	Contains the value relating to channel objects.
MQOT_AUTH_INFO (7)	Contains the value relating to authentication information objects.
MQOT_TOPIC (8)	Contains the value relating to topic objects.

Note: An array of 13 MQCFIL or MQCFIL64 values are returned but only those listed are meaningful.

2. This parameter relates to IBM MQ messages. This parameter is an array of values (MQCFIL or MQCFIL64) indexed by the following constants:

Table 29. Array indexed by persistence value	
Constant	Value
1	Contains the value for nonpersistent messages.
2	Contains the value for persistent messages.

Note: The index for each of these arrays starts at zero, so an index of 1 refers to the second row of the array. Elements of these arrays not listed in these tables contain no accounting or statistics information.

Application activity trace

Application activity trace produces detailed information about the behavior of applications connected to a queue manager. It traces the behavior of an application and provides a detailed view of the parameters used by an application as it interacts with IBM MQ resources. It also shows the sequence of MQI calls issued by an application.

Use Application activity trace when you require more information than is provided by Event monitoring, Message monitoring, Accounting and statistics messages, and Real-time monitoring.

Note: Activity trace is generated on the IBM MQ connection for each application; so if accounting messages are enabled the activity trace operations count towards the MQI accounting information for each application.

IBM MQ supports two methods of collecting application activity trace data.

- Central collection of application activity trace information, where application activity trace is collected and read by writing activity trace PCF messages to the system queue `SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE`.
- Subscription to activity trace data, written to special IBM MQ system topics.

Note that activity trace is not supported by IBM MQ for z/OS.

 As you can now specify application names on most of the programming languages that IBM MQ supports, refer to [Specifying the application name in supported programming languages](#) for more information.

Configuring central collection of application activity trace information

An application activity trace message is a PCF message. You configure activity trace using a configuration file. To configure central collection of application activity trace information, you set the ACTVTRC queue manager attribute. You can override this setting at connection level using MQCONN options, or at application stanza level using the activity trace configuration file.

About this task

Activity trace messages are composed of an MQMD structure: a PCF (MQCFH) header structure, followed by a number of PCF parameters. A sequence of ApplicationTraceData PCF groups follows the PCF parameters. These PCF groups collect information about the MQI operations that an application performs while connected to a queue manager. You configure activity trace using a configuration file called `mqat.ini`.

To control whether or not application activity trace information is collected, you configure one or more of the following settings:

1. The ACTVTRC queue manager attribute.
2. The ACTVCONO settings (in the MQCNO structure passed in MQCONN).
3. The matching stanza for the application in the activity trace configuration file `mqat.ini`.

The previous sequence is significant. The ACTVTRC attribute is overridden by the ACTVCONO settings, which are overridden by the settings in the `mqat.ini` file.

Trace entries are written after each operation has completed, unless otherwise stated. These entries are first written to the system queue `SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE`, then written to application activity trace messages when the application disconnects from the queue manager. For long running applications, intermediate messages are written if any of the following events occurs:

- The lifetime of the connection reaches a defined timeout value.
- The number of operations reaches a specified number.
- The amount of data collected in memory reaches the maximum message length allowed for the queue.

You set the timeout value using the **ActivityInterval** parameter. You set the number of operations using the **ActivityCount** parameter. Both parameters are specified in the activity trace configuration file `mqat.ini`.

Enabling application activity trace can affect performance. The overhead can be reduced by tuning the **ActivityCount** and the **ActivityInterval** settings. See [“Tuning the performance impact of application activity trace” on page 230](#).

The simplest way to view the contents of application activity trace messages is to use the [“amqsact sample program” on page 231](#).

Procedure

1. [“Setting ACTVTRC to control collection of activity trace information” on page 224](#).
2. [“Setting MQCONN options to control collection of activity trace information” on page 224](#).
3. [“Configuring activity trace behavior using mqat.ini” on page 225](#).
4. [“Tuning the performance impact of application activity trace” on page 230](#).

Setting ACTVTRC to control collection of activity trace information

Use the queue manager attribute ACTVTRC to control the collection of MQI application activity trace information

About this task

Application activity trace messages are generated only for connections that begin after application activity trace is enabled. The **ACTVTRC** parameter can have the following values:

ON

API activity trace collection is enabled

OFF

API activity trace collection is disabled

Note: The **ACTVTRC** setting can be overridden by the queue manager **ACTVCONO** parameter. If you set the **ACTVCONO** parameter to **ENABLED**, then the **ACTVTRC** setting can be overridden for a given connection using the **Options** field in the MQCNO structure. See [“Setting MQCONN options to control collection of activity trace information” on page 224](#).

Example

To change the value of the **ACTVTRC** parameter, you use the MQSC command `ALTER QMGR`. For example, to enable MQI application activity trace information collection use the following MQSC command:

```
ALTER QMGR ACTVTRC(ON)
```

What to do next

The simplest way to view the contents of application activity trace messages is to use the [“amqsact sample program” on page 231](#).

Enabling application activity trace can affect performance. The overhead can be reduced by tuning the **ActivityCount** and the **ActivityInterval** settings. See [“Tuning the performance impact of application activity trace” on page 230](#).

Setting MQCONN options to control collection of activity trace information

If the queue manager attribute **ACTVCONO** is set to **ENABLED**, you can use the **ConnectOpts** parameter on the MQCONN call to enable or disable application activity reports on a per connection basis. These options override the activity trace behavior defined by the queue manager attribute **ACTVTRC**, and can be overridden by settings in the activity trace configuration file `mqat.ini`.

Procedure

1. Set the queue manager attribute **ACTVCONO** to ENABLED.

Note: If an application attempts to modify the accounting behavior of an application using the **ConnectOpts** parameter, and the QMGR attribute **ACTVCONO** is set to DISABLED, then no error is returned to the application, and activity trace collection is defined by the queue manager attributes or the activity trace configuration file `mqat.ini`.

2. Set the **ConnectOpts** parameter on the MQCONN call to MQCNO_ACTIVITY_TRACE_ENABLED.

The **ConnectOpts** parameter on the MQCONN call can have the following values:

MQCNO_ACTIVITY_TRACE_DISABLED

Activity trace is disabled for the connection.

MQCNO_ACTIVITY_TRACE_ENABLED

Activity trace is enabled for the connection.

Note: If an application selects both MQCNO_ACTIVITY_TRACE_ENABLED and MQCNO_ACTIVITY_TRACE_DISABLED for MQCONN, the call fails with a reason code of MQRC_OPTIONS_ERROR.

3. Check that these activity trace settings are not being overridden by settings in the activity trace configuration file `mqat.ini`.

See [“Configuring activity trace behavior using mqat.ini” on page 225](#).

What to do next

The simplest way to view the contents of application activity trace messages is to use the [“amqsact sample program” on page 231](#).

Enabling application activity trace can affect performance. The overhead can be reduced by tuning the **ActivityCount** and the **ActivityInterval** settings. See [“Tuning the performance impact of application activity trace” on page 230](#).

Configuring activity trace behavior using mqat.ini

Activity trace behavior is configured using a configuration file called `mqat.ini`. This file is used to define the level and frequency of reporting activity trace data. The file also provides a way to define rules to enable and disable activity trace based on the name of an application.

About this task

  On AIX and Linux systems, `mqat.ini` is located in the queue manager data directory, which is the same location as the `qm.ini` file.

 On Windows systems, `mqat.ini` is located in the queue manager data directory `C:\Program Files\IBM\WebSphere MQ\qmgsr\queue_manager_name`. Users running applications to be traced need permission to read this file.

Note: Queue managers migrated from IBM WebSphere MQ 7.1 or earlier will have the `mqat.ini` file missing. In such cases, the `mqat.ini` file needs to be created manually and 660 permissions need to be set on the file.

When the `mqat.ini` file is modified, newly created IBM MQ connections will be processed according to the modified version. Existing connections will continue to use the previous version unless the queue manager parameters are altered, for example following an [ALTER QMGR](#) command.

This file follows the same stanza key and parameter-value pair format as the `mqs.ini` and `qm.ini` files.

The file consists of a single stanza, **AllActivityTrace**, to configure the level and frequency of reporting activity trace data by default for all activity trace.

The file can also contain multiple **ApplicationTrace** stanzas. Each one of these, defines a rule for the trace behavior for one or more connections, based on matching the application name of the connections to the rule.

AllActivityTrace stanza

A single **AllActivityTrace** stanza defines settings for the activity trace that is applied to all IBM MQ connections, unless overridden.

Individual values in the **AllActivityTrace** stanza can be overridden by more specific information in an **ApplicationTrace** stanza.

If more than one **AllActivityTrace** stanza is specified then the values in the last stanza is used. Parameters missing from the chosen **AllActivityTrace** take default values. Parameters and values from previous **AllActivityTrace** stanzas are ignored.

The following parameters can be specified under the **AllActivityTrace** stanza:

Table 30. Parameter/value pairs that can be used in the activity trace configuration file		
Name	Values (default in bold type)	Description
ActivityInterval	0-999999999 (1)	Approximate time interval in seconds between trace messages. All activity performed by a connection in that interval will be written in a single message. If this value is 0, the trace message is written when the connection disconnects (or when the activity count is reached).
ActivityCount	0-999999999 (100)	Number of MQI or XA operations between trace messages. If this value is 0, the trace message is written when the connection disconnects (or when the activity interval has elapsed).
TraceLevel	LOW / MEDIUM / HIGH	Amount of parameter detail traced for each operation. The description of individual operations details which parameters are included for each trace level.
TraceMessageData	0 - 104 857 600 (maximum 100 MB)	Amount of message data traced in bytes for MQGET, MQPUT, MQPUT1, and Callback operations
StopOnGetTraceMsg	ON / OFF	Using activity trace, to trace applications that are also processing activity trace messages, is not advisable due to possible looping occurring.
SubscriptionDelivery	BATCHED / 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.

ApplicationTrace stanza

An ApplicationTrace stanza contains a rule which defines which IBM MQ connections will be traced or not trace based on the application name. Optionally, the default behaviour defined under the Allsettings which override the global trace level and frequency settings.

This stanza can include ApplName, ApplFunction and ApplClass parameters which are used according to the matching rules defined in Connection Matching Rules to determine whether the stanza applies to a particular connection or not.

The stanza must include the Trace parameter to determine if this rule turns activity trace on or off for matching connections.

An off rule can be used to explicitly disable trace for more specific application names and to override the ACTVTRC setting of the queue manager or activity trace connection options.

The following parameters can be specified under the **ApplicationTrace** stanza:

Table 31. Parameter/value pairs that can be used in the application trace configuration file		
Name	Values (default in bold type)	Description
Trace	ON / OFF (Required parameter - no default value)	Activity trace switch. This switch can be used in the application-specific stanza to determine whether activity trace is active for the scope of the current application stanza. Note that this value overrides ACTVTRC and ACTVCONO settings for the queue manager.
ApplName	Character string (Required parameter - no default)	This value is used to determine which applications the ApplicationTrace stanza applies to. It is matched to the ApplName value from the API exit context structure (which is equivalent to the MQMD.PutApplName). The content of the ApplName value varies according to the application environment. For platforms other than z/OS, only the filename portion of the MQAXC.ApplName is matched to the value in the stanza. Characters to the left of the rightmost path separator are ignored when the comparison is made.  z/OS For z/OS applications, the entire MQAXC.ApplName is matched to the value in the stanza. A single wildcard character (*) can be used at the end of the ApplName value to match any number of characters after that point. If the ApplName value is set to a single wildcard character (*) then the ApplName value matches all applications.

Table 31. Parameter/value pairs that can be used in the application trace configuration file (continued)

Name	Values (default in bold type)	Description
  ApplFunction	Character string (default value *)	This value is used to qualify which application programs the ApplicationTrace stanza and App1Name value applies to. The stanza is optional, but is only valid for IBM i queue managers. A single wildcard character (*) can be used at the end of the App1Name value to match any number of characters. For example, an ApplicationTrace stanza specifying App1Name = * and App1Function = AMQSPUT0 applies to all invocations of the AMQSPUT0 program from any job.
App1Class	USER / MCA / ALL	The class of application. See the following table for an explanation of how the AppType values correspond to IBM MQ connections.

The following table shows how the *AppClass* values correspond to the **APICallerType** and **APIEnvironment** fields in the connection API exit context structure.

Table 32. Appclass values and how they correspond to the APICallerType and APIEnvironment fields

APPLCLASS	API Caller Type:	API Environment:	Description
USER	MQXACT_EXTERNAL	MQXE_OTHER	Only user applications are traced
MCA	(Any value)	MQXE_MCA MQXE_MCA_CLNTCONN MQXE_MCA_SVRCONN	Clients and channels (amqrmppa)
ALL	(Any value)	(Any value)	All connections are traced



Attention: You must use an **APPLCLASS** of **MCA** for client user applications, as a class of *USER* does not match these.

For example, to trace the **amqspu** sample application, you could use the following code:

```

ApplicationTrace:
App1Class=MCA                                     # Application type
                                                # Values: (USER | MCA | INTERNAL | ALL)
                                                # Default: USER
App1Name=amqspu      # Application name (may be wildcarded)
                                                # (matched to app name without path)
                                                # Default: *
Trace=ON                                             # Activity trace switch for application
                                                # Values: ( ON | OFF )
                                                # Default: OFF
ActivityInterval=30                                # Time interval between trace messages
                                                # Values: 0-99999999 (0=off)
                                                # Default: 0
ActivityCount=1                                     # Number of operations between trace msgs
                                                # Values: 0-99999999 (0=off)
                                                # Default: 0
TraceLevel=MEDIUM                                 # Amount of data traced for each operation
                                                # Values: LOW | MEDIUM | HIGH
                                                # Default: MEDIUM
TraceMessageData=1000                             # Amount of message data traced
                                                # Values: 0-1000000000
                                                # Default: 0
    
```

The default `mqat.ini` generated when a queue manager is created, contains a single rule to explicitly disable activity trace for the supplied activity trace sample, **amqsact**.

Connection Matching Rules

The queue manager applies the following rules to determine which stanzas settings to use for a connection.

1. A value specified in the **AllActivityTrace** stanza is used for the connection unless the value also occurs in an **ApplicationTrace** stanza and the stanza fulfills the matching criteria for the connection described in points 2, 3, and 4.
2. The **App1Class** is matched against the type of the IBM MQ connection. If the **App1Class** does not match the connection type then the stanza is ignored for this connection.
3. The *App1Name* value in the stanza is matched against the file name portion of the **App1Name** field from the API exit context structure (MQAXC) for the connection.

The file name portion is derived from the characters to the right of the final path separator (/ or \) character. If the stanza **App1Name** includes a wildcard (*) then only the characters to the left of the wildcard are compared with the equivalent number of characters from the **App1Name** of the connection.

For example, if a stanza value of "FRE*" is specified then only the first three characters are used in the comparison, so "path/FREEDOM" and "path\FREDDY" match, but "path/FRIEND" does not. If the *App1Name* value of the stanza does not match the connection **App1Name**, the stanza is ignored for this connection.

4. If more than one stanza matches the **App1Name** and **App1Class** of the connection, then the stanza with the most specific **App1Name** is used.

The most specific **App1Name** is defined as the one that uses the most characters to match the **App1Name** of the connection.

For example, if the ini file contains a stanza with **App1Name** = "FRE*" and another stanza with **App1Name** = "FREE*" then the stanza with **App1Name** = "FREE*" is chosen as the best match for a connection with **App1Name** = "path/FREEDOM" because it matches four characters (whereas **App1Name** = "FRE*" matches only three).

5. If after applying the rules in points 2, 3, and 4, there is more than one stanza that matches the connections **App1Name** and **App1Class** of the connection, the values from the last matching will be used and all other stanzas will be ignored.

Overriding default settings for each rule

Optionally, the global trace level and frequency settings under the **AllActivityTrace** stanza can be overridden for those connections matching an **ApplicationTrace** stanza.

The following parameters can be set under an **ApplicationTrace** stanza. If they are not set, the value is inherited from the **AllActivityTrace** stanza settings:

- **ActivityInterval**
- **ActivityCount**
- **TraceLevel**
- **TraceMessageData**
- **StopOnTraceMsg**

mqat.ini syntax

The syntax rules for the format of the mqat.ini file are:

- Text beginning with a hash or semicolon is considered to be a comment that extends to the end of the line.
- The first significant (non-comment) line must be a stanza key.
- A stanza key consists of the name of the stanza followed by a colon.
- A parameter-value pair consists of the name of a parameter followed by an equals sign and then the value.

- Only a single parameter-value pair can appear on a line. (A parameter-value must not wrap onto another line).
- Leading and trailing whitespace is ignored. There is no limit on the amount of white space between stanza names, parameter names and values, or parameter/value pairs. Line breaks are significant and not ignored
- The maximum length for any line is 2048 characters
- The stanza keys, parameter names, and constant parameter values are not case-sensitive, but the variable parameter values (*ApplName* and *DebugPath*) are case-sensitive.

Application Activity Trace File Example

The following example shows how the configuration data is specified in the Activity Trace ini file.

```
AllActivityTrace:
ActivityInterval=1
ActivityCount=100
TraceLevel=MEDIUM
TraceMessageData=0
StopOnGetTraceMsg=ON

ApplicationTrace:
ApplName=amqs*
Trace=ON
TraceLevel=HIGH
TraceMessageData=1000

ApplicationTrace:
ApplName=amqsact*
Trace=OFF
```

The above **AllActivityTrace** stanza defines how activity trace will perform by default when enabled, either through ApplicationTrace rules or through the queue manager ACTVTRC attribute or programmatically enabled by an application.

The first **ApplicationTrace** stanza defines a rule that will result in any MQI activity by an application whose name starts with "amqs" being traced. Trace generated for these applications will be of high detail and include up to 1000 bytes of message data. The activity interval and count parameters will be inherited

The second **ApplicationTrace** stanza defines a rule that turns trace off for applications with names starting "amqsact" (the activity trace sample). This rule will override the earlier 'on' rule for the amqsact application, resulting in no trace for that application.

An example is also shipped as a sample called `mqat.ini` in the C samples directory (the same directory as the `amqsact.c` file). This file can be copied to the queue manager data directory, for queue managers that have been migrated from an earlier release of IBM MQ.

What to do next

Enabling application activity trace can affect performance. The overhead can be reduced by tuning the **ActivityCount** and the **ActivityInterval** settings. See [“Tuning the performance impact of application activity trace” on page 230](#).

Tuning the performance impact of application activity trace

Enabling application activity trace can incur a performance penalty. This can be reduced by only tracing the applications that you need, by increasing the number of applications draining the queue, and by tuning **ActivityInterval**, **ActivityCount** and **TraceLevel** in `mqat.ini`.

About this task

Enabling application activity trace selectively for an application or for all queue manager applications can result in additional messaging activity, and in the queue manager requiring additional storage space. In environments where messaging performance is critical, for example, in high workload applications

or where a service level agreement (SLA) requires a minimum response time from the messaging provider, it might not be appropriate to collect application activity trace or it might be necessary to adjust the detail or frequency of trace activity messages that are produced. The preset values of **ActivityInterval**, **ActivityCount** and **TraceLevel** in the `mqat.ini` file give a default balance of detail and performance. However, you can tune these values to meet the precise functional and performance requirements of your system.

Procedure

- Only trace the applications that you need.

Do this by creating an ApplicationTrace application-specific stanza in `mqat.ini`, or by changing the application to specify `MQCNO_ACTIVITY_TRACE_ENABLED` in the options field on the **MQCNO** structure on an MQCONN call. See [“Configuring activity trace behavior using mqat.ini”](#) on page 225 and [“Setting MQCONN options to control collection of activity trace information”](#) on page 224.

- Before starting trace, check that at least one application is running and is ready to retrieve the activity trace message data from the `SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE`.
- Keep the queue depth as low as possible, by increasing the number of applications draining the queue.
- Set the **TraceLevel** value in the `mqat.ini` file to collect the minimum amount of data required.

`TraceLevel=LOW` has the lowest impact to messaging performance. See [“Configuring activity trace behavior using mqat.ini”](#) on page 225.

- Tune the **ActivityCount** and **ActivityInterval** values in `mqat.ini`, to adjust how often activity trace messages are generated.

If you are tracing multiple applications, the activity trace messages might be being produced faster than they can be removed from the `SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE`. However, when you reduce how often activity trace messages are generated, you are also increasing the storage space required by the queue manager and the size of the messages when they are written to the queue.

What to do next

The simplest way to view the contents of application activity trace messages is to use the [“amqsact sample program”](#) on page 231.

amqsact sample program

amqsact formats Application Activity Trace messages for you and is provided with IBM MQ.

The compiled program is located in the samples directory:


-   On AIX and Linux `MQ_INSTALLATION_PATH/samp/bin`
-  On Windows `MQ_INSTALLATION_PATH\tools\c\Samples\Bin`

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.

You can also control the trace period displayed and specify whether the activity trace messages are removed or retained after display.

amqsact



Required parameters for display mode

-m *QMgrName*

Name of the queue manager.

-d *Depth*

Number of records to display.

-w *Timeout*

Time to wait, in seconds. If no trace messages appear in the specified period, **amqsact** exits.

-s *StartTime*

Start time of record to process.

-e *EndTime*

End time of record to process.

Optional parameters for display mode

-q *QName*

Specify a specific queue to override the default queue name

-t *TopicString*

Subscribe to an event topic

-b

Browse records only

-v

Verbose output

Example output for display mode

Use **amqsact** on queue manager *TESTQM*, with verbose output, on an MQCONN API call:

```
amqsact -m TESTQM -v
```

The preceding command gives the following example output:

```
MonitoringType: MQI Activity Trace
Correl_id:
00000000: 414D 5143 5445 5354 514D 2020 2020 2020 'AMQCTESTQM'
00000010: B5F6 4251 2000 E601
QueueManager: 'TESTQM'
Host Name: 'ADMINIB-1VTJ6N1'
IntervalStartDate: '2014-03-15'
IntervalStartTime: '12:08:10'
IntervalEndDate: '2014-03-15'
IntervalEndTime: '12:08:10'
CommandLevel: 750
SeqNumber: 0
ApplicationName: 'IBM MQ_1\bin\amqsput.exe'
Application Type: MQAT_WINDOWS_7
ApplicationPid: 14076
UserId: 'Emma_Bushby'
API Caller Type: MQXACT_EXTERNAL
```



```


API Environment: MQXE_OTHER
Application Function: ''
Appl Function Type: MQFUN_TYPE_UNKNOWN
Trace Detail Level: 2
Trace Data Length: 0
Pointer size: 4
Platform: MQPL_WINDOWS_7
MQI Operation: 0
Operation Id: MQXF_CONN
ApplicationTid: 1
OperationDate: '2014-03-15'
OperationTime: '12:08:10'
ConnectionId:
00000000: 414D 5143 5445 5354 514D 2020 2020 2020 'AMQCTESTQM'
00000010: FFFFFFFB5FFFFFFF6 4251 2000 FFFFFFFE601
QueueManager: 'TESTQM'
Completion Code: MQCC_OK
Reason Code: 0

```

Dynamic mode

You can enable dynamic mode by specifying an application name, a channel name, or a connection identifier as an argument to **amqsact**. Note that 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 nondurable 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, with each instance receiving a copy of any activity trace data.



Required parameters for dynamic mode

-m *QMgrName*

Name of the queue manager.

-w *Timeout*

Time to wait, in seconds. If no trace messages appear in the specified period, **amqsact** exits.

Optional parameters for dynamic mode

-a *Application name*

Specify an application name 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

Verbose output

Example output for dynamic mode

The following command generates and displays activity trace messages for any connections 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.

```
amqsact -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.

```
amqsact -m QMGR1 -w 10 -c QMGR1.TO.QMGR2
```

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.

```
amqsact -m QMGR1 -w 60 -i 414D5143514D475231202020202020206B576B5420000701 -v
```

Subscribing to application activity trace information

You can dynamically subscribe to application activity trace information as an alternative to collecting the information through queue manager level configuration.

About this task

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.

In addition to writing trace data to the system queue, the product provides the ability to dynamically subscribe to activity trace data, written to special IBM MQ system topics, rather than collecting the information through queue manager level configuration.

Note that from IBM MQ 9.0 the product does not use exits for this purpose. If you have previously used exits to trace application activity, you must switch to using the replacement method for collecting application activity trace.

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.

Procedure

- [“Subscriptions to application activity trace” on page 234](#)
- [“Creating subscriptions to application activity trace” on page 235](#)
- [“Using amqsact to view trace messages” on page 236](#)
- [“Configuring trace levels using mqat.ini” on page 238](#)

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.

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](#).

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 amqspu^t running on a Windows system:

```
$SYS/MQ/INFO/QMGR/QMGR1/ActivityTrace/AplName/amqsput.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/  
414D5143514D4752312020202020206B576B5420000701
```

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/AplName/amqs*
```

Related concepts

[“System topics for monitoring and activity trace” on page 308](#)

System topics in queue manager topic trees are used for resource monitoring (some of which is similar to the content of statistics messages) and as a way of consuming application activity trace.

Using amqsact to view trace messages

You can use the **amqsact** program to generate and view trace messages.

The **amqsact** program is an IBM MQ sample. To use this sample, 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](#). 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.

```
amqsact -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.

```
amqsact -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.

```
amqsact -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.

```
amqsact -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 EXTCNN of "414D5143514D47523120202020202020". After a minute of inactivity, the **amqsact** program ends, and no new activity trace data is generated.

```
amqsact -m QMGR1 -w 60 -i 414D5143514D475231202020202020206B576B5420000701 -v
```

Configuring trace levels using mqat.ini

You configure trace levels for a queue manager by setting values for the AllActivityTrace stanza of the mqat.ini configuration file.

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.

Application activity trace message reference

Use this page to obtain an overview of the format of application activity trace messages and the information returned in these messages

Application activity trace messages are standard IBM MQ messages containing a message descriptor and message data. The message data contains information about the MQI operations performed by IBM MQ applications, or information about the activities occurring in an IBM MQ system.

Message descriptor

- An MQMD structure

Message data

- A PCF header (MQCFH)
- Application activity trace message data that is always returned
- Application activity trace message data that is operation-specific

Application activity trace message MQMD (message descriptor)

Use this page to understand the differences between the message descriptor of application activity trace messages and the message descriptor of event messages

The parameters and values in the message descriptor of application activity trace message are the same as in the message descriptor of event messages, with the following exception:

Format

Description: Format name of message data.

Value: **MQFMT_ADMIN**
Admin message.

CorrelId

Description: Correlation identifier.

Value: Initialized with the ConnectionId of the application

MQCFH (PCF Header)

Use this page to view the PCF values contained by the MQCFH structure for an activity trace message

For an activity trace message, the MQCFH structure contains the following values:

Type

Description: Structure type that identifies the content of the message.

Data type: MQLONG.
Value: MQCFT_APP_ACTIVITY

StrucLength

Description: Length in bytes of MQCFH structure.
Data type: MQLONG.
Value: MQCFH_STRUC_LENGTH

Version

Description: Structure version number.
Data type: MQLONG.
Values: MQCFH_VERSION_3

Command

Description: Command identifier. This field identifies the category of the message.
Data type: MQLONG.
Values: MQCMD_ACTIVITY_TRACE

MsgSeqNumber

Description: Message sequence number. This field is the sequence number of the message within a group of related messages.
Data type: MQLONG.
Values: 1

Control

Description: Control options.
Data type: MQLONG.
Values: MQCFC_LAST.

CompCode

Description: Completion code.
Data type: MQLONG.
Values: MQCC_OK.

Reason

Description: Reason code qualifying completion code.
Data type: MQLONG.
Values: MQRC_NONE.

ParameterCount

Description: Count of parameter structures. This field is the number of parameter structures that follow the MQCFH structure. A group structure (MQCFGR), and its included parameter structures, are counted as one structure only.

Data type: MQLONG.
 Values: 1 or greater

Application activity trace message data

Immediately following the PCF header is a set of parameters describing the time interval for the activity trace. These parameters also indicate the sequence of messages in the event of messages being written. The order and number of fields following the header is not guaranteed, allowing additional information to be added in the future.

Message name:	Activity trace message.
System queue:	SYSTEM.ADMIN.TRACE.ACTIVITY.QUEUE.

QueueManager

Description: The name of the queue manager.
 Identifier: MQCA_Q_MGR_NAME
 Data type: MQCFST
 Maximum length: MQ_Q_MGR_NAME_LENGTH

QSGName



Description: The name of the queue sharing group that the Queue Manager is a member of (z/OS only).
 Identifier: MQCA_QSG_NAME
 Data type: MQCFST
 Maximum length: MQ_Q_MGR_NAME_LENGTH

HostName

Description: The host name of the machine the Queue Manager is running on.
 Identifier: MQCACF_HOST_NAME
 Data type: MQCFST

IntervalStartDate

Description: The date of the start of the monitoring period.
 Identifier: MQCAMO_START_DATE
 Data type: MQCFST
 Maximum length: MQ_DATE_LENGTH

IntervalStartTime

Description: The time of the start of the monitoring period.
 Identifier: MQCAMO_START_TIME
 Data type: MQCFST
 Maximum length: MQ_TIME_LENGTH

IntervalEndDate

Description: The date of the end of the monitoring period.
Identifier: MQCAMO_END_DATE
Data type: MQCFST
Maximum length: MQ_DATE_LENGTH

IntervalEndTime

Description: The time of the end of the monitoring period.
Identifier: MQCAMO_END_TIME
Data type: MQCFST
Maximum length: MQ_TIME_LENGTH

CommandLevel

Description: The IBM MQ command level.
Identifier: MQIA_COMMAND_LEVEL
Data type: MQCFIN

SeqNumber

Description: The sequence number normally zero. This value is incremented for each subsequent record for long running connections.
Identifier: MQIACF_SEQUENCE_NUMBER
Data type: MQCFIN

ApplicationName

Description: The name of the application (program name).
Identifier: MQCACF_APPL_NAME
Data type: MQCFST
Maximum length: MQ_APPL_NAME_LENGTH

ApplClass

Description: Type of application that performed the activity. Possible values: MQAT_*
Identifier: MQIA_APPL_TYPE
Data type: MQCFIN

ApplicationPid

Description: The operating system Process ID of the application.
Identifier: MQIACF_PROCESS_ID
Data type: MQCFIN

UserId

Description: The user identifier context of the application.
Identifier: MQCACF_USER_IDENTIFIER

Data type: MQCFST
Maximum length: MQ_USER_ID_LENGTH

APICallerType

Description: The type of the application. Possible values: MQXACT_EXTERNAL or MQXACT_INTERNAL.
Identifier: MQIACF_API_CALLER_TYPE
Data type: MQCFIN

Environment

Description: The runtime environment of the application. Possible values: MQXE_*
Identifier: MQIACF_API_ENVIRONMENT
Data type: MQCFIN

ChannelName

Description: The channel name associated with the connection. This parameter is only returned when the value of the Environment parameter is either MQXE_MCA or MQXE_MCA_SVRCONN.
Identifier: MQCACH_CHANNEL_NAME
Data type: MQCFST
Maximum length: MQ_CHANNEL_NAME_LENGTH

ConnectionName

Description: The network connection name associated with the connection. This parameter is only returned when the value of the Environment parameter is either MQXE_MCA or MQXE_MCA_SVRCONN.
Identifier: MQCACH_CONNECTION_NAME
Data type: MQCFST
Maximum length: MQ_CONN_NAME_LENGTH

ChannelType

Description: The type of channel associated with the connection. This parameter is only returned when the value of the Environment parameter is either MQXE_MCA or MQXE_MCA_SVRCONN. Possible values: MQCHT_*
Identifier: MQIACH_CHANNEL_TYPE
Data type: MQCFIN

RemoteProduct

Description: The remote product identifier associated with the connection. This parameter is only returned when the value of the Environment parameter is either MQXE_MCA or MQXE_MCA_SVRCONN.
Identifier: MQCACH_REMOTE_PRODUCT
Data type: MQCFST
Maximum length: MQ_REMOTE_PRODUCT_LENGTH

RemoteVersion

Description:	The remote product version associated with the connection. This parameter is only returned when the value of the Environment parameter is either MQXE_MCA or MQXE_MCA_SVRCONN.
Identifier:	MQCACH_REMOTE_VERSION
Data type:	MQCFST
Maximum length:	MQ_REMOTE_VERSION_LENGTH

FunctionName

Description:	The name of the last high-level function initiated by the initial thread.
Identifier:	MQCACF_APPL_FUNCTION
Data type:	MQCFST

FunctionType

Description:	The type of the last high-level function initiated by the initial thread. Possible values: MQFUN_*
Identifier:	MQIACF_APPL_FUNCTION_TYPE
Data type:	MQCFIN

Detail

Description:	The detail level that is recorded for the connection. Possible values: 1=LOW 2=MEDIUM 3=HIGH
Identifier:	MQIACF_TRACE_DETAIL
Data type:	MQCFIN

TraceDataLength

Description:	The length of message data (in bytes) that is traced for this connection.
Identifier:	MQIACF_TRACE_DATA_LENGTH
Data type:	MQCFIN

PointerSize

Description:	The length (in bytes) of pointers on the platform the application is running (to assist in interpretation of binary structures).
Identifier:	MQIACF_POINTER_SIZE
Data type:	MQCFIN

Platform

Description:	The platform on which the queue manager is running. Possible values: MQPL_*
Identifier:	MQIA_PLATFORM
Data type:	MQCFIN

Variable parameters for application activity MQI operations

The application activity data MQCFGR structure is followed by the set of PCF parameters which corresponds to the operation being performed . The parameters for each operation are defined in the following section.

The trace level indicates the level of trace granularity that is required for the parameters to be included in the trace. The possible trace level values are:

1. Low

The parameter is included when "low", "medium" or "high" activity tracing is configured for an application. This setting means that a parameter is always included in the AppActivityData group for the operation. This set of parameters is sufficient to trace the MQI calls an application makes, and to see if they are successful.

2. Medium

The parameter is only included in the AppActivityData group for the operation when "medium" or "high" activity tracing is configured for an application. This set of parameters adds information about the resources, for example, queue and topic names used by the application.

3. High

The parameter is only included in the AppActivityData group for the operation when "high" activity tracing is configured for an application. This set of parameters includes memory dumps of the structures passed to the MQI and XA functions. For this reason, it contains more information about the parameters used in MQI and XA calls. The structure memory dumps are shallow copies of the structures. To avoid erroneous attempts to dereference pointers, the pointer values in the structures are set to NULL.

Note: The version of the structure that is dumped is not necessarily identical to the version used by an application. The structure can be modified by an API crossing exit, by the activity trace code, or by the queue manager. A queue manager can modify a structure to a later version, but the queue manager never changes it to an earlier version of the structure. To do so, would risk losing data.

MQBACK

Application has started the MQBACK MQI function

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type	MQCFIN

QMGrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client.
--------------	--

Note: The accuracy of this timer varies according to the platform that your enterprise uses.

PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQBEGIN

Application has started the MQBEGIN MQI function

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type	MQCFIN

MQBO

Description:	The MQBEGIN options structure. This parameter is not included if a NULL pointer is used on the MQBEGIN call.
PCF Parameter:	MQBACF_MQBO_STRUCT
Trace level:	3
Type	MQCFBS
Length:	The length in bytes of the MQBO structure.

QMGrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQCALLBACK

Application has started the MQCALLBACK function

ObjectHandle

Description:	The object handle
PCF Parameter:	MQIACF_HOBJ
Trace level:	1

Type	MQCFIN
------	--------

CallType

Description:	Why function has been called. One of the MQCBCT_* values
--------------	--

PCF Parameter:	MQIACF_CALL_TYPE
----------------	------------------

Trace level:	1
--------------	---

Type	MQCFIN
------	--------

MsgBuffer

Description:	Message data.
--------------	---------------

PCF Parameter:	MQBACF_MESSAGE_DATA
----------------	---------------------

Trace level:	1
--------------	---

Type	MQCFBS
------	--------

Length:	Length is governed by the TRACEDATA() parameter set in the APPTRACE configuration. If TRACEDATA=NONE then this parameter is omitted.
---------	--

MsgLength

Description:	Length of the message. (Taken from the DataLength field in the MQCBC structure).
--------------	--

PCF Parameter:	MQIACF_MSG_LENGTH
----------------	-------------------

Trace level:	1
--------------	---

Type	MQCFIN
------	--------

HighResTime

Description:	Time of operation in microseconds since midnight, January 1st 1970 (UTC)
--------------	--

Note: The accuracy of this timer varies according to platform support for high a resolution timer

PCF Parameter:	MQIAMO64_HIGHRES_TIME
----------------	-----------------------

Trace level:	2
--------------	---

Type	MQCFIN64
------	----------

ReportOptions

Description:	Options for report messages
--------------	-----------------------------

PCF Parameter:	MQIACF_REPORT
----------------	---------------

Trace level:	2
--------------	---

Type	MQCFIN
------	--------

MsgType

Description:	Type of message
--------------	-----------------

PCF Parameter:	MQIACF_MSG_TYPE
----------------	-----------------

Trace level:	2
--------------	---

Type	MQCFIN
------	--------

Expiry

Description:	Message lifetime
PCF Parameter:	MQIACF_EXPIRY
Trace level:	2
Type	MQCFIN

Format

Description:	Format name of message data
PCF Parameter:	MQCACH_FORMAT_NAME
Trace level:	2
Type	MQCFST
Length:	MQ_FORMAT_LENGTH

Priority

Description:	Message priority
PCF Parameter:	MQIACF_PRIORITY
Trace level:	2
Type	MQCFIN

Persistence

Description:	Message persistence
PCF Parameter:	MQIACF_PERSISTENCE
Trace level:	2
Type	MQCFIN

MsgId

Description:	Message identifier
PCF Parameter:	MQBACF_MSG_ID
Trace level:	2
Type	MQCFBS
Length:	MQ_MSG_ID_LENGTH

CorrelId

Description:	Correlation identifier
PCF Parameter:	MQBACF_CORREL_ID
Trace level:	2
Type	MQCFBS
Length:	MQ_CORREL_ID_LENGTH

ObjectName

Description:	The name of the opened object.
--------------	--------------------------------

PCF Parameter:	MQCACF_OBJECT_NAME
Trace level:	2
Type	MQCFST
Length:	MQ_Q_NAME_LENGTH

ResolvedQName

Description:	The local name of the queue from which the message was retrieved.
PCF Parameter:	MQCACF_RESOLVED_Q_NAME
Trace level:	2
Type	MQCFST
Length:	MQ_Q_NAME_LENGTH

ReplyToQueue

Description:	MQ_Q_NAME_LENGTH
PCF Parameter:	MQCACF_REPLY_TO_Q
Trace level:	2
Type	MQCFST

ReplyToQMgr

Description:	MQ_Q_MGR_NAME_LENGTH
PCF Parameter:	MQCACF_REPLY_TO_Q_MGR
Trace level:	2
Type	MQCFST

CodedCharSetId

Description:	Character set identifier of message data
PCF Parameter:	MQIA_CODED_CHAR_SET_ID
Trace level:	2
Type	MQCFIN

Encoding

Description:	Numeric encoding of message data.
PCF Parameter:	MQIACF_ENCODING
Trace level:	2
Type	MQCFIN

PutDate

Description:	MQ_PUT_DATE_LENGTH
PCF Parameter:	MQCACF_PUT_DATE
Trace level:	2
Type	MQCFST

PutTime

Description:	MQ_PUT_TIME_LENGTH
PCF Parameter:	MQCACF_PUT_TIME
Trace level:	2
Type	MQCFST

ResolvedQName

Description:	The queue name referred to by the ObjectHandle, when ResolvedType is MQOT_Q.
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level:	2
Type	MQCFST
Length:	MQ_Q_NAME_LENGTH.

ResObjectString

Description:	The object name referred to by the ObjectHandle, when ResolvedType is MQOT_TOPIC.
PCF Parameter:	MQCACF_RESOLVED_OBJECT_STRING
Trace level:	2
Type	MQCFST
Length:	Length varies.

ResolvedType

Description:	The type of the object referred to by the ObjectHandle. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter:	MQIACF_RESOLVED_TYPE
Trace level:	2
Type	MQCFIN

PolicyName

Description:	The policy name that was applied to this message. Note: AMS protected messages only
PCF Parameter:	MQCA_POLICY_NAME
Trace level:	2
Type	MQCFST
Length:	MQ_OBJECT_NAME_LENGTH

XmitqMsgId

Description:	The message ID of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQBACF_XQH_MSG_ID

Trace level:	2
Type	MQCFBS
Length:	MQ_MSG_ID_LENGTH

XmitqCorrelId

Description:	The correlation ID of the message in the transmission queue header.
--------------	---

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter:	MQBACF_XQH_CORREL_ID
Trace level:	2
Type	MQCFBS
Length:	MQ_CORREL_ID_LENGTH

XmitqPutTime

Description:	The put time of the message in the transmission queue header.
--------------	---

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter:	MQCACF_XQH_PUT_TIME
Trace level:	2
Type	MQCFST
Length:	MQ_PUT_TIME_LENGTH

XmitqPutDate

Description:	The put date of the message in the transmission queue header.
--------------	---

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter:	MQCACF_XQH_PUT_DATE
Trace level:	2
Type	MQCFST
Length:	MQ_PUT_DATE_LENGTH

XmitqRemoteQName

Description:	The remote queue destination of the message in the transmission queue header.
--------------	---

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter:	MQCACF_XQH_REMOTE_Q_Name
Trace level:	2
Type	MQCFST
Length:	MQ_Q_NAME_LENGTH

XmitqRemoteQMgr

Description:	The message ID of the message in the transmission queue header.
--------------	---

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter:	MQCACF_XQH_REMOTE_Q_MGR
----------------	-------------------------

Trace level:	2
Type	MQCFST
Length:	MQ_MSG_ID_LENGTH

MsgDescStructure

Description:	The MQMD structure. This parameter is omitted if a version 4 MQGMO was used to request that a Message Handle be returned instead of an MQMD
PCF Parameter:	MQBACF_MQMD_STRUCT
Trace level:	3
Type	MQCFBS
Length:	The length in bytes of the MQMD structure (actual size is dependent on structure version)

GetMsgOptsStructure

Description:	The MQGMO structure.
PCF Parameter:	MQBACF_MQGMO_STRUCT
Trace level:	3
Type	MQCFBS
Length:	The length in bytes of the MQGMO structure (actual size is dependent on structure version)

MQCBCContextStructure

Description:	The MQCBC structure.
PCF Parameter:	MQBACF_MQCBC_STRUCT
Trace level:	3
Type	MQCFBS
Length:	The length in bytes of the MQCBC structure (actual size is dependent on structure version)

QMgrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQCB

Application has started the manage callback MQI function

CallbackOperation

Description:	The manage callback function operation. Set to one of the MQOP_* values
--------------	---

PCF Parameter:	MQIACF_MQCB_OPERATION
Trace level:	1
Type	MQCFIN

CallbackType

Description:	The type of the callback function (CallbackType field from the MQCBD structure). Set to one of the MQCBT_* values
PCF Parameter:	MQIACF_MQCB_TYPE
Trace level:	1
Type	MQCFIN

CallbackOptions

Description:	The callback options. Set to one of the MQCBDO_* values
PCF Parameter:	MQIACF_MQCB_OPTIONS
Trace level:	1
Type	MQCFIN

CallbackFunction

Description:	The pointer to the callback function if started as a function call.
PCF Parameter:	MQBACF_MQCB_FUNCTION
Trace level:	1
Type	MQCFBS
Length:	Size of MQPTR

CallbackName

Description:	The name of the callback function if started as a dynamically linked program.
PCF Parameter:	MQCACF_MQCB_NAME
Trace level:	1
Type	MQCFST
Length:	Size of MQCHAR128

ObjectHandle

Description:	The object handle
PCF Parameter:	MQIACF_HOBJ
Trace level:	1
Type	MQCFIN

MaxMsgLength

Description:	Maximum message length. Set to an integer, or the special value MQCBD_FULL_MSG_LENGTH
PCF Parameter:	MQIACH_MAX_MSG_LENGTH
Trace level:	2

Type	MQCFIN
------	--------

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type	MQCFIN

ResolvedQName

Description:	The queue name referred to by the ObjectHandle, when ResolvedType is MQOT_Q.
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level:	2
Type	MQCFST
Length:	MQ_Q_NAME_LENGTH.

ResObjectString

Description:	The object name referred to by the ObjectHandle, when ResolvedType is MQOT_TOPIC.
PCF Parameter:	MQCACF_RESOLVED_OBJECT_STRING
Trace level:	2
Type	MQCFST
Length:	Length varies.

ResolvedType

Description:	The type of the object referred to by the ObjectHandle. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter:	MQIACF_RESOLVED_TYPE
Trace level:	2
Type	MQCFIN

Callback DescriptorStructure

Description:	The MQCBD structure. This parameter is omitted if a NULL MQCBC value is passed to the MQCB call.
PCF Parameter:	MQBACF_MQCBD_STRUCT
Trace level:	3
Type	MQCFBS

Length: The length in bytes of the MQCBC structure

MsgDescStructure

Description: The MQMD structure. The MsgDescStructure parameter is omitted if a NULL MQMD value is passed to the MQCB call.

PCF Parameter: MQBACF_MQMD_STRUCT

Trace level: 3

Type MQCFBS

Length: The length in bytes of the MQMD structure (actual size depends on structure version)

GetMsgOptsStructure

Description: The MQGMO structure. This parameter is omitted if a NULL MQGMO value is passed to the MQCB call.

PCF Parameter: MQBACF_MQGMO_STRUCT

Trace level: 3

Type MQCFBS

Length: The length in bytes of the MQGMO structure (actual size depends on structure version)

QMGrOpDuration

Description: Approximate API call duration, in microseconds, within the queue manager.
The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client.

Note: The accuracy of this timer varies according to the platform that your enterprise uses.

PCF Parameter: MQIAMO64_QMGR_OP_DURATION

Trace level: 2

Type MQCFIN64

MQCLOSE

Application has started the MQCLOSE MQI function

ObjectHandle

Description: The object handle

PCF Parameter: MQIACF_HOBJ

Trace level: 1

Type MQCFIN

CloseOptions

Description: Close options

PCF Parameter: MQIACF_CLOSE_OPTIONS

Trace level: 1

Type MQCFIN

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type	MQCFIN

ResolvedQName

Description:	The queue name referred to by the ObjectHandle, when ResolvedType is MQOT_Q.
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level:	2
Type	MQCFST
Length:	MQ_Q_NAME_LENGTH.

ResObjectString

Description:	The object name referred to by the ObjectHandle, when ResolvedType is MQOT_TOPIC.
PCF Parameter:	MQCACF_RESOLVED_OBJECT_STRING
Trace level:	2
Type	MQCFST
Length:	Length varies.

ResolvedType

Description:	The type of the object referred to by the ObjectHandle. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter:	MQIACF_RESOLVED_TYPE
Trace level:	2
Type	MQCFIN

QMgrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION

Trace level:	2
Type	MQCFIN64

MQCMIT

Application has started the MQCMIT MQI function

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type	MQCFIN

QMGrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQCONN and MQCONNX

Application has started the MQCONN or MQCONNX MQI function

ConnectionId

Description:	The Connection ID if available or MQCONNID_NONE if not
PCF Parameter:	MQBACF_CONNECTION_ID
Trace level:	1
Type:	MQCFBS
Maximum length:	MQ_CONNECTION_ID_LENGTH

QueueManagerName

Description:	The (unresolved) name of the queue manager used in the MQCONN(X) call
PCF Parameter:	MQCA_Q_MGR_NAME
Trace level:	1
Type:	MQCFST
Maximum length:	MQ_Q_MGR_NAME_LENGTH

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type:	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type:	MQCFIN

ConnectOptions

Description:	Connect Options Derived from MQCNO_* values Note: MQCONNX only
PCF Parameter:	MQIACF_CONNECT_OPTIONS
Trace level:	2
Type:	MQCFIN

ConnectionOptionsStructure

Description:	The MQCNO structure. Note: MQCONNX only)
PCF Parameter:	MQBACF_MQCNO_STRUCT
Trace level:	3
Type:	MQCFBS
Maximum length:	The length in bytes of the MQCNO structure (actual size depends on structure version)

ChannelDefinitionStructure

Description:	The MQCD structure. Note: Client connections only
PCF Parameter:	MQBACF_MQCD_STRUCT
Trace level:	3
Type:	MQCFBS
Maximum length:	The length in bytes of the MQCD structure (actual size depends on structure version)

QMgrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQCTL

Application has started the MQCTL MQI function

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type:	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type:	MQCFIN

CtlOperation

Description:	One of MQOP_* values
PCF Parameter:	MQIACF_CTL_OPERATION
Trace level:	1
Type:	MQCFIN

QMgrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQDISC

Application has started the MQDISC MQI function

CompCode

Description: The completion code indicating the result of the operation
PCF Parameter: MQIACF_COMP_CODE
Trace level: 1
Type: MQCFIN

Reason

Description: The reason code result of the operation
PCF Parameter: MQIACF_REASON_CODE
Trace level: 1
Type: MQCFIN

MQGET

Application has started the MQGET MQI function

ObjectHandle

Description: The object handle
PCF Parameter: MQIACF_HOBJ
Trace level: 1
Type: MQCFIN

GetOptions

Description: The get options from MQGMO.Options
PCF Parameter: MQIACF_GET_OPTIONS
Trace level: 1
Type: MQCFIN

CompCode

Description: The completion code indicating the result of the operation
PCF Parameter: MQIACF_COMP_CODE
Trace level: 1
Type: MQCFIN

Reason

Description: The reason code result of the operation
PCF Parameter: MQIACF_REASON_CODE
Trace level: 1
Type: MQCFIN

MsgBuffer

Description: Message data. If TRACEDATA=NONE then this parameter is omitted

PCF Parameter:	MQBACF_MESSAGE_DATA
Trace level:	1
Type:	MQCFBS
Maximum length:	Length is governed by the TRACEDATA() parameter set in the APPTTRACE configuration. (Included in the trace message as MQIACF_TRACE_DATA_LENGTH).

MsgLength

Description:	Length of the message.
PCF Parameter:	MQIACF_MSG_LENGTH
Trace level:	1
Type:	MQCFIN

HighResTime

Description:	Time of operation in microseconds since midnight, January 1 1970 (UTC) Note: The accuracy of this timer varies according to platform support for high a resolution timer
PCF Parameter:	MQIAMO64_HIGHRES_TIME
Trace level:	2
Type:	MQCFIN64

BufferLength

Description:	Length of the buffer provided by the application
PCF Parameter:	MQIACF_BUFFER_LENGTH
Trace level:	2
Type:	MQCFIN

ObjectName

Description:	The name of the opened object
PCF Parameter:	MQCACF_OBJECT_NAME
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

ResolvedQName

Description:	The local name of the queue from which the message was retrieved.
PCF Parameter:	MQCACF_RESOLVED_Q_NAME
Trace level:	2
Type:	MQCFST
Maximum length:	MQ_Q_NAME_LENGTH

ReportOptions

Description: Message report options
PCF Parameter: MQIACF_REPORT
Trace level: 2
Type: MQCFIN

MsgType

Description: Type of message
PCF Parameter: MQIACF_MSG_TYPE
Trace level: 2
Type: MQCFIN

Expiry

Description: Message lifetime
PCF Parameter: MQIACF_EXPIRY
Trace level: 2
Type: MQCFIN

Format

Description: Format name of message data
PCF Parameter: MQCACH_FORMAT_NAME
Trace level: 2
Type: MQCFST
Maximum length: MQ_FORMAT_LENGTH

Priority

Description: Message priority
PCF Parameter: MQIACF_PRIORITY
Trace level: 2
Type: MQCFIN

Persistence

Description: Message persistence
PCF Parameter: MQIACF_PERSISTENCE
Trace level: 2
Type: MQCFIN

MsgId

Description: Message identifier
PCF Parameter: MQBACF_MSG_ID
Trace level: 2

Type: MQCFBS
Maximum length: MQ_MSG_ID_LENGTH

CorrelId

Description: Correlation identifier
PCF Parameter: MQBACF_CORREL_ID
Trace level: 2
Type: MQCFBS
Maximum length: MQ_CORREL_ID_LENGTH

ReplyToQueue

Description:
PCF Parameter: MQCACF_REPLY_TO_Q
Trace level: 2
Type: MQCFST
Maximum length: MQ_Q_NAME_LENGTH

ReplyToQMgr

Description:
PCF Parameter: MQCACF_REPLY_TO_Q_MGR
Trace level: 2
Type: MQCFST
Maximum length: MQ_Q_MGR_NAME_LENGTH

CodedCharSetId

Description: Character set identifier of message data
PCF Parameter: MQIA_CODED_CHAR_SET_ID
Trace level: 2
Type: MQCFIN

Encoding

Description: Numeric encoding of message data.
PCF Parameter: MQIACF_ENCODING
Trace level: 2
Type: MQCFIN

PutDate

Description:
PCF Parameter: MQCACF_PUT_DATE
Trace level: 2
Type: MQCFST

Maximum length: MQ_PUT_DATE_LENGTH

PutTime

Description:

PCF Parameter: MQCACF_PUT_TIME

Trace level: 2

Type: MQCFST

Maximum length: MQ_PUT_TIME_LENGTH

ResolvedQName

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT_Q.

PCF Parameter: MQCACF_RESOLVED_LOCAL_Q_NAME

Trace level: 2

Type: MQCFST

Length: MQ_Q_NAME_LENGTH.

ResObjectString

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT_TOPIC.

PCF Parameter: MQCACF_RESOLVED_OBJECT_STRING

Trace level: 2

Type: MQCFST

Length: Length varies.

ResolvedType

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.

PCF Parameter: MQIACF_RESOLVED_TYPE

Trace level: 2

Type: MQCFIN

PolicyName

Description: The policy name that was applied to this message.

Note: AMS protected messages only

PCF Parameter: MQCA_POLICY_NAME

Trace level: 2

Type: MQCFST

Length: MQ_OBJECT_NAME_LENGTH

XmitqMsgId

Description:	The message ID of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQBACF_XQH_MSG_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_MSG_ID_LENGTH

XmitqCorrelId

Description:	The correlation ID of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQBACF_XQH_CORREL_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_CORREL_ID_LENGTH

XmitqPutTime

Description:	The put time of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQCACF_XQH_PUT_TIME
Trace level:	2
Type:	MQCFST
Length:	MQ_PUT_TIME_LENGTH

XmitqPutDate

Description:	The put date of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQCACF_XQH_PUT_DATE
Trace level:	2
Type:	MQCFST
Length:	MQ_PUT_DATE_LENGTH

XmitqRemoteQName

Description:	The remote queue destination of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQCACF_XQH_REMOTE_Q_NAME
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

XmitqRemoteQMGr

Description: The remote queue manager destination of the message in the transmission queue header.

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter: MQCACF_XQH_REMOTE_Q_MGR

Trace level: 2

Type: MQCFST

Length: MQ_Q_NAME_LENGTH

MsgDescStructure

Description: The MQMD structure.

PCF Parameter: MQBACF_MQMD_STRUCT

Trace level: 3

Type: MQCFBS

Maximum length: The length in bytes of the MQMD structure (actual size depends on structure version)

GetMsgOptsStructure

Description: The MQGMO structure.

PCF Parameter: MQBACF_MQGMO_STRUCT

Trace level: 3

Type: MQCFBS

Maximum length: The length in bytes of the MQGMO structure (actual size depends on structure version)

QMGrOpDuration

Description: Approximate API call duration, in microseconds, within the queue manager.

The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client.

Note: The accuracy of this timer varies according to the platform that your enterprise uses.

PCF Parameter: MQIAMO64_QMGR_OP_DURATION

Trace level: 2

Type: MQCFIN64

MQINQ

Application has started the MQINQ MQI function

ObjectHandle

Description: The object handle

PCF Parameter: MQIACF_HOBJ

Trace level: 1

Type: MQCFIN

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type:	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type:	MQCFIN

SelectorCount

Description:	The count of selectors that are supplied in the Selectors array.
PCF Parameter:	MQIACF_SELECTOR_COUNT
Trace level:	2
Type:	MQCFIN

Selectors

Description:	The list of attributes (integer or character) whose values must be returned by MQINQ.
PCF Parameter:	MQIACF_SELECTORS
Trace level:	2
Type:	MQCFIL

ResolvedQName

Description:	The queue name referred to by the ObjectHandle, when ResolvedType is MQOT_Q.
PCF Parameter:	MQCACF_RESOLVED_Q_NAME
Trace level:	2
Type:	MQCFST
Maximum length:	MQ_Q_NAME_LENGTH

ResObjectString

Description:	The object name referred to by the ObjectHandle, when ResolvedType is MQOT_TOPIC.
PCF Parameter:	MQCACF_RESOLVED_OBJECT_STRING
Trace level:	2
Type:	MQCFST
Maximum length:	Length varies

ResolvedType

Description:	The type of the object referred to by the ObjectHandle. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter:	MQIACF_RESOLVED_TYPE
Trace level:	2
Type:	MQCFIN

IntAttrCount

Description:	The number of integer attributes returned by the inquire operation
PCF Parameter:	MQIACF_INTATTR_COUNT
Trace level:	3
Type:	MQCFIN

IntAttrs

Description:	The integer attribute values returned by the inquire operation. This parameter is only present if IntAttrCount is > 0 when MQINQ returns.
PCF Parameter:	MQIACF_INT_ATTRS
Trace level:	3
Type:	MQCFIL

CharAttrs

Description:	The character attributes returned by the inquire operation. The values are concatenated together. This parameter is only included if CharAttrLength is > 0 when MQINQ returns.
PCF Parameter:	MQCACF_CHAR_ATTRS
Trace level:	3
Type:	MQCFST

QMGrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQOPEN

Application has started the MQOPEN MQI function

ObjectType

Description:	The object type passed in MQOT.ObjectType
PCF Parameter:	MQIACF_OBJECT_TYPE

Trace level: 1
Type: MQCFIN

ObjectName

Description: The name of the object passed to the MQI call before any queue name resolution is attempted.
PCF Parameter: MQCACF_OBJECT_NAME
Trace level: 1
Type: MQCFST
Maximum length: MQ_Q_NAME_LENGTH

ObjectQMgrName

Description: The name of the object queue manager passed to the MQI call before any queue name resolution is attempted.
PCF Parameter: MQCACF_OBJECT_Q_MGR_NAME
Trace level: 1
Type: MQCFST
Maximum length: MQ_Q_MGR_NAME_LENGTH

ObjectHandle

Description: The object handle
PCF Parameter: MQIACF_HOBJ
Trace level: 1
Type: MQCFIN

CompCode

Description: The completion code indicating the result of the operation
PCF Parameter: MQIACF_COMP_CODE
Trace level: 1
Type: MQCFIN

Reason

Description: The reason code result of the operation
PCF Parameter: MQIACF_REASON_CODE
Trace level: 1
Type: MQCFIN

OpenOptions

Description: Options used to open the object
PCF Parameter: MQIACF_OPEN_OPTIONS
Trace level: 1
Type: MQCFIN

AlternateUserId

Description: Only included if MQOO_ALTERNATE_USER_AUTHORITY is specified
PCF Parameter: MQCACF_ALTERNATE_USERID
Trace level: 2
Type: MQCFST
Maximum length: MQ_USER_ID_LENGTH

RecsPresent

Description: The number of object name records present. Only included if MQOD Version >= MQOD_VERSION_2
PCF Parameter: MQIACF_RECS_PRESENT
Trace level: 1
Type: MQCFIN

KnownDestCount

Description: Number of local queues opened successfully Only included if MQOD Version >= MQOD_VERSION_2
PCF Parameter: MQIACF_KNOWN_DEST_COUNT
Trace level: 1
Type: MQCFIN

UnknownDestCount

Description: Number of remote queues opened successfully Only included if MQOD Version >= MQOD_VERSION_2
PCF Parameter: MQIACF_UNKNOWN_DEST_COUNT
Trace level: 1
Type: MQCFIN

InvalidDestCount

Description: Number of queues that failed to open Only included if MQOD Version >= MQOD_VERSION_2
PCF Parameter: MQIACF_INVALID_DEST_COUNT
Trace level: 1
Type: MQCFIN

DynamicQName

Description: The dynamic queue name passed as input to the MQOPEN call.
PCF Parameter: MQCACF_DYNAMIC_Q_NAME
Trace level: 2
Type: MQCFST
Maximum length: MQ_Q_NAME_LENGTH

ResolvedLocalQName ^{1 2}

Description:	Contains the local queue name after name resolution has been carried out. (e.g. for remote queues this will be the name of the transmit queue)
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level:	2
Type:	MQCFST
Range:	If MQOD.Version is less than MQOD_VERSION_3 this contains the value of the MQOD.ObjectName field after the MQOPEN call has completed. If MQOD.Version is equal or greater than MQOD_VERSION_3 this contains the value in the MQOD.ResolvedQName field.
Maximum length:	MQ_Q_NAME_LENGTH

ResolvedLocalQMgrName ^{1 2}

Description:	The local queue manager name after name resolution has been performed.
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_MGR
Trace level:	2
Type:	MQCFST
Range:	Only if MQOD.Version >= MQOD_VERSION_3
Maximum length:	MQ_Q_MGR_NAME_LENGTH

ResolvedQName ^{1 2}

Description:	The queue name after name resolution has been carried out.
PCF Parameter:	MQCACF_RESOLVED_Q_NAME
Trace level:	2
Type:	MQCFST
Range:	If MQOD.Version is less than MQOD_VERSION_3 this contains the value of the MQOD.ObjectName field after the MQOPEN call has completed. If MQOD.Version is equal or greater than MQOD_VERSION_3 this contains the value in the MQOD.ResolvedQName field.
Maximum length:	MQ_Q_NAME_LENGTH

ResolvedQMgrName ^{1 2}

Description:	Contains the queue manager name after name resolution has been carried out. If MQOD.Version is less than MQOD_VERSION_3 this contains the value of the MQOD.ObjectQMgrName field after the MQOPEN call has completed. If MQOD.Version is equal or greater than MQOD_VERSION_3 this contains the value in the MQOD.ResolvedQMgrName field.
PCF Parameter:	MQCACF_RESOLVED_Q_MGR
Trace level:	2
Type:	MQCFST
Maximum length:	MQ_Q_MGR_NAME_LENGTH

AlternateSecurityId

Description:	Alternative security identifier. Only present if MQOD.Version is equal or greater than MQOD_VERSION_3, MQOO_ALTERNATE_USER_AUTHORITY is specified, and MQOD.AlternateSecurityId is not equal to MQSID_NONE.
PCF Parameter:	MQBACF_ALTERNATE_SECURITYID
Trace level:	2
Type:	MQCFBS
Maximum length:	MQ_SECURITY_ID_LENGTH

ObjectString

Description:	Long object name. Only included if MQOD.Version is equal or greater than MQOD_VERSION_4 and the VSLength field of MQOD.ObjectString is MQVS_NULL_TERMINATED or greater than zero.
PCF Parameter:	MQCACF_OBJECT_STRING
Trace level:	2
Type:	MQCFST
Maximum length:	Length varies.

SelectionString

Description:	Selection string. Only included if MQOD.Version is equal or greater than MQOD_VERSION_4 and the VSLength field of MQOD.SelectionString is MQVS_NULL_TERMINATED or greater than zero.
PCF Parameter:	MQCACF_SELECTION_STRING
Trace level:	2
Type:	MQCFST
Maximum length:	Length varies.

ResObjectString

Description:	The long object name after the queue manager resolves the name provided in the ObjectName field. Only included for topics and queue aliases that reference a topic object if MQOD.Version is equal or greater than MQOD_VERSION_4 and VSLength is MQVS_NULL_TERMINATED or greater than zero.
PCF Parameter:	MQCACF_RESOLVED_OBJECT_STRING
Trace level:	2
Type:	MQCFST
Maximum length:	Length varies.

ResolvedType

Description:	The type of the resolved (base) object being opened. Only included if MQOD.Version is equal or greater than MQOD_VERSION_4. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter:	MQIACF_RESOLVED_TYPE
Trace level:	2
Type:	MQCFIN

QMGrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

Application Activity Distribution List PCF Group Header Structure

If the MQOPEN function opens a distribution list, then the MQOPEN parameters includes one AppActivityDistList PCF group for each of the queues in the distribution list up to the number of structures numbered in RecsPresent. The AppActivityDistList PCF group combines information from the MQOR, and MQRR structures to identify the queue name, and indicate the result of the open operation on the queue. An AppActivityDistList group always starts with the following MQCFGR structure:

Table 33. AppActivityDistList group MQCFGR structure		
MQCFGR field	Value	Description
Type	MQCFT_GROUP	
StrucLength	Length in bytes of the MQCFGR structure	
Parameter	MQGACF_APP_DIST_LIST	Distribution list group parameter
ParameterCount	4	The number of parameter structures following the MQCFGR structure that are contained within this group.

ObjectName

Description:	The name of a queue in the distribution list MQ_Q_NAME_LENGTH. Only included if MQOR structures are provided.
PCF Parameter:	MQCACF_OBJECT_NAME
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH. Only included if MQOR structures are provided.

ObjectQMGrName

Description:	The name of the queue manager on which the queue named in ObjectName is defined.
PCF Parameter:	MQCACF_OBJECT_Q_MGR_NAME
Trace level:	2

¹ This parameter is only included if the object being opened resolves to a queue, and the queue is opened for MQOO_INPUT_*, MQOO_OUTPUT, or MQOO_BROWSE

² The ResolvedLocalQName parameter is only included if it is different from the ResolvedQName parameter.

Type:	MQCFST
Length:	MQ_Q_MGR_NAME_LENGTH. Only included if MQOR structures are provided.

CompCode

Description:	The completion code indicating the result of the open for this object. Only included if MQRR structures are provided and the reason code for the MQOPEN is MQRC_MULTIPLE_REASONS
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	2
Type:	MQCFIN

Reason

Description:	The reason code indicating the result of the open for this object. Only included if MQRR structures are provided and the reason code for the MQOPEN is MQRC_MULTIPLE_REASONS
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	2
Type:	MQCFIN

MQPUT

Application has started the MQPUT MQI function.

ObjectHandle

Description:	The object handle
PCF Parameter:	MQIACF_HOBJ
Trace level:	1
Type:	MQCFIN

PutOptions

Description:	The put options from MQPMO.Options
PCF Parameter:	MQIACF_PUT_OPTIONS
Trace level:	1
Type:	MQCFIN

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type:	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1

Type: MQCFIN

MsgBuffer

Description: Message data.

PCF Parameter: MQBACF_MESSAGE_DATA

Trace level: 1

Type: MQCFBS

Length: Length is governed by the TRACEDATA() parameter set in the APPTRACE configuration. If TRACEDATA=NONE then this parameter is omitted.

MsgLength

Description: Length of the message.

PCF Parameter: MQIACF_MSG_LENGTH

Trace level: 1

Type: MQCFIN

RecsPresent

Description: The number of put message records or response records present. Only included if MQPMO Version >= MQPMO_VERSION_2

PCF Parameter: MQIACF_RECS_PRESENT

Trace level: 1

Type: MQCFIN

KnownDestCount

Description: Number of messages sent successfully to local queues

PCF Parameter: MQIACF_KNOWN_DEST_COUNT

Trace level: 1

Type: MQCFIN

UnknownDestCount

Description: Number of messages sent successfully to remote queues

PCF Parameter: MQIACF_UNKNOWN_DEST_COUNT

Trace level: 1

Type: MQCFIN

InvalidDestCount

Description: Number of messages that could not be sent

PCF Parameter: MQIACF_INVALID_DEST_COUNT

Trace level: 1

Type: MQCFIN

HighResTime

Description:	Time of operation in microseconds since midnight, January 1st 1970 (UTC) Note: The accuracy of this timer varies according to platform support for high a resolution timer.
PCF Parameter:	MQIAMO64_HIGHRES_TIME
Trace level:	2
Type:	MQCFIN64

ObjectName

Description:	The name of the opened object.
PCF Parameter:	MQCACF_OBJECT_NAME
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

ResolvedQName

Description:	The name of the queue after queue name resolution has been performed.
PCF Parameter:	MQCACF_RESOLVED_Q_NAME
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

ResolvedQMgrName

Description:	The queue manager name after name resolution has been performed.
PCF Parameter:	MQCACF_RESOLVED_Q_MGR
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_MGR_NAME_LENGTH

ResolvedLocalQName ³

Description:	Contains the local queue name after name resolution has been carried out.
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level:	2
Type:	MQCFST

ResolvedLocalQMgrName ³

Description:	Contains the local queue manager name after name resolution has been carried out.
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_MGR
Trace level:	2
Type:	MQCFST

Length: MQ_Q_MGR_NAME_LENGTH

ReportOptions

Description: Message report options
PCF Parameter: MQIACF_REPORT
Trace level: 2
Type: MQCFIN

MsgType

Description: Type of message
PCF Parameter: MQIACF_MSG_TYPE
Trace level: 2
Type: MQCFIN

Expiry

Description: Message lifetime
PCF Parameter: MQIACF_EXPIRY
Trace level: 2
Type: MQCFIN

Format

Description: Format name of message data
PCF Parameter: MQCACH_FORMAT_NAME
Trace level: 2
Type: MQCFST
Length: MQ_FORMAT_LENGTH

Priority

Description: Message priority
PCF Parameter: MQIACF_PRIORITY
Trace level: 2
Type: MQCFIN

Persistence

Description: Message persistence
PCF Parameter: MQIACF_PERSISTENCE
Trace level: 2
Type: MQCFIN

MsgId

Description: Message identifier
PCF Parameter: MQBACF_MSG_ID

Trace level: 2
Type: MQCFBS
Length: MQ_MSG_ID_LENGTH

CorrelId

Description: Correlation identifier
PCF Parameter: MQBACF_CORREL_ID
Trace level: 2
Type: MQCFBS
Length: MQ_CORREL_ID_LENGTH

ReplyToQueue

Description:
PCF Parameter: MQCACF_REPLY_TO_Q
Trace level: 2
Type: MQCFST
Length: MQ_Q_NAME_LENGTH

ReplyToQMgr

Description:
PCF Parameter: MQCACF_REPLY_TO_Q_MGR
Trace level: 2
Type: MQCFST
Length: MQ_Q_MGR_NAME_LENGTH

CodedCharSetId

Description: Character set identifier of message data
PCF Parameter: MQIA_CODED_CHAR_SET_ID
Trace level: 2
Type: MQCFIN

Encoding

Description: Numeric encoding of message data.
PCF Parameter: MQIACF_ENCODING
Trace level: 2
Type: MQCFIN

PutDate

Description:
PCF Parameter: MQCACF_PUT_DATE
Trace level: 2

Type: MQCFST
Length: MQ_PUT_DATE_LENGTH

PutTime

Description:
PCF Parameter: MQCACF_PUT_TIME
Trace level: 2
Type: MQCFST
Length: MQ_PUT_TIME_LENGTH

ResolvedQName

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT_Q.
PCF Parameter: MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level: 2
Type: MQCFST
Length: MQ_Q_NAME_LENGTH.

ResObjectString

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT_TOPIC.
PCF Parameter: MQCACF_RESOLVED_OBJECT_STRING
Trace level: 2
Type: MQCFST
Length: Length varies.

ResolvedType

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter: MQIACF_RESOLVED_TYPE
Trace level: 2
Type: MQCFIN

PolicyName

Description: The policy name that was applied to this message.
Note: AMS protected messages only
PCF Parameter: MQCA_POLICY_NAME
Trace level: 2
Type: MQCFST
Length: MQ_OBJECT_NAME_LENGTH

XmitqMsgId

Description:	The message ID of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQBACF_XQH_MSG_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_MSG_ID_LENGTH

XmitqCorrelId

Description:	The correlation ID of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQBACF_XQH_CORREL_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_CORREL_ID_LENGTH

XmitqPutTime

Description:	The put time of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQCACF_XQH_PUT_TIME
Trace level:	2
Type:	MQCFST
Length:	MQ_PUT_TIME_LENGTH

XmitqPutDate

Description:	The put date of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQCACF_XQH_PUT_DATE
Trace level:	2
Type:	MQCFST
Length:	MQ_PUT_DATE_LENGTH

XmitqRemoteQName

Description:	The remote queue destination of the message in the transmission queue header. Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter:	MQCACF_XQH_REMOTE_Q_NAME
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

XmitqRemoteQMGr

Description: The remote queue manager destination of the message in the transmission queue header.

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter: MQCACF_XQH_REMOTE_Q_MGR

Trace level: 2

Type: MQCFST

Length: MQ_Q_NAME_LENGTH

PutMsgOptsStructure

Description: The MQPMO structure.

PCF Parameter: MQBACF_MQPMO_STRUCT

Trace level: 3

Type: MQCFBS

Length: The length in bytes of the MQPMO structure (actual size depends on structure version)

QMGrOpDuration

Description: Approximate API call duration, in microseconds, within the queue manager.

The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client.

Note: The accuracy of this timer varies according to the platform that your enterprise uses.

PCF Parameter: MQIAMO64_QMGR_OP_DURATION

Trace level: 2

Type: MQCFIN64

MQPUT Application Activity Distribution List PCF Group Header Structure

If the MQPUT function is putting to a distribution list, then the MQPUT parameters include one AppActivityDistList PCF group. For each of the queues in the distribution list, see “Application Activity Distribution List PCF Group Header Structure” on page 273. The AppActivityDistList PCF group combines information from the MQPMR, and MQRR structures to identify the PUT parameters, and indicate the result of the PUT operation on each queue. For MQPUT operations the AppActivityDistList group contains some or all of the following parameters (the CompCode and Reason is present if the reason code is MQRC_MULTIPLE_REASONS and the other parameters are determined by the MQPMO.PutMsgRecFields field):

CompCode

Description: The completion code indicating the result of the operation. Only included if MQRR structures are provided and the reason code for the MQPUT is MQRC_MULTIPLE_REASONS

PCF Parameter: MQIACF_COMP_CODE

³ The ResolvedLocalQName parameter is only included if it is different from the ResolvedQName parameter.

Trace level:	2
Type:	MQCFIN

Reason

Description:	The reason code indicating the result of the put for this object. Only included if MQRR structures are provided and the reason code for the MQPUT is MQRC_MULTIPLE_REASONS
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	2
Type:	MQCFIN

MsgId

Description:	Message identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_MSG_ID
PCF Parameter:	MQBACF_MSG_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_MSG_ID_LENGTH

CorrelId

Description:	Correlation identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_CORREL_ID
PCF Parameter:	MQBACF_CORREL_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_CORREL_ID_LENGTH

GroupId

Description:	Group identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_GROUP_ID
PCF Parameter:	MQBACF_GROUP_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_GROUP_ID_LENGTH

Feedback

Description:	Feedback. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_FEEDBACK
PCF Parameter:	MQIACF_FEEDBACK
Trace level:	2
Type:	MQCFIN

AccountingToken

Description:	AccountingToken. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_ACCOUNTING_TOKEN
PCF Parameter:	MQBACF_ACCOUNTING_TOKEN
Trace level:	2
Type:	MQCFBS
Length:	MQ_ACCOUNTING_TOKEN_LENGTH.

MQPUT1

Application has started the MQPUT1 MQI function

ObjectType

Description:	The object type passed in MQOT.ObjectType
PCF Parameter:	MQIACF_OBJECT_TYPE
Trace level:	1
Type:	MQCFIN

ObjectName

Description:	The name of the object passed to the MQI call before any queue name resolution is attempted.
PCF Parameter:	MQCACF_OBJECT_NAME
Trace level:	1
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

ObjectQMgrName

Description:	The name of the object queue manager passed to the MQI call before any queue name resolution is attempted.
PCF Parameter:	MQCACF_OBJECT_Q_MGR_NAME
Trace level:	1
Type:	MQCFST
Length:	MQ_Q_MGR_NAME_LENGTH

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type:	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1

Type: MQCFIN

PutOptions

Description: The put options from MQPMO.Options

PCF Parameter: MQIACF_PUT_OPTIONS

Trace level: 1

Type: MQCFIN

AlternateUserId

Description: Only included if MQPMO_ALTERNATE_USER_AUTHORITY is specified.

PCF Parameter: MQCACF_ALTERNATE_USERID

Trace level: 2

Type: MQCFST

Length: MQ_USER_ID_LENGTH

RecsPresent

Description: The number of object name records present

PCF Parameter: MQIACF_RECS_PRESENT

Trace level: 1

Type: MQCFIN

KnownDestCount

Description: Number of local queues opened successfully

PCF Parameter: MQIACF_KNOWN_DEST_COUNT

Trace level: 1

Type: MQCFIN

UnknownDestCount

Description: Number of remote queues opened successfully

PCF Parameter: MQIACF_UNKNOWN_DEST_COUNT

Trace level: 1

Type: MQCFIN

InvalidDestCount

Description: Number of queues that failed to open

PCF Parameter: MQIACF_INVALID_DEST_COUNT

Trace level: 1

Type: MQCFIN

MsgBuffer

Description: Message data.

PCF Parameter: MQBACF_MESSAGE_DATA

Trace level:	1
Type:	MQCFBS
Length:	Length is governed by the TRACEDATA() parameter set in the APPTRACE configuration. If TRACEDATA=NONE then this parameter is omitted.

MsgLength

Description:	Length of the message.
PCF Parameter:	MQIACF_MSG_LENGTH
Trace level:	1
Type:	MQCFIN

HighResTime

Description:	Time of operation in microseconds since midnight, January 1st 1970 (UTC) Note: The accuracy of this timer will vary according to platform support for high a resolution timer.
PCF Parameter:	MQIAMO64_HIGHRES_TIME
Trace level:	2
Type:	MQCFIN64

ResolvedQName

Description:	The name of the queue after queue name resolution has been performed.
PCF Parameter:	MQCACF_RESOLVED_Q_NAME
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

ResolvedQMgrName

Description:	The queue manager name after name resolution has been performed.
PCF Parameter:	MQCACF_RESOLVED_Q_MGR
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_MGR_NAME_LENGTH

ResolvedLocalQName ⁴

Description:	Contains the local queue name after name resolution has been carried out
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level:	2
Type:	MQCFST

ResolvedLocalQMgrName ⁴

Description:	Contains the local queue manager name after name resolution has been carried out.
--------------	---

PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_MGR
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_MGR_NAME_LENGTH

AlternateSecurityId

Description:	Alternate security identifier. Only present if MQOD.Version is equal or greater than MQOD_VERSION_3 and MQOD.AlternateSecurityId is not equal to MQSID_NONE.
PCF Parameter:	MQBACF_ALTERNATE_SECURITYID
Trace level:	2
Type:	MQCFBS
Length:	MQ_SECURITY_ID_LENGTH

ObjectString

Description:	Long object name. Only included if MQOD.Version is equal or greater than MQOD_VERSION_4 and the VSLength field of MQOD.ObjectString is MQVS_NULL_TERMINATED or greater than zero.
PCF Parameter:	MQCACF_OBJECT_STRING
Trace level:	2
Type:	MQCFST
Length:	Length varies.

ResObjectString

Description:	The long object name after the queue manager resolves the name provided in the ObjectName field. Only included for topics and queue aliases that reference a topic object if MQOD.Version is equal or greater than MQOD_VERSION_4 and VSLength is MQVS_NULL_TERMINATED or greater than zero.
PCF Parameter:	MQCACF_RESOLVED_OBJECT_STRING
Trace level:	2
Type:	MQCFST
Length:	Length varies.

ResolvedType

Description:	The type of the resolved (base) object being opened. Only included if MQOD.Version is equal or greater than MQOD_VERSION_4. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter:	MQIACF_RESOLVED_TYPE
Trace level:	2
Type:	MQCFIN

ReportOptions

Description:	Message report options
PCF Parameter:	MQIACF_REPORT

Trace level: 2
Type: MQCFIN

MsgType

Description: Type of message
PCF Parameter: MQIACF_MSG_TYPE
Trace level: 2
Type: MQCFIN

Expiry

Description: Message lifetime
PCF Parameter: MQIACF_EXPIRY
Trace level: 2
Type: MQCFIN

Format

Description: Format name of message data
PCF Parameter: MQCACH_FORMAT_NAME
Trace level: 2
Type: MQCFST
Length: MQ_FORMAT_LENGTH

Priority

Description: Message priority
PCF Parameter: MQIACF_PRIORITY
Trace level: 2
Type: MQCFIN

Persistence

Description: Message persistence
PCF Parameter: MQIACF_PERSISTENCE
Trace level: 2
Type: MQCFIN

MsgId

Description: Message identifier
PCF Parameter: MQBACF_MSG_ID
Trace level: 2
Type: MQCFBS
Length: MQ_MSG_ID_LENGTH

CorrelId

PCF Parameter: Correlation identifier
Description: MQBACF_CORREL_ID
Trace level: 2
Type: MQCFBS
Length: MQ_CORREL_ID_LENGTH

ReplyToQueue

Description:
PCF Parameter: MQCACF_REPLY_TO_Q
Trace level: 2
Type: MQCFST
Length: MQ_Q_NAME_LENGTH

ReplyToQMgr

Description:
PCF Parameter: MQCACF_REPLY_TO_Q_MGR
Trace level: 2
Type: MQCFST
Length: MQCFST

CodedCharSetId

Description: Character set identifier of message data
PCF Parameter: MQIA_CODED_CHAR_SET_ID
Trace level: 2
Type: MQCFIN

Encoding

Description: Numeric encoding of message data.
PCF Parameter: MQIACF_ENCODING
Trace level: 2
Type: MQCFIN

PutDate

Description:
PCF Parameter: MQCACF_PUT_DATE
Trace level: 2
Type: MQCFST
Length: MQ_PUT_DATE_LENGTH

PutTime

Description:
PCF Parameter: MQCACF_PUT_TIME
Trace level: 2
Type: MQCFST
Length: MQ_PUT_TIME_LENGTH

PolicyName

Description: The policy name that was applied to this message.
Note: AMS protected messages only
PCF Parameter: MQCA_POLICY_NAME
Trace level: 2
Type: MQCFST
Length: MQ_OBJECT_NAME_LENGTH

XmitqMsgId

Description: The message ID of the message in the transmission queue header.
Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter: MQBACF_XQH_MSG_ID
Trace level: 2
Type: MQCFBS
Length: MQ_MSG_ID_LENGTH

XmitqCorrelId

Description: The correlation ID of the message in the transmission queue header.
Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter: MQBACF_XQH_CORREL_ID
Trace level: 2
Type: MQCFBS
Length: MQ_CORREL_ID_LENGTH

XmitqPutTime

Description: The put time of the message in the transmission queue header.
Note: Only when Format is MQFMT_XMIT_Q_HEADER
PCF Parameter: MQCACF_XQH_PUT_TIME
Trace level: 2
Type: MQCFST
Length: MQ_PUT_TIME_LENGTH

XmitqPutDate

Description: The put date of the message in the transmission queue header.

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter: MQCACF_XQH_PUT_DATE

Trace level: 2

Type: MQCFST

Length: MQ_PUT_DATE_LENGTH

XmitqRemoteQName

Description: The remote queue destination of the message in the transmission queue header.

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter: MQCACF_XQH_REMOTE_Q_NAME

Trace level: 2

Type: MQCFST

Length: MQ_Q_NAME_LENGTH

XmitqRemoteQMGr

Description: The remote queue manager destination of the message in the transmission queue header.

Note: Only when Format is MQFMT_XMIT_Q_HEADER

PCF Parameter: MQCACF_XQH_REMOTE_Q_MGR

Trace level: 2

Type: MQCFST

Length: MQ_Q_NAME_LENGTH

PutMsgOptsStructure

Description: The MQPMO structure.

PCF Parameter: MQBACF_MQPMO_STRUCT

Trace level: 3

Type: MQCFBS

Length: The length in bytes of the MQPMO structure (actual size depends on structure version)

QMGrOpDuration

Description: Approximate API call duration, in microseconds, within the queue manager.

The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client.

Note: The accuracy of this timer varies according to the platform that your enterprise uses.

PCF Parameter: MQIAMO64_QMGR_OP_DURATION

Trace level: 2

MQPUT1 AppActivityDistList PCF Group Header Structure

If the MQPUT1 function is putting to a distribution list, then the variable parameters include one AppActivityDistList PCF group. For each of the queues in the distribution list, see “Application Activity Distribution List PCF Group Header Structure” on page 273. The AppActivityDistList PCF group combines information from the MQOR, MQPMR, and MQRR structures to identify the objects, and the PUT parameters, and indicate the result of the PUT operation on each queue. For MQPUT1 operations the AppActivityDistList group contains some or all of the following parameters (the CompCode, Reason, ObjectName, and ObjectQMGrName is present if the reason code is MQRC_MULTIPLE_REASONS and the other parameters is determined by the MQPMO.PutMsgRecFields field):

CompCode

Description:	The completion code indicating the result of the put for this object. Only included if MQRR structures are provided and the reason code for the MQPUT1 is MQRC_MULTIPLE_REASONS
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	2
Type:	MQCFIN

Reason

Description:	The reason code indicating the result of the put for this object. Only included if MQRR structures are provided and the reason code for the MQPUT1 is MQRC_MULTIPLE_REASONS
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	2
Type:	MQCFIN

ObjectName

Description:	The name of a queue in the distribution list. Only included if MQOR structures are provided.
PCF Parameter:	MQCACF_OBJECT_NAME
Trace level:	2
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

MsgId

Description:	Message identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_MSG_ID
PCF Parameter:	MQBACF_MSG_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_MSG_ID_LENGTH

⁴ The ResolvedLocalQName parameter is only included if it is different from the ResolvedQName parameter.

CorrelId

Description:	Correlation identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_CORREL_ID
PCF Parameter:	MQBACF_CORREL_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_CORREL_ID_LENGTH

GroupId

Description:	Group identifier. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_GROUP_ID
PCF Parameter:	MQBACF_GROUP_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_GROUP_ID_LENGTH

Feedback

Description:	Feedback. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_FEEDBACK
PCF Parameter:	MQIACF_FEEDBACK
Trace level:	2
Type:	MQCFIN

AccountingToken

Description:	AccountingToken. Only included if MQPMR structures are provided.and PutMsgRecFields includes MQPMRF_ACCOUNTING_TOKEN
PCF Parameter:	MQBACF_ACCOUNTING_TOKEN
Trace level:	2
Type:	MQCFBS
Length:	MQ_ACCOUNTING_TOKEN_LENGTH.

MQSET

Application has started the MQSET MQI function

ObjectHandle

Description:	The object handle
PCF Parameter:	MQIACF_HOBJ
Trace level:	1
Type:	MQCFIN

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE

Trace level: 1
Type: MQCFIN

Reason

Description: The reason code result of the operation
PCF Parameter: MQIACF_REASON_CODE
Trace level: 1
Type: MQCFIN

SelectorCount

Description: The count of selectors that are supplied in the Selectors array.
PCF Parameter: MQIACF_SELECTOR_COUNT
Trace level: 2
Type: MQCFIN

Selectors

Description: The list of attributes (integer or character) whose values are being updated by MQSET.
PCF Parameter: MQIACF_SELECTORS
Trace level: 2
Type: MQCFIL

ResolvedQName

Description: The queue name referred to by the ObjectHandle, when ResolvedType is MQOT_Q.
PCF Parameter: MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level: 2
Type: MQCFST
Length: MQ_Q_NAME_LENGTH.

ResObjectString

Description: The object name referred to by the ObjectHandle, when ResolvedType is MQOT_TOPIC.
PCF Parameter: MQCACF_RESOLVED_OBJECT_STRING
Trace level: 2
Type: MQCFST
Length: Length varies.

ResolvedType

Description: The type of the object referred to by the ObjectHandle. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter: MQIACF_RESOLVED_TYPE
Trace level: 2

Type	MQCFIN
------	--------

IntAttrCount

Description:	The number of integer attributes to be updated by the set operation.
PCF Parameter:	MQIACF_INTATTR_COUNT
Trace level:	3
Type:	MQCFIN

IntAttrs

Description:	The integer attribute values
PCF Parameter:	MQIACF_INT_ATTRS
Trace level:	3
Type:	MQCFIL
Range:	This parameter is only present if IntAttrCount is > 0

CharAttrs

Description:	The character attributes to be updated by the set operation. The values are concatenated together.
PCF Parameter:	MQCACF_CHAR_ATTRS
Trace level:	3
Type:	MQCFST
Range:	This parameter is only included if CharAttrLength is > 0

QMGrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQSUB

Application has started the MQSUB MQI function

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type:	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type:	MQCFIN

SubHandle

Description:	The subscription handle
PCF Parameter:	MQIACF_HSUB
Trace level:	1
Type:	MQCFIN

ObjectHandle

Description:	The object handle
PCF Parameter:	MQIACF_HOBJ
Trace level:	1
Type:	MQCFIN

Options

Description:	Subscription options
PCF Parameter:	MQIACF_SUB_OPTIONS
Trace level:	1
Type:	MQCFIN

ObjectName

Description:	The name of the object.
PCF Parameter:	MQCACF_OBJECT_NAME
Trace level:	1
Type:	MQCFST
Length:	MQ_Q_NAME_LENGTH

ObjectString

Description:	Long object name.
PCF Parameter:	MQCACF_OBJECT_STRING
Trace level:	1
Type:	MQCFST
Range:	Only included if the VSLength field of MQSD.ObjectString is greater than zero or MQVS_NULL_TERMINATED.
Length:	Length varies.

AlternateUserId

Description:	
--------------	--

PCF Parameter:	MQCACF_ALTERNATE_USERID
Trace level:	2
Type:	MQCFST
Range:	Only included if MQSO_ALTERNATE_USER_AUTHORITY is specified.
Length:	MQ_USER_ID_LENGTH

AlternateSecurityId

Description:	Alternate security identifier.
PCF Parameter:	MQBACF_ALTERNATE_SECURITYID
Trace level:	2
Type:	MQCFBS
Range:	Only present if MQSO_ALTERNATE_USER_AUTHORITY is specified and MQSD.AlternateSecurityId is not equal to MQSID_NONE.
Length:	MQ_SECURITY_ID_LENGTH

SubName

Description:	Subscription Name
PCF Parameter:	MQCACF_SUB_NAME
Trace level:	2
Type:	MQCFST
Range:	Only included if the VSLength field of MQSD.SubName is greater than zero or MQVS_NULL_TERMINATED.
Length:	Length varies.

SubUserData

Description:	Subscription User Data
PCF Parameter:	MQCACF_SUB_USER_DATA
Trace level:	2
Type:	MQCFST
Range:	Only included if the VSLength field of MQSD.SubName is greater than zero or MQVS_NULL_TERMINATED.
Length:	Length varies.

SubCorrelId

Description:	Subscription Correlation identifier
PCF Parameter:	MQBACF_SUB_CORREL_ID
Trace level:	2
Type:	MQCFBS
Length:	MQ_CORREL_ID_LENGTH

SelectionString

Description:	Selection string.
--------------	-------------------

PCF Parameter:	MQCACF_SELECTION_STRING
Trace level:	2
Type:	MQCFST
Range:	Only included if the VSLength field of MQSD. SelectionString is MQVS_NULL_TERMINATED or greater than zero.
Length:	Length varies.

ResolvedQName

Description:	The queue name referred to by the ObjectHandle, when ResolvedType is MQOT_Q.
PCF Parameter:	MQCACF_RESOLVED_LOCAL_Q_NAME
Trace level:	2
Type	MQCFST
Length:	MQ_Q_NAME_LENGTH.

ResObjectString

Description:	The object name referred to by the ObjectHandle, when ResolvedType is MQOT_TOPIC.
PCF Parameter:	MQCACF_RESOLVED_OBJECT_STRING
Trace level:	2
Type	MQCFST
Length:	Length varies.

ResolvedType

Description:	The type of the object referred to by the ObjectHandle. Possible values are MQOT_Q, MQOT_TOPIC, or MQOT_NONE.
PCF Parameter:	MQIACF_RESOLVED_TYPE
Trace level:	2
Type	MQCFIN

SubDescriptorStructure

Description:	The MQSD structure.
PCF Parameter:	MQBACF_MQSD_STRUCT
Trace level:	3
Type:	MQCFBS
Length:	The length in bytes of the MQSD structure.

QMgrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
--------------	---

PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQSUBRQ

Application has started the MQSUBRQ MQI function

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type:	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type:	MQCFIN

SubHandle

Description:	The subscription handle
PCF Parameter:	MQIACF_HSUB
Trace level:	1
Type:	MQCFIN

SubOptions

Description:	The sub options from MQSB.Options
PCF Parameter:	MQIACF_SUBRQ_OPTIONS
Trace level:	2
Type:	MQCFIN

Action

Description:	The subscription request action (MQSR_*)
PCF Parameter:	MQIACF_SUBRQ_ACTION
Trace level:	2
Type:	MQCFIN

NumPubs

Description:	The number of publications sent as a result of this call (from MQSB.NumPubs)
PCF Parameter:	MQIACF_NUM_PUBS
Trace level:	2
Type:	MQCFIN

QMGrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

MQSTAT

Application has started the MQSTAT MQI function

CompCode

Description:	The completion code indicating the result of the operation
PCF Parameter:	MQIACF_COMP_CODE
Trace level:	1
Type:	MQCFIN

Reason

Description:	The reason code result of the operation
PCF Parameter:	MQIACF_REASON_CODE
Trace level:	1
Type:	MQCFIN

Type

Description:	Type of status information being requested
PCF Parameter:	MQIACF_STATUS_TYPE
Trace level:	2
Type:	MQCFIN

StatusStructure

Description:	The MQSTS structure.
PCF Parameter:	MQBACF_MQSTS_STRUCT
Trace level:	3
Type:	MQCFBS
Length:	The length in bytes of the MQSTS structure (actual size depends on structure version)

QMGrOpDuration

Description:	Approximate API call duration, in microseconds, within the queue manager. The duration does not include the time spent outside of the queue manager. For example, the time taken as an IBM MQ client. Note: The accuracy of this timer varies according to the platform that your enterprise uses.
PCF Parameter:	MQIAMO64_QMGR_OP_DURATION
Trace level:	2
Type	MQCFIN64

Variable Parameters for Application Activity XA Operations

XA operations are API calls that applications can make to enable MQ to participate in a transaction. The parameters for each operation are defined in the following section.

The trace level indicates the level of trace granularity that is required for the parameters to be included in the trace. The possible trace level values are:

1. Low

The parameter is included when "low", "medium" or "high" activity tracing is configured for an application. This setting means that a parameter is always included in the AppActivityData group for the operation. This set of parameters is sufficient to trace the MQI calls an application makes, and to see if they are successful.

2. Medium

The parameter is only included in the AppActivityData group for the operation when "medium" or "high" activity tracing is configured for an application. This set of parameters adds information about the resources, for example, queue and topic names used by the application.

3. High

The parameter is only included in the AppActivityData group for the operation when "high" activity tracing is configured for an application. This set of parameters includes memory dumps of the structures passed to the MQI and XA functions. For this reason, it contains more information about the parameters used in MQI and XA calls. The structure memory dumps are shallow copies of the structures. To avoid erroneous attempts to dereference pointers, the pointer values in the structures are set to NULL.

Note: The version of the structure that is dumped is not necessarily identical to the version used by an application. The structure can be modified by an API crossing exit, by the activity trace code, or by the queue manager. A queue manager can modify a structure to a later version, but the queue manager never changes it to an earlier version of the structure. To do so, would risk losing data.

AXREG

Application has started the AXREG AX function

XID

Description:	The XID structure
PCF Parameter:	MQBACF_XA_XID
Trace level:	1
Type:	MQCFBS
Length:	Sizeof(XID)

Rmid

Description: Resource manager identifier
PCF Parameter: MQIACF_XA_RMID
Trace level: 1
Type: MQCFIN

Flags

Description: Flags
PCF Parameter: MQIACF_XA_FLAGS
Trace level: 1
Type: MQCFIN

XARetCode

Description: Return code
PCF Parameter: MQIACF_XA_RETCODE
Trace level: 1
Type: MQCFIN

AXUNREG

Application has started the AXUNREG AX function

Rmid

Description: Resource manager identifier
PCF Parameter: MQIACF_XA_RMID
Trace level: 1
Type: MQCFIN

Flags

Description: Flags
PCF Parameter: MQIACF_XA_FLAGS
Trace level: 1
Type: MQCFIN

XARetCode

Description: Return code
PCF Parameter: MQIACF_XA_RETCODE
Trace level: 1
Type: MQCFIN

XACLOSE

Application has started the XACLOSE AX function

Xa_info

Description: Information used to initialize the resource manager.

PCF Parameter: MQCACF_XA_INFO
Trace level: 1
Type: MQCFST

Rmid

Description: Resource manager identifier
PCF Parameter: MQIACF_XA_RMID
Trace level: 1
Type: MQCFIN

Flags

Description: Flags
PCF Parameter: MQIACF_XA_FLAGS
Trace level: 1
Type: MQCFIN

XARetCode

Description: Return code
PCF Parameter: MQIACF_XA_RETCODE
Trace level: 1
Type: MQCFIN

XACOMMIT

Application has started the XACOMMIT AX function

XID

Description: The XID structure
PCF Parameter: MQBACF_XA_XID
Trace level: 1
Type: MQCFBS
Length: Sizeof(XID)

Rmid

Description: Resource manager identifier
PCF Parameter: MQIACF_XA_RMID
Trace level: 1
Type: MQCFIN

Flags

Description: Flags
PCF Parameter: MQIACF_XA_FLAGS
Trace level: 1
Type: MQCFIN

XARetCode

Description:	Return code
PCF Parameter:	MQIACF_XA_RETCODE
Trace level:	1
Type:	MQCFIN

XACOMPLETE

Application has started the XACOMPLETE AX function

Handle

Description:	Handle to async operation
PCF Parameter:	MQIACF_XA_HANDLE
Trace level:	1
Type:	MQCFIN

Retval

Description:	Return value of the asynchronous function
PCF Parameter:	MQIACF_XA_RETVAL
Trace level:	1
Type:	MQCFINMQCFBS

Rmid

Description:	Resource manager identifier
PCF Parameter:	MQIACF_XA_RMID
Trace level:	1
Type:	MQCFIN

Flags

Description:	Flags
PCF Parameter:	MQIACF_XA_FLAGS
Trace level:	1
Type:	MQCFIN

XARetCode

Description:	Return code
PCF Parameter:	MQIACF_XA_RETCODE
Trace level:	1
Type:	MQCFIN

XAEND

Application has started the XAEND AX function

XID

Description:	The XID structure
--------------	-------------------

PCF Parameter:	MQBACF_XA_XID
Trace level:	1
Type:	MQCFBS
Length:	Sizeof(XID)

Rmid

Description:	Resource manager identifier
PCF Parameter:	MQIACF_XA_RMID
Trace level:	1
Type:	MQCFIN

Flags

Description:	Flags
PCF Parameter:	MQIACF_XA_FLAGS
Trace level:	1
Type:	MQCFIN

XARetCode

Description:	Return code
PCF Parameter:	MQIACF_XA_RETCODE
Trace level:	1
Type:	MQCFIN

XAFORGET

Application has started the AXREG AX function

XID

Description:	The XID structure
PCF Parameter:	MQBACF_XA_XID
Trace level:	1
Type:	MQCFBS
Length:	Sizeof(XID)

Rmid

Description:	Resource manager identifier
PCF Parameter:	MQIACF_XA_RMID
Trace level:	1
Type:	MQCFIN

Flags

Description:	Flags
PCF Parameter:	MQIACF_XA_FLAGS
Trace level:	1

Type: MQCFIN

XARetCode

Description: Return code
PCF Parameter: MQIACF_XA_RETCODE
Trace level: 1
Type: MQCFIN

XAOPEN

Application has started the XAOPEN AX function

Xa_info

Description: Information used to initialize the resource manager.
PCF Parameter: MQCACF_XA_INFO
Trace level: 1
Type: MQCFST

Rmid

Description: Resource manager identifier
PCF Parameter: MQIACF_XA_RMID
Trace level: 1
Type: MQCFIN

Flags

Description: Flags
PCF Parameter: MQIACF_XA_FLAGS
Trace level: 1
Type: MQCFIN

XARetCode

Description: Return code
PCF Parameter: MQIACF_XA_RETCODE
Trace level: 1
Type: MQCFIN

XAPREPARE

Application has started the XAPREPARE AX function

XID

Description: The XID structure
PCF Parameter: MQBACF_XA_XID
Trace level: 1
Type: MQCFBS
Length: Sizeof(XID)

Rmid

Description:	Resource manager identifier
PCF Parameter:	MQIACF_XA_RMID
Trace level:	1
Type:	MQCFIN

Flags

Description:	Flags
PCF Parameter:	MQIACF_XA_FLAGS
Trace level:	1
Type:	MQCFIN

XARetCode

Description:	Return code
PCF Parameter:	MQIACF_XA_RETCODE
Trace level:	1
Type:	MQCFIN

XARECOVER

Application has started the XARECOVER AX function

Count

Description:	Count of XIDs
PCF Parameter:	MQIACF_XA_COUNT
Trace level:	1
Type:	MQCFIN

XIDs

Description:	The XID structures
	Note: There are multiple instances of this PCF parameter - one for every XID structure up to Count XIDs
PCF Parameter:	MQBACF_XA_XID
Trace level:	1
Type:	MQCFBS
Length:	Sizeof(XID)

Rmid

Description:	Resource manager identifier
PCF Parameter:	MQIACF_XA_RMID
Trace level:	1
Type:	MQCFIN

Flags

Description:	Flags
PCF Parameter:	MQIACF_XA_FLAGS
Trace level:	1
Type:	MQCFIN

XARetCode

Description:	Return code
PCF Parameter:	MQIACF_XA_RETCODE
Trace level:	1
Type:	MQCFIN

XAROLLBACK

Application has started the XAROLLBACK AX function

XID

Description:	The XID structure
PCF Parameter:	MQBACF_XA_XID
Trace level:	1
Type:	MQCFBS
Length:	Sizeof(XID)

Rmid

Description:	Resource manager identifier
PCF Parameter:	MQIACF_XA_RMID
Trace level:	1
Type:	MQCFIN

Flags

Description:	Flags
PCF Parameter:	MQIACF_XA_FLAGS
Trace level:	1
Type:	MQCFIN

XARetCode

Description:	Return code
PCF Parameter:	MQIACF_XA_RETCODE
Trace level:	1
Type:	MQCFIN

XASTART

Application has started the XASTART AX function

XID

Description:	The XID structure
PCF Parameter:	MQBACF_XA_XID
Trace level:	1
Type:	MQCFBS
Length:	Sizeof(XID)

Rmid

Description:	Resource manager identifier
PCF Parameter:	MQIACF_XA_RMID
Trace level:	1
Type:	MQCFIN

Flags

Description:	Flags
PCF Parameter:	MQIACF_XA_FLAGS
Trace level:	1
Type:	MQCFIN

XARetCode

Description:	Return code
PCF Parameter:	MQIACF_XA_RETCODE
Trace level:	1
Type:	MQCFIN

Multi

System topics for monitoring and activity trace

System topics in queue manager topic trees are used for resource monitoring (some of which is similar to the content of statistics messages) and as a way of consuming application activity trace.

\$SYS/MQ branch of queue manager topic tree

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 application activity trace and monitoring. For more information, see [Topic Trees](#).

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.

- You 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.
- 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.

Performance monitoring API available through publish/subscribe

You can monitor queue manager statistics by using the simple publish/subscribe mechanism and either the amqsrua sample application, or your own application. Statistics are published to a system topic under \$SYS/MQ/INFO/QMGR that help users monitor resources. These statistics can be viewed either by running the amqsrua sample application, or by writing an application that subscribes to the resource monitoring system topic in a similar way to **amqsrua**.

Related concepts

[“Application activity trace” on page 223](#)

Application activity trace produces detailed information about the behavior of applications connected to a queue manager. It traces the behavior of an application and provides a detailed view of the parameters used by an application as it interacts with IBM MQ resources. It also shows the sequence of MQI calls issued by an application.

Monitoring system resource usage by using the amqsrua command

You can use the **amqsrua** command to query performance data that is related to the system resource usage of a queue manager.

About this task

The amqsrua sample application showcases a way to consume IBM MQ monitoring publications and display performance data 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.

You can run the command with just the queue manager name and interactively step through to choose the **CLASS**, then **TYPE** and then **object** parameters that are available for the queue manager at each step. If you know the **CLASS**, **TYPE** and **object** names for which you want to see information, you can specify them when you run the **amqsrua** command.

By default, the amqsrua application looks for statistics that are published by the queue manager under the topic tree \$SYS/MQ/INFO/QMGR. Other components or applications can use a similar mechanism to publish under a different topic starting point. From IBM MQ 9.1.0, you can use the **-p** parameter to specify where amqsrua looks for the statistics for these other components on both Linux and Windows.

-m

Queue manager name. The queue manager must be running. If you do not specify a queue manager name, statistics for the default queue manager are displayed.

-c

CLASS name. IBM MQ resource usage publications are associated with a class. The classes represent the top level in the tree of metadata that describes the available resource usage information.

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.

STATAPP

Returns information about usage statistics for the application specified.

See [“Developing your own resource monitoring application” on page 312](#) for information on how you specify the topic tree for metadata, and [Using the application name in supported programming languages](#) for the characters you can use in the names of applications.

Note: You use the **-o** option to specify the application name being monitored for STATAPP, and the queue name for STATQ.

-t

TYPE name. IBM MQ resource usage publications are associated with a type within a class. Each publication includes the class and type that allows the class/type/element definitions to be found and the resulting publications to be handled. The class/type/element descriptions are published as metadata at queue manager startup.

-o

Object name. Resource usage publications are PCF messages that consist of a sequence of PCF elements. The PCF elements that are published for each class/type pair are advertised in the metadata. A leaf is stored in the tree that describes each element, thus allowing the elements to be processed.

-p

Metadata prefix. Specify a topic tree starting point where **amqsrua** can look for statistics that are published by queue managers. The default topic tree is \$SYS/MQ/INFO/QMGR but other components or applications might publish statistics under a different topic tree starting point.

-n

Publication count. You can specify how many reports are returned before the command ends. The data is published approximately 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.

-s

Model queue. Optionally specify a model queue to use. (By default **amqsrua** uses the SYSTEM.DEFAULT.MODEL.QUEUE queue.)

-h

Usage

Procedure

1. From the samples directory, issue the following command to display the available data for the queue manager:

-  On Linux, `MQ_INSTALLATION_PATH/samp/bin:`

```
./amqsrua -m QMgrName
```

-  On Windows, `MQ_INSTALLATION_PATH\tools\c\Samples\Bin64:`

```
amqsrua -m QMgrName
```

where *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.

The following options are available:

```
CPU : Platform central processing units
DISK : Platform persistent data stores
STATMQI : API usage statistics
STATQ : API per-queue usage statistics
Enter Class selection
==>
```

2. From the list of CLASS options, enter STATMQI.

```
==> STATMQI
CONNDISC : MQCONN and MQDISC
OPENCLOSE : MQOPEN and MQCLOSE
INQSET : MQINQ and MQSET
PUT : MQPUT
GET : MQGET
SYNCPPOINT : Commit and rollback
SUBSCRIBE : Subscribe
PUBLISH : Publish
Enter Type selection
==>
```

3. From the list of TYPE options, enter PUT.

```
==>PUT
Publication received PutDate:20170329 PutTime:17045485 Interval:4 minutes,13.978 seconds
Interval total MQPUT/MQPUT1 count 22
Interval total MQPUT/MQPUT1 byte count 25284 100/sec
Non-persistent message MQPUT count 22
Persistent message MQPUT count 0
Failed MQPUT count 0
Non-persistent message MQPUT1 count 0
Persistent message MQPUT1 count 0
Failed MQPUT1 count 0
Put non-persistent messages - byte count 25284 100/sec
Put persistent messages - byte count 0
MQSTAT count 0

Publication received PutDate:20170329 PutTime:17050485 Interval:10.001 seconds
Interval total MQPUT/MQPUT1 count 1
Interval total MQPUT/MQPUT1 byte count 524 52/sec
Non-persistent message MQPUT count 1
Persistent message MQPUT count 0
Failed MQPUT count 0
Non-persistent message MQPUT1 count 0
Persistent message MQPUT1 count 0
Failed MQPUT1 count 0
Put non-persistent messages - byte count 524 52/sec
Put persistent messages - byte count 0
MQSTAT count 0
```

Results

You used the **amqsrua** sample application interactively to look at statistics that queue managers publish on the system topics under the metadata prefix `$SYS/MQ/INFO/QMGR`.

Note: Classes and types of resource publications that are available for queue managers might be different depending on their configuration, version, and platform. Use **amqsrua** interactively to find the classes, types, and elements that are available for your specific queue manager.

What to do next

To develop your own monitoring application, see [“Developing your own resource monitoring application”](#) on page 312.

Related tasks

[Preparing and running sample programs on Windows](#)

[Preparing and running sample programs on UNIX and Linux](#)

Developing your own resource monitoring application

You can develop your own application 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 metadata to discover what resource usage information is available, and on what topics, and then subscribe to the advertised topics.

The topic tree for the metadata has the following structure:

```
$SYS/MQ/INFO/QMGR/QMGR-NAME/Monitor/class[/instance]/type]
```

Use the following structure for the metadata of the topic tree:

```
$SYS/MQ/INFO/QMGR/QMGR-NAME/Monitor/class[/resourceid]/type]
```

to clarify that you are specifying the resource being monitored, which has the maximum of 28 character application name, and to avoid confusion with internal metadata type attributes.

For a list of possible classes, see [“Monitoring system resource usage by using the amqsrua command”](#) on page 309.

The source code for the amqsrua application is provided as an IBM MQ sample. You can use this sample application as a guide for creating your own monitoring application. 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 AIX and Linux platforms, `MQ_INSTALLATION_PATH/samp/`
-  On Windows platforms, `MQ_INSTALLATION_PATH\tools\c\Samples\`

The amqsrua application subscribes to IBM MQ resource usage topics and formats the resulting published PCF data. The application source provides a basic example of how to subscribe to and consume this type of administrative data. The amqsrua application 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.

Multi Metrics published on the system topics

Metrics are categorized into classes, and sub-categorized into types. There are various metrics published under each metric class and type.

Index

- [“CPU \(platform central processing units\)” on page 313](#)
- [“DISK \(platform persistent data stores\)” on page 314](#)
- [“STATMQI \(API usage statistics\)” on page 315](#)
- [“STATQ \(API per-queue usage statistics\)” on page 317](#)
- [“STATAPP \(per-application usage statistics\)” on page 319](#)
- [“NHAREPLICA \(per-instance Native HA statistics\)” on page 320](#)

  See [“Monitoring system resource usage by using the amqsrua command” on page 309](#) for information on how you collect data for the options listed, with the exception of NHAREPLICA.

You can use also the [ALTER QMGR](#) command to monitor STATMQI and STATQ at the queue manager level, or the local queue attribute STATQ to monitor individual queues; see [ALTER QUEUES](#) for this option.

CPU (platform central processing units)

Introduction

Where statistics refer to the *current interval*, this is the interval defined in the published message by the MQIAMO64_MONITOR_INTERVAL parameter.

Statistics are usually published every 10 seconds, the *published interval*, as long as there is at least one active subscriber, but the precise interval should always be taken from the message.

Important: Unless otherwise specified, metrics are otherwise absolute values for the point in time when captured.

SystemSummary (CPU performance - platform wide)

User CPU time percentage *X%*

The average percentage of time (taken over the last 10 second interval) used by the CPU when it was in non-privileged code.

System CPU time percentage *X%*

The average percentage of time (taken over the last 10 second interval) used by the CPU when it was in privileged code.

CPU load - one minute average *X*

The one minute load average. "Load average" is a industry-wide term, but the exact value reported might differ across platforms.

CPU load - five minute average *X*

The five minute load average. "Load average" is a industry-wide term, but the exact value reported might differ across platforms.

CPU load - fifteen minute average *X*

The fifteen minute load average. "Load average" is a industry-wide term, but the exact value reported might differ across platforms.

CPU system summary

RAM free percentage *X%*

RAM total bytes XMB

RAM free percentage X%

RAM total bytes XMB

QMGrSummary (CPU performance - running queue manager)

User CPU time - percentage estimate for queue manager X%

The average percentage of time (taken over the last 10 second interval) used by the CPU when this queue manager's processes were in non-privileged code.

System CPU time - percentage estimate for queue manager X%

The average percentage of time (taken over the last 10 second interval) used by the CPU when this queue manager's processes were in privileged code.

RAM total bytes - estimate for queue manager XMB

This is an approximation of the memory used by the queue manager.

DISK (platform persistent data stores)

The *SystemSummary* and *QMGrSummary* are absolute values at the point of time of capture. See the [Introduction](#) for details of the *published interval*.

SystemSummary (disk usage - platform wide)

MQ errors file system - bytes in use XMB

MQ errors file system - free space X%

MQ FDC file count X

MQ trace file system - bytes in use XMB

MQ trace file system - free space X%

QMGrSummary (disk usage - running queue managers)

Queue Manager file system - bytes in use XMB

Queue Manager file system - free space X%

Log (disk usage - queue manager recovery log)

Log - bytes in use X

Log - bytes max X

The maximum number of bytes that can be written to the log if all the primary and secondary extents were full. This is less than the size of the log filesystem.

Log file system - bytes in use X

Log file system - bytes max X

 Log file system – free space X%

 Log – disk written log sequence number X

The LSN written and forced to disk as a 64-bit number.

Log - physical bytes written for the current interval X.

See the [Introduction](#) for the definition of *current interval*.

Log - logical bytes written for the *current interval* X

Log - write latency X uSec

A rolling average that represents the time that a single write to disk takes.

Where **LogWriteIntegrity**=*TripleWrite*, the physical number of bytes written to disk is greater than the logical bytes written.

Log - write size X, also rolling average.

Log - occupied by extents waiting to be archived X.

Only published when **logtype**= linear and **LogManagement** = archive. See [Log stanza of the gm.ini file](#) for more information.

Log - space in MB required for media recovery X.

Only published when **logtype**= linear.

Log - space in MB occupied by reusable extents X

Only published when **logtype**= linear and **LogManagement** = automatic. See [Log stanza of the gm.ini file](#) for more information.

Log - current primary space in use X%.

Log file space in use as a percentage of primary logs. This value can be more than 100%.

Log - workload primary space utilization X%.

The percentage log file space in use as a rolling average over recent history.

  Log – quorum log sequence number X

The LSN that has been replicated between a quorum of instances in the HA group as a 64-bit number (only returned if the queue manager is configured for Native HA).

  Log - slowest write since restart

The highest latency individual log write since the queue manager was started (in microseconds).

  Log - timestamp of slowest write

The UTC timestamp when the highest latency individual log write occurred (expressed as microseconds since the epoch - 1970-01-01T00:00:00Z).

STATMQI (API usage statistics)

All API usage statistics reflect occurrences and/or percentages for the published interval. See the [Introduction](#) for the definition of *published interval*.

The statistics outputs a count of the number of failed MQI calls, but not every failed MQI call appears in these statistics - indeed the failures of not every MQI call have their statistics recorded. This is because many reasons that MQI calls fail are diagnosed before the MQI call reached the internals of the queue manager where the statistics are recorded.

An example of this is MQRC_HCONN_ERROR returned to a client application. If a client application passes a bad **hconn**, the MQ client diagnoses that error and returns MQRC_HCONN_ERROR without passing the MQI call onto the queue manager. Hence, the failed MQI call never appears in the statistics recorded by the queue manager.

Statistics of failed MQI calls are useful because they enable customers to troubleshoot poorly-written applications that generate unnecessary failed MQI calls, thereby impacting performance. Some examples of failing reasons for various MQI calls that are recorded in the statistics:

- MQCONN/MQCONN/MQOPEN returns 2035 MQRC_NOT_AUTHORIZED when diagnosed by the queue manager, not the client. For example running **amqspu**t as nobody.
- MQPUT/MQPUT1 returns 2053 MQRC_Q_FULL because MAXDEPTH has been exceeded.
- MQGET returns 2033 MQRC_NO_MSG_AVAILABLE when browsing or destructively getting from an empty queue.
- MQSUBRQ returns 2437 MQRC_NO_RETAINED_MSG because there is no retained message.

CONNDISC (MQCONN and MQDISC)

MQCONN/MQCONN count X

Failed MQCONN/MQCONN count X

Concurrent connections - high water mark X

MQDISC count *X*

OPENCLOSE (MQOPEN and MQCLOSE)

MQOPEN count *X Y/sec*

Failed MQOPEN count *X*

MQCLOSE count *X Y/sec*

Failed MQCLOSE count *X*

INQSET (MQINQ and MQSET)

MQINQ count *X*

Failed MQINQ count *X*

MQSET count *X*

Failed MQSET count *X*

PUT (MQPUT)

Interval total MQPUT/MQPUT1 count *X*

Interval total MQPUT/MQPUT1 byte count *X Y/sec*

Non-persistent message MQPUT count *X*

Persistent message MQPUT count *X*

Failed MQPUT count *X*

Non-persistent message MQPUT1 count *X*

Persistent message MQPUT1 count *X*

Failed MQPUT1 count *X*

Put non-persistent messages - byte count *X Y/sec*

Put persistent messages - byte count *X*

MQSTAT count *X*

GET (MQGET)

Interval total destructive get- count *X*

Interval total destructive get - byte count *X Y/sec*

Non-persistent message destructive get - count *X*

Persistent message destructive get - count *X*

Failed MQGET - count *X*

Got non-persistent messages - byte count *X Y/sec*

Got persistent messages - byte count *X*

Non-persistent message browse - count *X*

Persistent message browse - count *X*

Failed browse count *X*

Non-persistent message browse - byte count *X Y/sec*

Persistent message browse - byte count *X*

Expired message count *X*

Purged queue count *X*

MQCB count *X*

Failed MQCB count *X*

MQCTL count *X*

SYNCPOINT (commit and rollback)

Commit count *X*

Rollback count *X*

SUBSCRIBE (subscribe)

Create durable subscription count *X*
Alter durable subscription count *X*
Resume durable subscription count *X*
Create non-durable subscription count *X*
Failed create/alter/resume subscription count *X*
Delete durable subscription count *X*
Delete non-durable subscription count *X*
Subscription delete failure count *X*
MQSUBRQ count *X*
Failed MQSUBRQ count *X*
Durable subscriber - high water mark *X*
Durable subscriber - low water mark *X*
Non-durable subscriber - high water mark *X*
Non-durable subscriber - low water mark *X*

PUBLISH (publish)

Topic MQPUT/MQPUT1 interval total *X*
Interval total topic bytes put *X* *Y*/sec
Published to subscribers - message count *X*
Published to subscribers - byte count *X*
Non-persistent - topic MQPUT/MQPUT1 count *X*
Persistent - topic MQPUT/MQPUT1 count *X*
Failed topic MQPUT/MQPUT1 count *X*

STATQ (API per-queue usage statistics)

GENERAL (General)

messages expired *X* (moved from `GET` for IBM MQ 9.3.0 and later CD versions)
queue purged count *X* (moved from `GET` for IBM MQ 9.3.0 and later CD versions)
average queue time *X* uSec (moved from `GET` for IBM MQ 9.3.0 and later CD versions)
Queue depth *X* (moved from `GET` for IBM MQ 9.3.0 and later CD versions)

 **V 9.4.2**  **V 9.4.0.10** open input count

Number of queue handles open at the end of the interval for input (MQGET). This value is the same as the IPPROCS value reported by DISPLAY QLOCAL / QSTATUS.

 **V 9.4.2**  **V 9.4.0.10** open output count

Number of queue handles at the end of the interval that are open for output (MQPUT). This value is the same as the OPPROCS value reported by DISPLAY QLOCAL / QSTATUS.

 **V 9.4.2**  **V 9.4.0.10** open browse count

Number of queue handles open at the end of the interval that included the MQOO_BROWSE option on their call to MQOPEN. Note: these handles are also included in the open input count.

 **V 9.4.2**  **V 9.4.0.10** open publish count

Number of queue handles open at the end of the interval that were opened by queue manager processes to put messages to subscriptions that specified this queue as their destination. Note: these handles are also included in the open output count.

OPENCLOSE (MQOPEN and MQCLOSE)MQOPEN count *X*MQCLOSE count *X***INQSET (MQINQ and MQSET)**MQINQ count *X*MQSET count *X***PUT (MQPUT and MQPUT1)**MQPUT/MQPUT1 count *X*MQPUT byte count *X*MQPUT non-persistent message count *X*MQPUT persistent message count *X*rolled back MQPUT count *X*MQPUT1 non-persistent message count *X*MQPUT1 persistent message count *X*non-persistent byte count *X*persistent byte count *X*lock contention *X*%

The percentage of attempts to lock the queue that resulted in waiting for another process to release the lock first. Decreasing lock contention is likely to increase the maximum throughput of your system because taking a lock that is not currently locked is a more efficient than waiting for a lock to be released.

queue avoided puts *X*%

If a message is put to a queue when there is a waiting getter, the message might not need to be queued as it can be possible for it to be passed to the getter immediately. So this message is said to have avoided the queue, and "queue avoided puts" is the count of such messages. Increasing queue avoidance is likely to increase the maximum throughput of your system because it avoids the cost of putting the message onto the queue and getting it off again.

queue avoided bytes *X*%

If a message is put to a queue when there is a waiting getter, the message might not need to be queued as it can be possible for it to be passed to the getter immediately. So this message is said to have avoided the queue, and "queue avoided bytes" is the count of such bytes. Increasing queue avoidance is likely to increase the maximum throughput of your system because it avoids the cost of putting the message onto the queue and getting it off again.

GET (MQGET)MQGET count *X*MQGET byte count *X*destructive MQGET non-persistent message count *X*destructive MQGET persistent message count *X*rolled back MQGET count *X*destructive MQGET non-persistent byte count *X*destructive MQGET persistent byte count *X*MQGET browse non-persistent message count *X*MQGET browse persistent message count *X*MQGET browse non-persistent byte count *X*MQGET browse persistent byte count *X*messages expired *X* (moved to GENERAL from IBM MQ 9.3)

queue purged count *X* (moved to GENERAL from IBM MQ 9.3)
average queue time *X* uSec (moved to GENERAL from IBM MQ 9.3)
Queue depth *X* (moved to GENERAL from IBM MQ 9.3)
destructive MQGET fails *X*
destructive MQGET fails with MQRC_NO_MSG_AVAILABLE *X*
destructive MQGET fails with MQRC_TRUNCATED_MSG_FAILED *X*
MQGET browse fails *X*
MQGET browse fails with MQRC_NO_MSG_AVAILABLE *X*
MQGET browse fails with MQRC_TRUNCATED_MSG_FAILED *X*

V 9.4.2 **V 9.4.0.10** **EXTENDED**

msg search count

Number of MQGETs where the queue manager searched for a message (this will be every MQGET that was not satisfied by directly passing an MQPUT to a waiting getter - see "queue avoided puts").

msg not found count

Number of MQGETs where the queue manager failed to find a message.

msg examine count

Number of (matching and unmatching) messages examined by searches. Reasons for match failures are described in the following statistics:

intran get skipped count

Messages examined but skipped because they were locked by an uncommitted MQGET transaction.

put skipped count

Messages examined but skipped because they had been put in a transaction that had not been committed.

selection mismatch count

Messages that were checked and did not match properties required by selector.

correlid mismatch short count

Messages examined because of MQMO_MATCH_CORREL_ID and skipped because quick CorrelId hash did not match requested Id.

correlid mismatch long count

Messages examined and matching quick CorrelId hash check, but failing full CorrelId comparison.

msgid mismatch count

Messages examined because of MQMO_MATCH_MSG_ID and skipped because MsgId did not match requested Id.

load msg dtl count

Messages or message headers that needed to be loaded from the Q file to check for a match.

STATAPP (per-application usage statistics)

INSTANCE (instance statistics)

Instance count *X* absolute

Movable instance count *X* absolute

Instance shortfall count *X* absolute

Instances started *X* interval

Initiated outbound instance moves X interval
Completed outbound instance moves X interval
Instances ended during reconnect X interval
Instances ended X interval

NHAREPLICA (per-instance Native HA statistics)

REPLICATION (replication statistics)

Average network round-trip time X *uSec*
Synchronous log bytes sent X
Catch-up log bytes sent X
Synchronous compressed log bytes sent X
Catch-up compressed log bytes sent X
Synchronous uncompressed log bytes sent X
Catch-up uncompressed log bytes sent X
Synchronous log data average compression time X *uSec*
Catch-up log data average compression time X *uSec*
Synchronous log bytes decompressed
Catch-up log bytes decompressed
Synchronous log data average decompression time X *uSec*
Catch-up log data average decompression time X *uSec*
Log write average acknowledgment latency X *uSec*
Log write average acknowledgment size X
Backlog bytes X
Backlog average bytes X
 Backlog long-term average bytes X
The long-term rolling average number of bytes the replica is behind.
 Acknowledged log sequence number X
The LSN the instance has acknowledged as a 64-bit number.
 Log file system - bytes in use X
The number of bytes in use by the log file system on the instance.
 Log file system – free space $X\%$
The amount of free space.
 Queue Manager file system - bytes in use X *MB*
The number of MB in use by the queue manager file system.
 Queue Manager file system - free space $X\%$
The amount of free space.
 MQ FDC file count X
The number of FDCs present on the instance.
 Catch-up time percentage $X\%$
The percentage of time in this interval the replica spent catching-up.
 Throttling time percentage $X\%$

The percentage of time in this interval that the replica caused throttling of queue manager workload in an attempt to allow it to catch-up.

V 10.0.0 Log - write latency X uSec

A rolling average that represents the time that a single write to disk takes.

V 10.0.0 Log - write size X , also rolling average

V 10.0.0 Log - slowest write since restart

The highest latency individual log write since the queue manager was started (in microseconds).

V 10.0.0 Log - timestamp of slowest write

The UTC timestamp when the highest latency individual log write occurred (expressed as microseconds since the epoch - 1970-01-01T00:00:00Z).

V 10.0.0 **RECOVERY (recovery group statistics)**

Average network round-trip time X uSec

Compressed log bytes sent

Log data average compression time X uSec

Log bytes decompressed

Log data average decompression time X uSec

Log bytes sent X

The number of bytes sent to the group in this interval.

Backlog bytes X

The number of bytes the group is behind.

Backlog average bytes X

The short-term rolling average number of bytes the group is behind.

Rebase count X

The number of times the group has been rebased.

Recovery LSN X

The LSN the group could recover from as a 64-bit number.

Related information

[Metrics published when using the IBM MQ Operator](#)

Real-time monitoring

Real-time monitoring is a technique that allows you to determine the current state of queues and channels within a queue manager. The information returned is accurate at the moment the command was issued.

A number of commands are available that when issued return real-time information about queues and channels. Information can be returned for one or more queues or channels and can vary in quantity. Real-time monitoring can be used in the following tasks:

- Helping system administrators understand the steady state of their IBM MQ system. This helps with problem diagnosis if a problem occurs in the system.
- Determining the condition of your queue manager at any moment, even if no specific event or problem has been detected.
- Assisting with determining the cause of a problem in your system.

With real-time monitoring, information can be returned for either queues or channels. The amount of real-time information returned is controlled by queue manager, queue, and channel attributes.

- You monitor a queue by issuing commands to ensure that the queue is being serviced properly. Before you can use some of the queue attributes, you must enable them for real-time monitoring.
- You monitor a channel by issuing commands to ensure that the channel is running properly. Before you can use some of the channel attributes, you must enable them for real-time monitoring.

Real-time monitoring for queues and channels is in addition to, and separate from, performance and channel event monitoring.

Attributes that control real-time monitoring

Some queue and channel status attributes hold monitoring information, if real-time monitoring is enabled. If real-time monitoring is not enabled, no monitoring information is held in these monitoring attributes. Examples demonstrate how you can use these queue and channel status attributes.

You can enable or disable real-time monitoring for individual queues or channels, or for multiple queues or channels. To control individual queues or channels, set the queue attribute MONQ or the channel attribute MONCHL, to enable or disable real-time monitoring. To control many queues or channels together, enable or disable real-time monitoring at the queue manager level by using the queue manager attributes MONQ and MONCHL. For all queue and channel objects with a monitoring attribute that is specified with the default value, QMGR, real-time monitoring is controlled at the queue manager level.

Automatically defined cluster-sender channels are not IBM MQ objects, so do not have attributes in the same way as channel objects. To control automatically defined cluster-sender channels, use the queue manager attribute, MONACLS. This attribute determines whether automatically defined cluster-sender channels within a queue manager are enabled or disabled for channel monitoring.

For real-time monitoring of channels, you can set the MONCHL attribute to one of the three monitoring levels: low, medium, or high. You can set the monitoring level either at the object level or at the queue manager level. The choice of level is dependent on your system. Collecting monitoring data might require some instructions that are relatively expensive computationally, such as obtaining system time. To reduce the effect of real-time monitoring, the medium and low monitoring options measure a sample of the data at regular intervals rather than collecting data all the time. [Table 34 on page 322](#) summarizes the monitoring levels available for real-time monitoring of channels:

<i>Table 34. Monitoring levels</i>		
Level	Description	Usage
Low	Measure a small sample of the data, at regular intervals.	For objects that process a high volume of messages.
Medium	Measure a sample of the data, at regular intervals.	For most objects.
High	Measure all data, at regular intervals.	For objects that process only a few messages per second, on which the most current information is important.

For real-time monitoring of queues, you can set the MONQ attribute to one of the three monitoring levels, low, medium or high. However, there is no distinction between these values. The values all enable data collection, but do not affect the size of the sample.

Examples

The following examples demonstrate how to set the necessary queue, channel, and queue manager attributes to control the level of monitoring. For all of the examples, when monitoring is enabled, queue and channel objects have a medium level of monitoring.

1. To enable both queue and channel monitoring for all queues and channels at the queue manager level, use the following commands:

```
ALTER QMGR MONQ(MEDIUM) MONCHL(MEDIUM)
ALTER QL(Q1) MONQ(QMGR)
ALTER CHL(QM1.TO.QM2) CHLTYPE(SDR) MONCHL(QMGR)
```

2. To enable monitoring for all queues and channels, with the exception of local queue, Q1, and sender channel, QM1 . TO . QM2, use the following commands:

```
ALTER QMGR MONQ(MEDIUM) MONCHL(MEDIUM)
ALTER QL(Q1) MONQ(OFF)
ALTER CHL(QM1.TO.QM2) CHLTYPE(SDR) MONCHL(OFF)
```

3. To disable both queue and channel monitoring for all queues and channels, with the exception of local queue, Q1, and sender channel, QM1 . TO . QM2, use the following commands:

```
ALTER QMGR MONQ(OFF) MONCHL(OFF)
ALTER QL(Q1) MONQ(MEDIUM)
ALTER CHL(QM1.TO.QM2) CHLTYPE(SDR) MONCHL(MEDIUM)
```

4. To disable both queue and channel monitoring for all queues and channels, regardless of individual object attributes, use the following command:

```
ALTER QMGR MONQ(NONE) MONCHL(NONE)
```

5. To control the monitoring capabilities of automatically defined cluster-sender channels use the following command:

```
ALTER QMGR MONACLS(MEDIUM)
```

6. To specify that automatically defined cluster-sender channels are to use the queue manager setting for channel monitoring, use the following command:

```
ALTER QMGR MONACLS(QMGR)
```

Related concepts

[“Real-time monitoring” on page 321](#)

Real-time monitoring is a technique that allows you to determine the current state of queues and channels within a queue manager. The information returned is accurate at the moment the command was issued.

[“Using IBM MQ online monitoring” on page 342](#)

You can collect monitoring data for queues and channels (including automatically defined cluster-server channels) by setting the MONQ, MONCHL, and MONACLS attributes.

Related tasks

[“Displaying queue and channel monitoring data” on page 323](#)

To display real-time monitoring information for a queue or channel, use either the IBM MQ Explorer or the appropriate MQSC command. Some monitoring fields display a comma-separated pair of indicator values, which help you to monitor the operation of your queue manager. Examples demonstrate how you can display monitoring data.

[Displaying and altering queue manager attributes](#)

[Monitoring \(MONCHL\)](#)

Displaying queue and channel monitoring data

To display real-time monitoring information for a queue or channel, use either the IBM MQ Explorer or the appropriate MQSC command. Some monitoring fields display a comma-separated pair of indicator values,

which help you to monitor the operation of your queue manager. Examples demonstrate how you can display monitoring data.

About this task

Monitoring fields that display a pair of values separated by a comma provide short term and long term indicators for the time measured since monitoring was enabled for the object, or from when the queue manager was started:

- The short term indicator is the first value in the pair and is calculated in a way such that more recent measurements are given a higher weighting and will have a greater effect on this value. This gives an indication of recent trend in measurements taken.
- The long term indicator is the second value in the pair and is calculated in a way such that more recent measurements are not given such a high weighting. This gives an indication of the longer term activity on performance of a resource.

These indicator values are most useful to detect changes in the operation of your queue manager. This requires knowledge of the times these indicators show when in normal use, in order to detect increases in these times. By collecting and checking these values regularly you can detect fluctuations in the operation of your queue manager. This can indicate a change in performance.

Obtain real-time monitoring information as follows:

Procedure

1. To display real-time monitoring information for a queue, use either the IBM MQ Explorer or the MQSC command `DISPLAY QSTATUS`, specifying the optional parameter `MONITOR`.
2. To display real-time monitoring information for a channel, use either the IBM MQ Explorer or the MQSC command `DISPLAY CHSTATUS`, specifying the optional parameter `MONITOR`.

Example

The queue, Q1, has the attribute `MONQ` set to the default value, `QMGR`, and the queue manager that owns the queue has the attribute `MONQ` set to `MEDIUM`. To display the monitoring fields collected for this queue, use the following command:

```
DISPLAY QSTATUS(Q1) MONITOR
```

The monitoring fields and monitoring level of queue, Q1 are displayed as follows:

```
QSTATUS(Q1)
TYPE(Queue)
MONQ(MEDIUM)
QTIME(11892157,24052785)
MSGAGE(37)
LPUTDATE(2005-03-02)
LPUTTIME(09.52.13)
LGETDATE(2005-03-02)
LGETTIME(09.51.02)
```

The sender channel, QM1.TO.QM2, has the attribute `MONCHL` set to the default value, `QMGR`, and the queue manager that owns the queue has the attribute `MONCHL` set to `MEDIUM`. To display the monitoring fields collected for this sender channel, use the following command:

```
DISPLAY CHSTATUS(QM1.TO.QM2) MONITOR
```

The monitoring fields and monitoring level of sender channel, QM1.TO.QM2 are displayed as follows:

```
CHSTATUS(QM1.TO.QM2)
XMITQ(Q1)
CONNNAME(127.0.0.1)
CURRENT
CHLTYPE(SDR)
STATUS(RUNNING)
SUBSTATE(MQGET)
```

```
MONCHL(MEDIUM)
XQTIME(755394737,755199260)
NETTIME(13372,13372)
EXITTIME(0,0)
XBATCHSZ(50,50)
COMPTIME(0,0)
STOPREQ(NO)
RQMNAME(QM2)
```

Related concepts

[“Real-time monitoring” on page 321](#)

Real-time monitoring is a technique that allows you to determine the current state of queues and channels within a queue manager. The information returned is accurate at the moment the command was issued.

Related reference

[DISPLAYQSTATUS](#)

Monitoring queues

Use this page to view tasks that help you to resolve a problem with a queue and the application that services that queue. Various monitoring options are available to determine the problem

Frequently, the first sign of a problem with a queue that is being serviced is that the number of messages on the queue (CURDEPTH) increases. If you expect an increase at certain times of day or under certain workloads, an increasing number of messages might not indicate a problem. However, if you have no explanation for the increasing number of messages, you might want to investigate the cause.

You might have an application queue where there is a problem with the application, or a transmission queue where there is a problem with the channel. Additional monitoring options are available when the application that services the queue is a channel.

The following examples investigate problems with a particular queue, called Q1, and describe the fields that you look at in the output of various commands:

Determining whether your application has the queue open

If you have a problem with a queue, check whether your application has the queue open

About this task

Perform the following steps to determine whether your application has the queue open:

Procedure

1. Ensure that the application that is running against the queue is the application that you expect. Issue the following command for the queue in question:

```
DISPLAY QSTATUS(Q1) TYPE(HANDLE) ALL
```

In the output, look at the APPLTAG field, and check that the name of your application is shown. If the name of your application is not shown, or if there is no output at all, start your application.

2. If the queue is a transmission queue, look in the output at the CHANNEL field.

If the channel name is not shown in the CHANNEL field, determine whether the channel is running.

3. Ensure that the application that is running against the queue has the queue open for input. Issue the following command:

```
DISPLAY QSTATUS(Q1) TYPE(QUEUE) ALL
```

In the output, look at the IPPROCS field to see if any application has the queue open for input. If the value is 0 and this is a user application queue, make sure that the application opens the queue for input to get the messages off the queue.

Checking that messages on the queue are available

If you have a large number of messages on the queue and your application is not processing any of those messages, check whether the messages on the queue are available to your application

About this task

Perform the following steps to investigate why your application is not processing messages from the queue:

Procedure

1. Ensure that your application is not asking for a specific message ID or correlation ID when it should be processing all the messages on the queue.
2. Although the current depth of the queue might show that there is an increasing number of messages on the queue, some messages on the queue might not be available to be got by an application, because they are not committed; the current depth includes the number of uncommitted MQPUTs of messages to the queue. Issue the following command:

```
DISPLAY QSTATUS(Q1) TYPE(Queue) ALL
```

In the output, look at the UNCOM field to see whether there are any uncommitted messages on the queue.

3. If your application is attempting to get any messages from the queue, check whether the putting application is committing the messages correctly. Issue the following command to find out the names of applications that are putting messages to this queue:

```
DISPLAY QSTATUS(Q1) TYPE(HANDLE) OPENTYPE(OUTPUT)
```

4. Then issue the following command, inserting in *appltag* the APPLTAG value from the output of the previous command:

```
DISPLAY CONN(*) WHERE(APPLTAG EQ appltag) UOWSTDA UOWSTTI
```

This shows when the unit of work was started and will help you discover whether the application is creating a long running unit of work. If the putting application is a channel, you might want to investigate why a batch is taking a long time to complete.

Checking whether your application is getting messages off the queue

If you have a problem with a queue and the application that services that queue, check whether your application is getting messages off the queue

About this task

To check whether your application is getting messages off the queue, perform the following checks:

Procedure

1. Ensure that the application that is running against the queue is actually processing messages from the queue. Issue the following command:

```
DISPLAY QSTATUS(Q1) TYPE(Queue) ALL
```

In the output, look at the LGETDATE and LGETTIME fields which show when the last get was done from the queue.

2. If the last get from this queue was longer ago than expected, ensure that the application is processing messages correctly.

If the application is a channel, check whether messages are moving through that channel

Determining whether the application can process messages fast enough

If messages are building up on the queue, but your other checks have not found any processing problems, check that the application can process messages fast enough. If the application is a channel, check that the channel can process messages fast enough.

About this task

To determine whether the application is processing messages fast enough, perform the following tests:

Procedure

1. Issue the following command periodically to gather performance data about the queue:

```
DISPLAY QSTATUS(Q1) TYPE(QQUEUE) ALL
```

If the values in the QTIME indicators are high, or are increasing over the period, and you have already ruled out the possibility of long running Units of Work by checking that messages on the queue are available, the getting application might not be keeping up with the putting applications.

2. If your getting application cannot keep up with the putting applications, consider adding another getting application to process the queue.

Whether you can add another getting application depends on the design of the application and whether the queue can be shared by more than one application. Features such as message grouping or getting by correlation ID might help to ensure that two applications can process a queue simultaneously.

Checking the queue when the current depth is not increasing

Even if the current depth of your queue is not increasing, it might still be useful to monitor the queue to check whether your application is processing messages correctly.

About this task

To gather performance data about the queue: Issue the following command periodically:

Procedure

Issue the following command periodically:

```
DISPLAY QSTATUS(Q1) TYPE(QQUEUE) MSGAGE QTIME
```

In the output, if the value in MSGAGE increases over the period of time, and your application is designed to process all messages, this might indicate that some messages are not being processed at all.

Monitoring channels

Use this page to view tasks that help you to resolve a problem with a transmission queue and the channel that services that queue. Various channel monitoring options are available to determine the problem.

Frequently, the first sign of a problem with a queue that is being serviced is that the number of messages on the queue (CURDEPTH) increases. If you expect an increase at certain times of day or under certain workloads, an increasing number of messages might not indicate a problem. However, if you have no explanation for the increasing number of messages, you might want to investigate the cause.

You might have a problem with the channel that services a transmission queue. Various channel monitoring options are available to help you to determine the problem.

The following examples investigate problems with a transmission queue called QM2 and a channel called QM1.TO.QM2. This channel is used to send messages from queue manager, QM1, to queue manager, QM2. The channel definition at queue manager QM1 is either a sender or server channel, and the channel definition at queue manager, QM2, is either a receiver or requester channel.

Determining whether the channel is running

If you have a problem with a transmission queue, check whether the channel is running.

About this task

Perform the following steps to check the status of the channel that is servicing the transmission queue:

Procedure

1. Issue the following command to find out which channel you expect to process the transmission queue QM2:

```
DIS CHANNEL(*) WHERE(XMITQ EQ QM2)
```

In this example, the output of this command shows that the channel servicing the transmission queue is QM1.TO.QM2

2. Issue the following command to determine the status of the channel, QM1.TO.QM2:

```
DIS CHSTATUS(QM1.TO.QM2) ALL
```

3. Inspect the STATUS field of the output from the **CHSTATUS** command:

- If the value of the STATUS field is **RUNNING**, check that the channel is moving messages
- If the output from the command shows no status, or the value of the STATUS field is **STOPPED**, **RETRY**, **BINDING**, or **REQUESTING**, perform the appropriate step, as follows:

4. Optional: If the value of the STATUS field shows no status, the channel is inactive, so perform the following steps:

- a) If the channel should have been started automatically by a trigger, check that the messages on the transmission queue are available.

If there are messages available on the transmission queue, check that the trigger settings on the transmission queue are correct.

- b) Issue the following command to start the channel again manually:

```
START CHANNEL(QM1.TO.QM2)
```

5. Optional: If the value of the STATUS field is **STOPPED**, perform the following steps:

- a) Check the error logs to determine why the channel stopped. If the channel stopped owing to an error, correct the problem.

Ensure also that the channel has values specified for the retry attributes: *SHORTRTY* and *LONGRTY*. In the event of transient failures such as network errors, the channel will then attempt to restart automatically.

- b) Issue the following command to start the channel again manually:

```
START CHANNEL(QM1.TO.QM2)
```

 On IBM MQ for z/OS, you can detect when a user stops a channel by using command event messages.

6. Optional: If the value of the STATUS field is **RETRY**, perform the following steps:

- a) Check the error logs to identify the error, then correct the problem.

- b) Issue the following command to start the channel again manually:

```
START CHANNEL(QM1.TO.QM2)
```

or wait for the channel to connect successfully on its next retry.

7. Optional: If the value of the STATUS field is **BINDING** or **REQUESTING**, the channel has not yet successfully connected to the partner. Perform the following steps:

- a) Issue the following command, at both ends of the channel, to determine the substate of the channel:

```
DIS CHSTATUS(QM1.TO.QM2) ALL
```

Note:

- i) In some cases there might be a substate at one end of the channel only.
 - ii) Many substates are transitory, so issue the command a few times to detect whether a channel is stuck in a particular substate.
- b) Check [Table 35 on page 329](#) to determine what action to take:

Table 35. Substates seen with status binding or requesting		
Initiating MCA substate ¹	Responding MCA substate ²	Notes
NAMESERVER		The initiating MCA is waiting for a name server request to complete. Ensure that the correct host name has been specified in the channel attribute, CONNAME, and that your name servers are set up correctly.
SCYEXIT	SCYEXIT	The MCAs are currently <i>in conversation</i> through a security exit. For more information, see “Determining whether the channel can process messages fast enough” on page 332 .
	CHADEXIT	The channel autodefinition exit is currently executing. For more information, see “Determining whether the channel can process messages fast enough” on page 332 .
RCVEXIT SENDEXIT MSGEXIT MREXIT	RCVEXIT SENDEXIT MSGEXIT MREXIT	Exits are called at channel startup for MQXR_INIT. Review the processing in this part of your exit if this takes a long time. For more information, see “Determining whether the channel can process messages fast enough” on page 332 .
SERIALIZE	SERIALIZE	This substate only applies to channels with a disposition of SHARED.
NETCONNECT		This substate is shown if there is a delay in connecting due to incorrect network configuration.
SSLHANDSHAKE	SSLHANDSHAKE	A TLS handshake consists of a number of sends and receives. If network times are slow, or connection to lookup CRLs are slow, this affects the time taken to do the handshake.  On IBM MQ for z/OS this substate can also be indicative of not having enough SSLTASKS.

Notes:

- i) The initiating MCA is the end of the channel which started the conversation. This can be senders, cluster-senders, fully-qualified servers and requesters. In a server-requester pair, it is the end from which you started the channel.
- ii) The responding MCA is the end of the channel which responded to the request to start the conversation. This can be receivers, cluster-receivers, requesters (when the server or sender is started), servers (when the requester is started) and senders (in a requester-sender call-back pair of channels).

Checking that the channel is moving messages

If you have a problem with a transmission queue, check that the channel is moving messages

Before you begin

Issue the command `DIS CHSTATUS(QM1.TO.QM2) ALL`. If the value of the STATUS field is RUNNING, the channel has successfully connected to the partner system.

Check that there are no uncommitted messages on the transmission queue, as described in [“Checking that messages on the queue are available”](#) on page 326.

About this task

If there are messages available for the channel to get and send, perform the following checks:

Procedure

1. In the output from the display channel status command, `DIS CHSTATUS(QM1.TO.QM2) ALL`, look at the following fields:

MSGS

Number of messages sent or received (or, for server-connection channels, the number of MQI calls handled) during this session (since the channel was started).

BUFSENT

Number of transmission buffers sent. This includes transmissions to send control information only.

BYTSENT

Number of bytes sent during this session (since the channel was started). This includes control information sent by the message channel agent.

LSTMSGDA

Date when the last message was sent or MQI call was handled, see LSTMSGTI.

LSTMSGTI

Time when the last message was sent or MQI call was handled. For a sender or server, this is the time the last message (the last part of it if it was split) was sent. For a requester or receiver, it is the time the last message was put to its target queue. For a server-connection channel, it is the time when the last MQI call completed.

CURMSGGS

For a sending channel, this is the number of messages that have been sent in the current batch. For a receiving channel, it is the number of messages that have been received in the current batch. The value is reset to zero, for both sending and receiving channels, when the batch is committed.

STATUS

The status of the channel, which can be Starting, Binding, Initializing, Running, Stopping, Retrying, Paused, Stopped, or Requesting.

SUBSTATE

The action that the channel is currently performing.

INDOUBT

Whether the channel is currently in doubt. This is only YES while the sending Message Channel Agent is waiting for an acknowledgment that a batch of messages that it has sent has been successfully received. It is NO at all other times, including the period during which messages are being sent, but before an acknowledgment has been requested. For a receiving channel, the value is always NO.

2. Determine whether the channel has sent any messages since it started. If any have been sent, determine when the last message was sent.
3. The channel might have started a batch that has not yet completed, as indicated by a non-zero value in CURMSGGS. If INDOUBT is YES, the channel is waiting to receive acknowledgment that the other end of the channel received the batch. Look at the SUBSTATE field in the output and refer to [Table 36](#) on page 331:

Table 36. Sender and receiver MCA substates		
Sender SUBSTATE	Receiver SUBSTATE	Notes
MQGET	RECEIVE	Normal states of a channel at rest.
SEND	RECEIVE	SEND is usually a transitory state. If SEND is seen it indicates that the communication protocol buffers have filled. This can indicate a network problem.
RECEIVE		If the sender is seen in RECEIVE substate for any length of time, it is waiting on a response, either to a batch completion or a heartbeat. You might want to check why a batch takes a long time to complete.

Note: You might also want to determine whether the channel can process messages fast enough, especially if the channel has a substate associated with exit processing.

Checking why a batch takes a long time to complete

Reasons why a batch can take a long time to complete include a slow network or a channel is using message retry processing.

About this task

When a sender channel has sent a batch of messages it waits for confirmation of that batch from the receiver, unless the channel is pipelined. The factors described in this task can affect how long the sender channel waits. See also [“Determining whether the channel can process messages fast enough” on page 332](#) for help with this task.

Procedure

- Check whether the network is slow.

The NETTIME value is the amount of time, displayed in microseconds, taken to send an end of batch request to the remote end of the channel and receive a response minus the time to process the end of batch request. This value can be large for either of the following reasons:

- The network is slow. A slow network can affect the time it takes to complete a batch. The measurements that result in the indicators for the NETTIME field are measured at the end of a batch. However, the first batch affected by a slowdown in the network is not indicated with a change in the NETTIME value because it is measured at the end of the batch.
- Requests are queued at the remote end, for example a channel can be retrying a put, or a put request may be slow due to page set I/O. Once any queued requests have completed, the duration of the end of batch request is measured. So if you get a large NETTIME value, check for unusual processing at the remote end.

Note: You must enable channel monitoring (MONCHL) to get NETTIME values displayed, see [“Attributes that control real-time monitoring” on page 322](#). For more details about NETTIME, see [Channel status properties](#).

- Check whether the channel is using message retry.

If the receiver channel fails to put a message to a target queue, it might use message retry processing, rather than put the message to a dead-letter queue immediately. Retry processing can cause the batch to slow down. In between MQPUT attempts, the channel will have STATUS(PAUSED), indicating that it is waiting for the message retry interval to pass.

Determining whether the channel can process messages fast enough

If there messages are building up on the transmission queue, but you have found no processing problems, determine whether the channel can process messages fast enough.

Before you begin

Issue the following command repeatedly over a period of time to gather performance data about the channel:

```
DIS CHSTATUS(QM1.TO.QM2) ALL
```

About this task

Confirm that there are no uncommitted messages on the transmission queue, as described in [“Checking that messages on the queue are available”](#) on page 326, then check the XQTIME field in the output from the display channel status command. When the values of the XQTIME indicators are consistently high, or increase over the measurement period, the indication is that the channel is not keeping pace with the putting applications.

Perform the following tests:

Procedure

1. Check whether exits are processing.

If exits are used on the channel that is delivering these messages, they might add to the time spent processing messages. To identify if this is the case, do the following checks:

a) In the output of the command DIS CHSTATUS(QM1.TO.QM2) ALL, check the EXITTIME field.

If the time spent in exits is higher than expected, review the processing in your exits for any unnecessary loops or extra processing, especially in message, send, and receive exits. Such processing affects all messages moved across the channel.

b) In the output of the command DIS CHSTATUS(QM1.TO.QM2) ALL, check the SUBSTATE field.

If the channel has of one of the following substates for a significant time, review the processing in your exits:

- SCYEXIT
- RCVEXIT
- SENDEXIT
- MSGEXIT
- MREXIT

For more information on channel substates, see the table [Table 36 on page 331](#).

2. Check whether the network is slow.

If messages are not moving fast enough across a channel, it might be because the network is slow. To identify if this is the case, do the following checks:

a) In the output of the command DIS CHSTATUS(QM1.TO.QM2) ALL, check the NETTIME field.

These indicators are measured when the sending channel asks its partner for a response. This happens at the end of each batch and, when a channel is idle during heartbeating.

b) If this indicator shows that round trips are taking longer than expected, use other network monitoring tools to investigate the performance of your network.

3. Check whether the channel is using compression.

If the channel is using compression, this adds to the time spent processing messages. If the channel is using only one compression algorithm, do the following checks:

a) In the output of the command DIS CHSTATUS(QM1.TO.QM2) ALL, check the COMPTIME field.

These indicators show the time spent during compression or decompression.

- b) If the chosen compression is not reducing the amount of data to send by the expected amount, change the compression algorithm.
- 4. If the channel is using multiple compression algorithms, do the following checks:
 - a) In the output of the command `DIS CHSTATUS(QM1.TO.QM2) ALL`, check the `COMPTIME`, `COMPHDR`, and `COMPMSG` fields.
 - b) Change the compression algorithms specified on the channel definition, or consider writing a message exit to override the channel's choice of compression algorithm for particular messages if the rate of compression, or choice of algorithm, is not providing the required compression or performance.

Solving problems with cluster channels

If you have a build up of messages on the `SYSTEM.CLUSTER.TRANSMIT.QUEUE` queue, the first step in diagnosing the problem is discovering which channel, or channels, are having a problem delivering messages.

About this task

To discover which channel, or channels, using the `SYSTEM.CLUSTER.TRANSMIT.QUEUE` are having a problem delivering messages. Perform the following checks:

Procedure

1. Issue the following command:

```
DIS CHSTATUS(*) WHERE(XQMSGSA GT 1)
```

Note: If you have a busy cluster that has many messages moving, consider issuing this command with a higher number to eliminate the channels that have only a few messages available to deliver.

2. Look through the output for the channel, or channels, that have large values in the field `XQMSGSA`. Determine why the channel is not moving messages, or is not moving them fast enough. Use the tasks outlined in [“Monitoring channels” on page 327](#) to diagnose the problems with the channels found to be causing the build up.

Monitoring clusters

Within a cluster you can monitor application messages, control messages, and logs. There are special monitoring considerations when the cluster load balances between two or more instances of a queue.

Monitoring application messages in the cluster

Typically, all cluster messages that leave the queue manager pass through the `SYSTEM.CLUSTER.TRANSMIT.QUEUE`, irrespective of which cluster sender channel is being used to transmit the message. Each channel is draining messages targeted for that channel in parallel with all other cluster sender channels. A growing build-up of messages on this queue can indicate a problem with one or more channels and must be investigated:

- The depth of the queue must be monitored appropriately for the cluster design.
- The following command returns all channels that have more than one message that is waiting on the transmit queue:

```
DIS CHSTATUS(*) WHERE(XQMSGSA GT 1)
```

With all cluster messages on a single queue, it is not always easy to see which channel has problems when it begins to fill up. Using this command is an easy way to see which channel is responsible.

You can configure a cluster queue manager to have multiple transmission queues. If you change the queue manager attribute `DEFCLXQ` to `CHANNEL`, every cluster-sender channel is associated with a different cluster transmit queue. Alternatively you can configure separate transmission queues manually.

To display all the cluster transmit queues that are associated with cluster-sender channels, run the command:

```
DISPLAY CLUSQMGR (qmgrName) XMITQ
```

Define cluster transmission queues so that they follow the pattern of having the fixed stem of the queue name on the left. You can then query the depth of all the cluster transmission queues returned by the **DISPLAY CLUSMGR** command, by using a generic queue name:

```
DISPLAY QUEUE (qname *) CURDEPTH
```

Monitoring control messages in the cluster

The `SYSTEM.CLUSTER.COMMAND.QUEUE` queue is used for processing all cluster control messages for a queue manager, either generated by the local queue manager or sent to this queue manager from other queue managers in the cluster. When a queue manager is correctly maintaining its cluster state, this queue tends toward zero. There are situations where the depth of messages on this queue can temporarily grow however:

- Having lots of messages on the queue indicates churn in the cluster state.
- When making significant changes, allow the queue to settle in between those changes. For example, when moving repositories, allow the queue to reach zero before moving the second repository.

While a backlog of messages exists on this queue, updates to the cluster state or cluster-related commands are not processed. If messages are not being removed from this queue for a long time, further investigation is required, initially through inspection of the queue manager error logs (or CHINIT logs on z/OS) which might explain the process that is causing this situation.

The `SYSTEM.CLUSTER.REPOSITORY.QUEUE` holds the cluster repository cache information as a number of messages. It is usual for messages to always exist on this queue, and more for larger clusters. Therefore, the depth of messages on this queue is not an issue for concern.

Monitoring logs

Problems that occur in the cluster might not show external symptoms to applications for many days (and even months) after the problem originally occurs due to the caching of information and the distributed nature of clustering. However, the original problem is often reported in the IBM MQ error logs (and CHINIT logs on z/OS). For this reason, it is vital to actively monitor these logs for any messages written that relate to clustering. These messages must be read and understood, with any action taken where necessary.

For example: A break in communications with a queue manager in a cluster can result in knowledge of certain cluster resources that are being deleted due to the way that clusters regularly revalidate the cluster resources by republishing the information. A warning of such an event potentially occurring is reported by the message [AMQ9465](#) or [CSQX465I](#) on z/OS systems. This message indicates that the problem needs to be investigated.

Special considerations for load balancing

When the cluster load balances between two or more instances of a queue, consuming applications must be processing messages on each of the instances. If one or more of those consuming applications terminates or stops processing messages, it is possible that clustering might continue to send messages to those instances of the queue. In this situation, those messages are not processed until the applications are functioning correctly again. For this reason the monitoring of the applications is an important part of the solution and action must be taken to reroute messages in that situation. An example of a mechanism to automate such monitoring can be found in this sample: [The Cluster Queue Monitoring sample program \(AMQSCLM\)](#).

Related concepts

[“Tuning distributed publish/subscribe networks” on page 399](#)

Use the tuning tips in this section to help improve the performance of your IBM MQ distributed publish/subscribe clusters and hierarchies.

[“Balancing producers and consumers in publish/subscribe networks” on page 405](#)

An important concept in asynchronous messaging performance is *balance*. Unless message consumers are balanced with message producers, there is the danger that a backlog of unconsumed messages might build up and seriously affect the performance of multiple applications.

Monitoring transmission queue switching

It is important that you monitor the process of cluster-sender channels switching transmission queues so that the impact on your enterprise is minimized. For example, you should not attempt this process when the workload is high or by switching many channels simultaneously.

The process of switching channels

The process used to switch channels is:

1. The channel opens the new transmission queue for input and starts getting messages from it (using get by correlation ID)
2. A background process is initiated by the queue manager to move any messages queued for the channel from its old transmission queue to its new transmission queue. While messages are being moved any new messages for the channel are queued to the old transmission queue to preserve sequencing. This process might take a while to complete if there are a large number of messages for the channel on its old transmission queue, or new messages are rapidly arriving.
3. When no committed or uncommitted messages remain queued for the channel on its old transmission queue then the switch is completed. New messages are now put directly to the new transmission queue.

To avoid the eventuality of numerous channels switching simultaneously IBM MQ provides the ability to switch the transmission queue of one or more channels that are not running. On:

- IBM MQ for Multiplatforms the command is called **runswchl**
- IBM MQ for z/OS the CSQUTIL utility can be used to process a SWITCH CHANNEL command instead

Monitoring the status of switch operations

To understand the status of switch operations administrators can perform the following actions:

- Monitor the queue manager error log (AMQERR01.LOG) where messages are output to indicate the following stages during the operation:
 - The switch operation has started
 - The moving of messages has started
 - Periodic updates on how many messages are left to move (if the switch operation does not complete quickly)
 - The moving of messages has completed
 - The switch operation has completed

On z/OS, these messages are output to the queue manager job log, not the channel initiator job log, although a single message is output by a channel to the channel initiator job log if it initiates a switch when starting.

- Use the DISPLAY CLUSQMGR command to query the transmission queue that each cluster-sender channel is currently using.
- Run the **runswchl** command (or CSQUTIL on z/OS) in query mode to ascertain the switching status of one or more channels. The output of this command identifies the following for each channel:
 - Whether the channel has a switch operation pending
 - Which transmission queue the channel is switching from and to

- How many messages remain on the old transmission queue

Each command is really useful, because in one invocation you can determine the status of every channel, the impact a configuration change has had and whether all switch operations have completed.

Potential issues that might occur

See [Potential issues when switching transmission queues](#) for a list of some issues that might be encountered when switching transmission queue, their causes, and most likely solutions.

Related concepts

“[Tuning distributed publish/subscribe networks](#)” on page 399

Use the tuning tips in this section to help improve the performance of your IBM MQ distributed publish/subscribe clusters and hierarchies.

“[Balancing producers and consumers in publish/subscribe networks](#)” on page 405

An important concept in asynchronous messaging performance is *balance*. Unless message consumers are balanced with message producers, there is the danger that a backlog of unconsumed messages might build up and seriously affect the performance of multiple applications.

Multi Monitoring application balancing

You can use the **DISPLAY APSTATUS** command to monitor the state of application balancing across a uniform cluster, and to investigate why the application is not balanced if that is unexpected.

Monitoring the current state of applications across queue managers in a cluster

From any queue manager in a uniform cluster you can get an overview of the current state of applications across all the queue managers of a cluster by running the [DIS APSTATUS](#) command.

From IBM MQ 9.2.0 the **TYPE** field is also displayed in the output.

For example, after a queue manager has just started you might see output like the following:

```
1 : DIS APSTATUS(*) type(APPL)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)          CLUSTER(UNIDEMO)
COUNT(8)                MOVCOUNT(8)
BALANCED(UNKNOWN)
TYPE (APPL)
```

This shows that there is an application called MYAPP in the uniform cluster, and currently there are eight instances, all eight of which are considered movable around the uniform cluster. The balanced value of UNKNOWN is a temporary value, which indicates that the queue manager has not yet been through an attempt to rebalance the application if necessary.

After a short time, you are more likely to see the following output:

```
1 : DIS APSTATUS(*) type(APPL)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)          CLUSTER(UNIDEMO)
COUNT(8)                MOVCOUNT(8)
BALANCED(NO)
TYPE (APPL)
```

This output shows that the application has eight instances, but that they are not balanced across the uniform cluster. At this point, it is worth looking at the distribution of applications in the cluster.

To do this, run the **DIS APSTATUS** command again. Note that you can run this command against any queue manager in the uniform cluster:

```
1 : DIS APSTATUS(*) type(QMGR)
AMQ8932I: Display application status details.
```



```

APPLNAME(MYAPP)          ACTIVE(YES)
COUNT(6)                MOVCOUNT(6)
BALSTATE(HIGH)           LMSGDATE(2019-05-24)
LMSGTIME(13:11:10)       QMNAME(UNID001)
QMID(UNID001_2019-05-24_13.09.35)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)          ACTIVE(YES)
COUNT(1)                MOVCOUNT(1)
BALSTATE(LOW)            LMSGDATE(2019-05-24)
LMSGTIME(13:11:03)       QMNAME(UNID002)
QMID(UNID002_2019-05-24_13.09.39)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)          ACTIVE(YES)
COUNT(1)                MOVCOUNT(1)
BALSTATE(LOW)            LMSGDATE(2019-05-24)
LMSGTIME(13:11:07)       QMNAME(UNID003)
QMID(UNID003_2019-05-24_13.09.43)
TYPE (QMGR)

```

From this, you can see at this point in time that queue manager UNID001 has six instances, but queue managers UNID0002 and UNID0003 only have one each. The BALSTATE output gives an indication of the balance state the last time the queue manager reported it. However, be aware that the instance count might be more recent than the BALSTATE field.

This output also gives a good indication that the uniform cluster is distributing information about this application. The application is known to all queue manager instances in the uniform cluster, and the last message date and time is very recent.

Also, the ACTIVE field indicates that all queue managers in the cluster are considered to be communicating with one another; if ACTIVE is set to NO on any queue manager, it indicates that regular communication with it has been interrupted.

If left to rebalance itself, the output would eventually end up something like:

```

1 : DIS APSTATUS(*) type(QMGR)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)          ACTIVE(YES)
COUNT(3)                MOVCOUNT(3)
BALSTATE(OK)             LMSGDATE(2019-05-24)
LMSGTIME(13:14:22)       QMNAME(UNID001)
QMID(UNID001_2019-05-24_13.09.35)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)          ACTIVE(YES)
COUNT(3)                MOVCOUNT(3)
BALSTATE(OK)             LMSGDATE(2019-05-24)
LMSGTIME(13:13:53)       QMNAME(UNID002)
QMID(UNID002_2019-05-24_13.09.39)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)          ACTIVE(YES)
COUNT(2)                MOVCOUNT(2)
BALSTATE(OK)             LMSGDATE(2019-05-24)
LMSGTIME(13:13:47)       QMNAME(UNID003)
QMID(UNID003_2019-05-24_13.09.43)
TYPE (QMGR)

```

The BALSTATE of the queue managers is now OK indicating that a steady state has been achieved.

Monitoring individual application instances

You can look at the individual application instances, but you do this for each queue manager. Looking at the output on UNID001:

```

1 : DIS APSTATUS(*) type(LOCAL)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)
CONNTAG(MQCT02DFE75C02EA0A20UNID001_2019-05-24_13.09.35MYAPP)
CONNS(1)                 IMMREASN(NONE)
IMMCOUNT(0)              IMMDATE( )
IMMTIME( )               MOVABLE(YES)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)
CONNTAG(MQCT02DFE75C02E50A20UNID001_2019-05-24_13.09.35MYAPP)

```

```

CONNS(1)                                IMMREASN(NONE)
IMMCOUNT(0)                             IMMDATE( )
IMMTIME( )                              MOVABLE(YES)
AMQ8932I: Display application status details.
APPLNAME(MYAPP)
CONNTAG(MQCT02DFE75C02E60A20UNID001_2019-05-24_13.09.35MYAPP)
CONNS(1)                                IMMREASN(NONE)
IMMCOUNT(0)                             IMMDATE( )
IMMTIME( )                              MOVABLE(YES)
TYPE (LOCAL)
One MQSC command read.

```

The presence of MOVABLE(YES) indicates this instance can be moved if needed to another queue manager in the cluster. In the following example, the application instance is not movable because it is not connecting as a client:

```

3 : DISPLAY APSTATUS('ServerApp') TYPE(LOCAL)
AMQ8932I: Display application status details.
APPLNAME(ServerApp)
CONNTAG(MQCT02DFE75C01800B20UNID001_2019-05-24_13.09.35ServerApp)
CONNS(1)                                IMMREASN(NOTCLIENT)
IMMCOUNT(0)                             IMMDATE( )
IMMTIME( )                              MOVABLE(NO)
TYPE (LOCAL)

```

You can use the **CONNTAG** field to see the individual queue manager connections from that instance, which can be useful if the application instance is IMMREASN(NOTRECONN). In the following example, the client application is not movable, and investigating the connect options shows it does not have MQCNO_RECONNECT in the **CONNOPTS** field.

```

1 : DISPLAY APSTATUS('ClientApp') TYPE(LOCAL)
AMQ8932I: Display application status details.
APPLNAME(ClientApp)
CONNTAG(MQCT02DFE75C01CB0B20UNID001_2019-05-24_13.09.35ClientApp)
CONNS(1)                                IMMREASN(NOTRECONN)
IMMCOUNT(0)                             IMMDATE( )
IMMTIME( )                              MOVABLE(NO)
TYPE (LOCAL)

2 : DISPLAY CONN(*) TYPE(CONN) WHERE(CONNTAG eq
'MQCT02DFE75C01CB0B20UNID001_2019-05-24_13.09.35ClientApp') ALL
AMQ8276I: Display Connection details.
CONN(02DFE75C01CB0B20)
EXTCONN(414D5143554E49443030312020202020)
TYPE (CONN)
PID(14656)                                TID(20)
APPLDESC(IBM MQ Channel)                 APPLTAG(ClientApp)
APPLTYPE(USER)                           ASTATE(NONE)
CHANNEL(SYSTEM.DEF.SVRCONN)              CLIENTID( )
CONNNAME(127.0.0.1)
CONNOPTS(MQCNO_HANDLE_SHARE_BLOCK,MQCNO_SHARED_BINDING)
USERID(MyUserId)                         UOWLOG( )
UOWSTDA( )                               UOWSTTI( )
UOWLOGDA( )                             UOWLOGTI( )
URTYPE(QMGR)
EXTURID(XA_FORMATID[] XA_GTRID[] XA_BQUAL[])
QMURID(0.0)                               UOWSTATE(NONE)
CONNTAG(MQCT02DFE75C01CB0B20UNID001_2019-05-24_13.09.35ClientApp)
TYPE (CONN)

```

Related concepts

[Automatic application balancing](#)

Monitoring performance and resource usage on z/OS

Use this topic to understand the facilities available to monitor the performance, and resource usage of your IBM MQ for z/OS subsystems.

Related tasks

[Configuring queue managers on z/OS](#)

Introduction to monitoring IBM MQ for z/OS

Use this topic as an overview of the monitoring facilities available for IBM MQ for z/OS. For example, obtaining snapshots, using IBM MQ trace, online monitoring, and events.

This topic describes how to monitor the performance and resource usage of IBM MQ.

- It outlines some of the information that you can retrieve and briefly describes a general approach to investigating performance problems. See [“Investigating performance problems” on page 345](#) for more information.
- It describes how you can collect statistics about the performance of IBM MQ by using SMF records.
- It describes how to gather accounting data to enable you to charge your customers for their use of your IBM MQ systems.
- It describes how to use IBM MQ events (alerts) to monitor your systems.

Here are some of the tools you might use to monitor IBM MQ; they are described in the sections that follow:

- Tools provided by IBM MQ:
 - [Using DISPLAY commands](#)
 - [“Using CICS adapter statistics” on page 340](#)
 - [“Using IBM MQ events” on page 342](#)
- z/OS service aids:
 - [“Using System Management Facility” on page 343](#)
- Other IBM licensed programs:
 - [Using the Resource Measurement Facility](#)
 - [Using Tivoli Decision Support for z/OS](#)
 - [Using the CICS monitoring facility](#)

Information about interpreting the data gathered by the performance statistics trace is given in [“Interpreting IBM MQ for z/OS performance statistics” on page 346](#).

Information about interpreting the data gathered by the accounting trace is given in [“Interpreting IBM MQ for z/OS accounting data” on page 385](#).

Getting snapshots of IBM MQ using the DISPLAY commands

IBM MQ provides the MQSC facility which can give a snapshot of the performance, and resource usage using the DISPLAY commands.

You can get an idea of the current state of IBM MQ by using the DISPLAY commands and, for the CICS adapter, the CICS adapter panels.

Using DISPLAY commands

You can use the IBM MQ MQSC DISPLAY or PCF Inquire commands to obtain information about the current state of IBM MQ. They provide information about the status of the command server, process definitions, queues, the queue manager, and its associated components. These commands are:

MQSC command	PCF command
DISPLAY ARCHIVE	Inquire Archive
DISPLAY AUTHINFO	Inquire Authentication Information Object
DISPLAY CFSTATUS	Inquire CF Structure Status
DISPLAY CFSTRUCT	Inquire CF Structure

MQSC command	PCF command
DISPLAY CHANNEL	Inquire Channel
DISPLAY CHINIT	Inquire Channel Initiator
DISPLAY CHSTATUS	Inquire Channel Status
DISPLAY CMDSERV	
DISPLAY CLUSQMGR	Inquire Cluster Queue Manager
DISPLAY CONN	Inquire Connection
DISPLAY GROUP	Inquire Group
DISPLAY LOG	Inquire Log
DISPLAY PROCESS	Inquire Process
DISPLAY QMGR	Inquire Queue Manager
DISPLAY QSTATUS	Inquire Queue Status
DISPLAY QUEUE	Inquire Queue
DISPLAY SECURITY	Inquire Security
DISPLAY STGCLASS	Inquire Storage Class
DISPLAY SYSTEM	Inquire System
DISPLAY TRACE	
DISPLAY USAGE	Inquire Usage

For the detailed syntax of each command, see [MQSC commands](#) or [PCF commands](#). All of the functions of these commands (except DISPLAY CMDSERV and DISPLAY TRACE) are also available through the operations and control panels.

These commands provide a snapshot of the system only at the moment the command was processed. If you want to examine trends in the system, you must start an IBM MQ trace and analyze the results over a period of time.

Using CICS adapter statistics

If you are an authorized CICS user, you can use the CICS adapter control panels to display CICS adapter statistics dynamically.

These statistics provide a snapshot of information related to CICS thread usage and situations when all threads are busy. The display connection panel can be refreshed by pressing the Enter key.

For more information about configuring the IBM MQ CICS adapter, see the [Configuring connections to MQ](#) section of the CICS documentation.

Using IBM MQ trace

You can record performance statistics and accounting data for IBM MQ by using the IBM MQ trace facility. Use this topic to understand how to control IBM MQ trace.

The data generated by IBM MQ is sent to:

- The System Management Facility (SMF), specifically as SMF record type 115, subtypes 1 and 2 for the performance statistics trace
- The SMF, specifically as SMF record type 116, subtypes zero, 1, and 2 for the accounting trace.

If you prefer, the data generated by the IBM MQ accounting trace can also be sent to the generalized trace facility (GTF).

Starting IBM MQ trace

You can start the IBM MQ trace facility at any time by issuing the IBM MQ [START TRACE](#) command.

Accounting data can be lost if the accounting trace is started or stopped while applications are running. To collect accounting data successfully, the following conditions must apply:

- The accounting trace must be active when an application starts, and it must still be active when the application finishes.
- If the accounting trace is stopped, any accounting data collection that was active stops.

You can also start collecting some trace information automatically if you specify YES on the SMFSTAT (SMF STATISTICS) and SMFACCT (SMF ACCOUNTING) parameters of the CSQ6SYSP macro.  These parameters are described in [Using CSQ6SYSP](#).

Before starting an IBM MQ trace, read [“Using System Management Facility”](#) on page 343.

Controlling IBM MQ trace

To control the IBM MQ trace data collection at start-up, specify values for the parameters in the CSQ6SYSP macro when you customize IBM MQ.  See [Using CSQ6SYSP](#) for details.

You can control IBM MQ tracing when the queue manager is running with these commands:

- START TRACE
- ALTER TRACE
- STOP TRACE

You can choose the destination to which trace data is sent. Possible destinations are:

SMF

System Management Facility

GTF

Generalized Trace Facility (accounting trace only)

SRV

Serviceability routine for diagnostic use by IBM service personnel

For daily monitoring, information is sent to SMF (the default destination). SMF data sets typically contain information from other systems; this information is not available for reporting until the SMF data set is dumped.

You can also send accounting trace information to the GTF. This information has an event identifier of 5EE.

 The [The MQI call and user parameter](#), and [z/OS generalized trace facility \(GTF\)](#) describes how to deal with IBM MQ trace information sent to the GTF.

For information about IBM MQ commands, see [MQSC commands](#).

Effect of trace on IBM MQ performance

Using the IBM MQ trace facility can have a significant effect on IBM MQ and transaction performance. For example, if you start a global trace for class 1 or for all classes, it is likely to increase processor usage and transaction response times by approximately 50%. However, if you start a global trace for classes 2 - 4 alone, the increase in processor usage and transaction response times is likely to be less than 1% additional processor cost to the cost of IBM MQ calls. The same applies for a statistics or accounting trace.

▶ z/OS Using IBM MQ online monitoring

You can collect monitoring data for queues and channels (including automatically defined cluster-server channels) by setting the MONQ, MONCHL, and MONACLS attributes.

Table 37 on page 342 summarizes the commands to set these attributes at different levels and to display the monitoring information.

Table 37. Setting and displaying attributes to control online monitoring			
Attribute	Applicable at this level	Set using command	Display monitoring information using command
MONQ	Queue	DEFINE QLOCAL DEFINE QMODEL ALTER QLOCAL ALTER QMODEL	DISPLAY QSTATUS
	Queue manager	ALTER QMGR	
MONCHL	Channel	DEFINE CHANNEL ALTER CHANNEL	DISPLAY CHSTATUS
	Queue manager	ALTER QMGR	
MONACLS	Queue manager	ALTER QMGR	

For full details of these commands, see [MQSC commands](#). For more information about online monitoring, see “Monitoring your IBM MQ network” on page 5.

▶ z/OS Using IBM MQ events

IBM MQ instrumentation events provide information about errors, warnings, and other significant occurrences in a queue manager. You can monitor the operation of all your queue managers by incorporating these events into your own system management application.

IBM MQ instrumentation events fall into the following categories:

Queue manager events

These events are related to the definitions of resources within queue managers. For example, an application attempts to put a message to a queue that does not exist.

Performance events

These events are notifications that a threshold condition has been reached by a resource. For example, a queue depth limit has been reached, or the queue was not serviced within a predefined time limit.

Channel events

These events are reported by channels as a result of conditions detected during their operation. For example, a channel instance is stopped.

Configuration events

These events are notifications that an object has been created, changed, or deleted.

When an event occurs, the queue manager puts an *event message* on the appropriate *event queue*, if defined. The event message contains information about the event that can be retrieved by a suitable IBM MQ application.

IBM MQ events can be enabled using the IBM MQ commands or the operations and control panels.

See “Event types” on page 33 for information about the IBM MQ events that generate messages, and for information about the format of these messages. See [Event message reference](#) for information about enabling the events.

Using System Management Facility

You can use SMF to collect statistics and accounting information. To use SMF, certain parameters must be set in z/OS and in IBM MQ.

What is SMF?

System management facility (SMF) is a z/OS service aid used to collect information from various z/OS subsystems. This information is dumped and reported periodically, for example, hourly. You can use SMF with the IBM MQ trace facility to collect data from IBM MQ. In this way you can monitor *trends*, for example, in system utilization and performance, and collect accounting information about each user ID using IBM MQ.

To record performance statistics (record type 115) to SMF specify the following in the SMFPRMxx member of SYS1.PARMLIB or with the SETSMF z/OS operator command.

```
SYS(TYPE(115))
```

To record accounting information (record type 116) to SMF specify the following in the SMFPRMxx member of SYS1.PARMLIB or with the SETSMF z/OS operator command.

```
SYS(TYPE(116))
```

To use the z/OS command **SETSMF**, either **PROMPT(ALL)** or **PROMPT(LIST)** must be specified in the SMFPRM xx member. For more information, see [SMFPRMxx \(system management facilities \(SMF\) parameters\)](#).

You can start collecting some trace information automatically if you specify YES on the **SMFSTAT (SMF STATISTICS)** and **SMFACCT (SMF ACCOUNTING)** parameters of the CSQ6SYSP macro; this is described in [Using CSQ6SYSP](#).

Specifying YES on the **SMFSTAT** and **SMFACCT** parameters enables you to collect trace information as a queue manager starts.

You can also start collection of the data when the queue manager is running with the **START TRACE** command, specifying **START TRACE(A)** or **START TRACE(S)**.

You can turn on or off the recording of accounting information at the queue or queue manager level using the **ACCTQ** parameter of the **DEFINE QLOCAL**, **DEFINE QMODEL**, **ALTER QLOCAL**, **ALTER QMODEL**, or **ALTER QMGR** commands.

You can control the collection of channel accounting data at the channel or queue manager level using the **STATCHL** parameter of the **DEFINE CHANNEL**, **ALTER CHANNEL** or **ALTER QMGR** commands.

You can specify the interval at which IBM MQ collects statistics and accounting data in one of these ways:

- You can collect statistics data and accounting data at different intervals, using **STATIME** (statistics data) and **ACCTIME** (accounting data) in your system parameters (described in [Using CSQ6SYSP](#)).
- You can collect statistics data and accounting data at the same interval by specifying a value for **STATIME** in your system parameters (described in [Using CSQ6SYSP](#)).
- You can collect statistics data and accounting data by specifying zero for **STATIME**.

SMF must be running before you can send data to it. For more information about SMF, see the [z/OS MVS System Management Facilities \(SMF\) manual](#).

For the statistics and accounting data to be reset, at least one MQI call must be issued during the accounting interval.

Allocating additional SMF buffers

When you start a trace, you must ensure that you allocate adequate SMF buffers. Specify SMF buffering on the **VSAM BUFSP** parameter of the access method services **DEFINE CLUSTER** statement. Specify CISZ(4096) and BUFSP(81920) on the **DEFINE CLUSTER** statement for each SMF VSAM data set.

If an SMF buffer shortage occurs, SMF rejects any trace records sent to it. IBM MQ sends a CSQW133I message to the z/OS console when this occurs. IBM MQ treats the error as temporary and remains active even though SMF data can be lost. When the shortage has been alleviated and trace recording has resumed, IBM MQ sends a CSQW123I message to the z/OS console.

Reporting data in SMF

You can use the SMF program IFASMFDP (or IFASMF DL if logstreams are being used) to dump SMF records to a sequential data set so that they can be processed.

There are several ways to report on this data, for example:

- Write an application program to read and report information from the SMF data set. You can then tailor the report to fit your exact needs.
- Use SupportPac MP1B or a monitoring product to interpret the data.

No interval CLASS(03) SMF accounting records produced during long running processes

You are collecting CLASS(3) SMF116 accounting records for IBM MQ, but are getting no records produced while a long running process runs.

The CLASS(3) SMF116 accounting records normally are produced only when a process ends. For long running processes, for example CICS, this might not produce a sufficient number of records, as the process can run for a month or longer. However, you might want to gather SMF116 records at set time intervals while a process is running.

To gather CLASS(3) SMF116 accounting records, you must set **SMFACCT = YES** and issue the following command:

```
START TRACE(ACCTG) DEST(SMF) CLASS(03)
```

If you have set CLASS(3) SMF116 active and an accounting collection broadcast occurs, an interval CLASS(3) SMF116 accounting record is produced for any process currently running that was also running at the time of the previous collection broadcast and has done any IBM MQ work since.

You can set the collection broadcast to occur on a regular time interval by setting **ACCTIME** in [CSQ6SYSP](#) as follows:

- If your **ACCTIME** has been set to a value greater than 0, that is your broadcast interval in minutes and seconds (in the format of 'mmm.ss').
- If your **ACCTIME** has been set to a value of -1 (default) then your **STATIME** value will be used for accounting data.
- If your **ACCTIME = 0**, the SMF broadcast of your system is used (**SMF INTVAL**).
- If your **ACCTIME = 0** and your **SMF INTVAL** is not set, no broadcast occurs and no interval records are produced.

Using other products with IBM MQ

You can use other products to help you to improve the presentation of, or to augment statistics related to, performance and accounting. For example, Resource Measurement Facility, Tivoli Decision Support, and CICS monitoring.

Using the Resource Measurement Facility

Resource Measurement Facility (RMF) is an IBM licensed program (program number 5685-029) that provides system-wide information about processor utilization, I/O activity, storage, and paging. You can use RMF to monitor the utilization of physical resources across the whole system dynamically. For more information, see the [*z/OS Resource Measurement Facility User's Guide*](#).

Using Tivoli Decision Support for z/OS

You can use Tivoli Decision Support for z/OS to interpret RMF and SMF records.

Tivoli Decision Support for z/OS is an IBM licensed program (program number 5698-B06) that enables you to manage the performance of your system by collecting performance data in a Db2® database and presenting the data in various formats for use in systems management. Tivoli Decision Support for can generate graphic and tabular reports using systems management data it stores in its Db2 database. It includes an administration dialog, a reporting dialog, and a log collector, all of which interact with a standard Db2 database.

This is described in the [*IBM Tivoli Decision Support for z/OS: Administration Guide and Reference*](#).

Using the CICS monitoring facility

The CICS monitoring facility provides performance information about each CICS transaction running. It can be used to investigate the resources used and the time spent processing transactions. For background information, see the [*CICS Performance Guide*](#) and [*Developing CICS System Programs*](#), together with the two companion reference manuals, formerly called the *CICS Customization Guide*.

Investigating performance problems

Performance problems can arise from various factors. For example, incorrect resource allocation, poor application design, and I/O restraints. Use this topic to investigate some of the possible causes of performance problems.

Performance can be adversely affected by:

- Buffer pools that are an incorrect size
- Lack of real storage
- I/O contention for page sets or logs
- Log buffer thresholds that are set incorrectly
- Incorrect setting of the number of log buffers
- Large messages
- Units of recovery that last a long time, incorporating many messages for each sync point
- Messages that remain on a queue for a long time
- RACF® auditing
- Unnecessary security checks
- Inefficient program design

When you analyze performance data, always start by looking at the overall system before you decide that you have a specific IBM MQ problem. Remember that almost all symptoms of reduced performance are magnified when there is contention. For example, if there is contention for DASD, transaction response

times can increase. Also, the more transactions there are in the system, the greater the processor usage and greater the demand for both virtual and real storage.

In such situations, the system shows heavy use of *all* its resources. However, the system is actually experiencing normal system stress, and this stress might be hiding the cause of a performance reduction. To find the cause of such a loss of performance, you must consider all items that might be affecting your active tasks.

Investigating the overall system

Within IBM MQ, the performance problem is either increased response time or an unexpected and unexplained heavy use of resources. First check factors such as total processor usage, DASD activity, and paging. An IBM tool for checking total processor usage is resource management facility (RMF). In general, you must look at the system in some detail to see why tasks are progressing slowly, or why a specific resource is being heavily used.


Start by looking at general task activity, then focus on particular activities, such as specific tasks or a specific time interval.

Another possibility is that the system has limited real storage; therefore, because of paging interrupts, the tasks progress more slowly than expected.

Investigating individual tasks

You can use the accounting trace to gather information about IBM MQ tasks. These trace records tell you a great deal about the activity that the task has performed, and about how much time the task spent suspended, waiting for latches. The trace record also includes information about how much Db2 and coupling facility activity were performed by the task.

Interpreting IBM MQ accounting data is described in [“Interpreting IBM MQ for z/OS accounting data” on page 385](#).

Long running units of work can be identified by the presence of message CSQR026I in the job log. This message indicates that a task has existed for more than three queue manager checkpoints and its log records have been shunted.  For a description of log record shunting, see [The log files](#).

Interpreting IBM MQ for z/OS performance statistics

Use this topic as an index to the different SMF records created by IBM MQ for z/OS.

IBM MQ for z/OS performance statistics are written as SMF type 115 records. Statistics records are produced periodically at a time interval specified by the **STATIME** parameter of the CSQ6SYSP system parameter module, or at the SMF global recording interval if you specify zero for **STATIME**. The information provided in the SMF records comes from the following components of IBM MQ:

Buffer manager	Manages the buffer pools in virtual storage and the writing of pages to page sets as the buffer pools become full. Also manages the reading of pages from page sets.
Coupling facility manager	Manages the interface with the coupling facility.
Data manager	Manages the links between messages and queues. It calls the buffer manager to process the pages with messages on them.
Db2 manager	Manages the interface with the Db2 database that is used as the shared repository.
Lock manager	Manages locks
Log manager	Manages the writing of log records, which are essential for maintaining the integrity of the system if there is a back out request, or for recovery, if there is a system or media failure.

Message manager	Processes all IBM MQ API requests.
Storage manager	Manages storage, for example, storage pool allocation, expansion, and deallocation.
Topic manager	Manages the topic and subscription information
Coupling facility SMDS manager	Manages the shared message data sets (SMDS) for large messages stored in the coupling facility.

IBM MQ statistics are written to SMF as SMF type 115 records. The following subtypes can be present:

1

System information, for example, related to the logs and storage.

2

Information about number of messages and paging information. Queue sharing group information related to the coupling facility and Db2.

5 and 6

Detailed information about internal storage usage in the queue manager address space. While you can view this information, some of it is intended only for IBM use.

7

Storage manager summary information. While you can view this information, some of it is intended only for IBM use.

201

Page set input/output information

215

Buffer pool information

216

Queue information

231

System information for the channel initiator address space.

Note that:

- Subtype 1, 2, 201, and 215 records are created with statistics trace class 1.
- Subtype 5, 6, and 7 records are created with statistics trace class 3.
- Subtype 231 records are created with statistics trace class 4.
- Subtype 216 records are created with statistics trace class 5.

The subtype is specified in the SM115STF field (shown in [Table 38 on page 350](#)).

Commands required to collect SMF type 115 and type 116 records

Use this topic as a reference for the commands needed to collect type 115 and type 116 SMF records.

Using the START TRACE command

Use the dynamic version of the [START TRACE](#) command with the following options to collect the records:

- `START TRACE(STAT) DEST(SMF) CLASS(*)` and `START TRACE(ACCTG) DEST(SMF) CLASS(*)` starts traces for classes 1 to 3
- `START TRACE(STAT) DEST(SMF) CLASS(4)` and `START TRACE(ACCTG) DEST(SMF) CLASS(4)` starts channel initiator statistics and channel accounting data respectively.

See [Planning for channel initiator SMF data](#) for more details on the class 4 channel initiator information.

- `START TRACE(STAT) DEST(SMF) CLASS(5)` starts queue statistics.



Attention: If you are using the console version of the command add `cpf` to the start of the command you issue.

For the CLASS(4) channel statistics, you need to set the STATCHL attribute on the CHANNEL definition. For more information, see [“Controlling channel statistics collection”](#) on page 165.

For the CLASS(5) queue statistics, you need to set the STATQ attribute on the QUEUE and/or QMGR definition. For more information, see [“Controlling queue statistics collection”](#) on page 164.

SMF type 115 records - IBM MQ statistics

To verify the current statistics gathering, issue the [DISPLAY TRACE](#) command `DISPLAY TRACE (STAT)`. You should see the following:

```
RESPONSE=MPX1
CSQW127I QML1 CURRENT TRACE ACTIVITY IS -
TNO TYPE CLASS DEST USERID RMID
02 STAT 01,02,03,04 SMF * *
END OF TRACE REPORT
```

Note: Setting the ZPARM SMFSTAT or ZPARM SMFACCT attribute to an asterisk now only controls classes 1, 2 and 3. You must turn on classes 4 and 5 using the START TRACE command.

SMF type 116 Class 3 & 4 data – IBM MQ task and channel accounting records

To verify the current accounting trace data collection, issue the [DISPLAY TRACE](#) command `DISPLAY TRACE (ACCT)`. You should see the following:

```
RESPONSE=MPX1
CSQW127I QML1 CURRENT TRACE ACTIVITY IS -
TNO TYPE CLASS DEST USERID RMID
03 ACCTG * SMF * *
05 ACCTG 04 SMF * *
END OF TRACE REPORT
CSQ9022I QML1 CSQWVC1 'DISPLAY TRACE' NORMAL COMPLETION
```

You can turn SMF116 data collection on and off dynamically as there is no cycle of the queue manager necessary, using the following commands:

```
START TRACE(ACCTG) DEST(SMF) CLASS(3)
START TRACE(ACCTG) DEST(SMF) CLASS(4)

STOP TRACE(ACCTG) DEST(SMF) CLASS(3)
STOP TRACE(ACCTG) DEST(SMF) CLASS(4)
```

The channel accounting data requires additional set-up to produce the SMF. The steps are:

- Issue the command to turn on gathering of the information for all sender receiver and client connections:

```
ALTER QMGR STATCHL(HIGH)
```

- Issue the command to turn on gathering of the information for auto-defined cluster channels if clustering is in use:

```
ALTER QMGR STATCLS(HIGH)
```

Note: In the preceding statements, the value of HIGH, MED, or LOW have the same effect.

- Verify that all channels have STATCHL set to QMGR:

```
DISPLAY CHANNEL(*) STATCHL
```

Verifying the STATIME



Attention: You should set the interval to be no more than fifteen minutes using the [SET SYSTEM](#) command, although the default is 30 minutes, because a shorter interval is advised.

If you are evaluating a queue sharing group (QSG), ensure that the STATIME is consistent across the QSG.

- Display the system settings using the following [DISPLAY SYSTEM](#) command:

```
DISPLAY SYSTEM
```

- If the STATIME value is zero, meaning that the statistics interval is set to the default LPAR value, then this is typically OK.

There have been instances of very high volume LPARs creating varying durations of the SMF data production. If in your evaluation of the data there have been widely varying durations, set the STATIME to a nonzero value for all queue managers.

- The following examples set the STATIME interval to five minutes:

```
SET SYSTEM STATIME(05)
```

or if at IBM MQ 9.2.4 or above:

```
SET SYSTEM STATIME(05.00)
```

Note: The new interval does not take effect until the current interval has elapsed, so you need to make the change prior to the start of data collection.

Verifying the ACCTIME value for queue managers that are at release level IBM MQ 9.2.4 or higher:

- If ACCTIME is set to -1, it is the same as the STATIME attribute.
- If this evaluation is for a QSG, ensure that the ACCTIME is consistent across the QSG. If it is not, set the values to be the same interval before starting to gather the data.
- If the ACCTIME is set to 30 or above, alter it to be no greater than 15 using the [SET SYSTEM](#) command. The following examples set the ACCTIME interval to five minutes:

```
SET SYSTEM ACCTIME(05.00)
```

Note: The new interval does not take effect until the current interval has elapsed, so you need to make the change prior to the start of data collection.

Verifying the ACCTQ setting on the queue managers

- Display the ACCTQ setting for the queue managers you are investigating using the following command:

```
DISPLAY QMGR ACCTQ
```

- If the value is ACCTQ(ON), no further action is needed. Otherwise, issue the following command:

```
ALTER QMGR ACCTQ(ON)
```

- If clustering is in use, ensure that accounting is turned on for the SYSTEM.CLUSTER.TRANSMIT.QUEUE and any other named cluster transmit queues hosted by the queue manager. To determine whether accounting is enabled for the cluster transmission queues:
 - Display the ACCTQ setting for the queue or queues, using the following command:

```
DISPLAY QL(SYSTEM.CLUSTER.TRANSMIT.QUEUE) ACCTQ
```

- If the value is ACCTQ(ON), no further action is needed. Otherwise, issue the following command:

```
ALTER QL(SYSTEM.CLUSTER.TRANSMIT.QUEUE) ACCTQ(ON)
```

- For all other high-volume queues, verify that the value of ACCTQ is ON, or set ACCTG to QMGR.

Layout of an SMF type 115 record

You can use this section as a reference for the format of an SMF type 115 record.

The standard layout for SMF records involves three parts:

SMF header

Provides format, identification, and time and date information about the record itself.

Self-defining section

Defines the location and size of the individual data records within the SMF record.

Data records

The actual data from IBM MQ that you want to analyze.

For more information about SMF record formats, see [z/OS MVS System Management Facilities \(SMF\)](#).

Related reference

[“The SMF header” on page 350](#)

Use this topic as a reference for the format of the SMF header.

[“Self-defining sections” on page 351](#)

Use this topic as a reference for format of the self-defining sections of the SMF record.

[“Examples of SMF statistics records” on page 352](#)

Use this topic to understand some example SMF records.

The SMF header

Use this topic as a reference for the format of the SMF header.

[Table 38 on page 350](#) shows the format of SMF record header (SM115).

Table 38. SMF record 115 header description						
Offset: Dec	Offset: Hex	Type	Len	Name	Description	Example
0	0	Structure	28	SM115	SMF record header.	
0	0	Integer	2	SM115LEN	SMF record length.	14A0
2	2		2		Reserved.	
4	4	Integer	1	SM115FLG	System indicator.	5E
5	5	Integer	1	SM115RTY	Record type. The SMF record type, for IBM MQ statistics records this is always 115 (X'73').	73
6	6	Integer	4	SM115TME	Time when SMF moved record.	00355575
10	A	Integer	4	SM115DTE	Date when SMF moved record.	0100223F
14	E	Character	4	SM115SID	z/OS subsystem ID. Defines the z/OS subsystem on which the records were collected.	D4E5F4F1 (MV41)
18	12	Character	4	SM115SSI	IBM MQ subsystem ID.	D4D8F0F7 (MQ07)
22	16	Integer	2	SM115STF	Record subtype.	0002
24	18	Character	3	SM115REL	Internal version number.	F9F4F6 (946)
27	1B		1		Reserved	

Table 38. SMF record 115 header description (continued)						
Offset: Dec	Offset: Hex	Type	Len	Name	Description	Example
28	1C	Character	0	SM115END	End of SMF header and start of self-defining section.	

Note:  If SMF116REL is 946 or greater, locate the QWHX header and take the IBM MQ version from the 6-character QWHXREL field. See [Self-defining sections](#).

Self-defining sections

Use this topic as a reference for format of the self-defining sections of the SMF record.

A self-defining section of a type 115 SMF record tells you where to find a statistics record, how long it is, and how many times that type of record is repeated (with different values). The self-defining sections follow the header, at fixed offsets from the start of the SMF record. Each statistics record can be identified by an eye-catcher string.

The following types of self-defining section are available to users for type 115 records. Each self-defining section points to statistics data related to one of the IBM MQ components. [Table 39 on page 351](#) summarizes the sources of the statistics, the eye-catcher strings, and the offsets of the self-defining sections from the start of the SMF record header.

Table 39. Offsets to self-defining sections				
Source of statistics	Record subtype (SM115STF)	Offset of self-defining section		Eye-catcher of data
		Dec	Hex	
Storage manager	1	100	X'64'	QSST
Log manager	1	116	X'74'	QJST
Message manager	2	36	X'24'	QMST
Data manager	2	44	X'2C'	QIST
No longer used. The self-defining section will be binary zeros.	2	52	X'34'	
Lock manager	2	60	X'3C'	QLST
Db2 manager	2	68	X'44'	Q5ST
Coupling Facility manager	2	76	X'4C'	QEST
Topic manager	2	84	X'54'	QTST
SMDS usage	2	92	X'5C'	QESD
Buffer manager - one for each buffer pool	215	36	X'24'	QPST
Channel initiator	231			QWSX
Data manager page set - one for each page set	201	36	X'24'	QIS1
Storage manager	5	36	X'24'	QSPH
Storage manager	6	36	X'24'	QSGM

Table 39. Offsets to self-defining sections (continued)				
Source of statistics	Record subtype (SM115STF)	Offset of self-defining section		Eye-catcher of data
		Dec	Hex	
Storage manager	7	36	X'24'	QSRS
Queues	216	36	X'24'	QQST

Note: Some of the storage manager information in subtype 5, 6 and 7 records is intended only for IBM use. Other self-defining sections that are not listed contain data for IBM use only.

Each self-defining section is two fullwords long and has this format:

```
ssssssssllllnnnn
```

where:

- ssssssss is a fullword containing the offset from the start of the SMF record.
- llll is a halfword giving the length of this data record.
- nnnn is a halfword giving the number of data records in this SMF record.

For more information see, [“Examples of SMF statistics records” on page 352.](#)

Note: Always use offsets in the self-defining sections to locate the statistics records.

QWHX header

V 10.0.0 - V 10.0.0

The QWHX header contains the queue manager release in the form *VVRRMM*, and the name of the queue sharing group that the queue manager is a member of. The QWHX is described in assembler macro `thlqual.SCSQMACS(CSQDQWHX)`.

The QWHX header is available in all SMF 115 subtypes and can be located from the self defining section, which is at offset 28 (X'1C') from the beginning of each record, as follows:

- Take the fullword offset, ssssssss, from the self defining section and add 52 bytes (X'34') to it, this offset is the start of the QWHX header.

z/OS Examples of SMF statistics records

Use this topic to understand some example SMF records.

Figure 24 on page 353 shows an example of part of the SMF record for subtype 1. Subtype 1 includes the storage manager and log manager statistics records. The SMF record header is shown underlined.

The self-defining section at offset X'64' refers to storage manager statistics and the self-defining section at offset X'74' refers to log manager statistics, both shown in **bold**.

The storage manager statistics record is located at offset X'0000011C' from the start of the header and is X'48' bytes long. There is one set of storage manager statistics, identified by the eye-catcher string QSST. The start of this statistics record is also shown in the example.

The log manager statistics record is located at offset X'00000164' from the start of the header and is X'78' bytes long. There is one set of log manager statistics, identified by the eye-catcher string QJST.


```

000000 02000000 5E730035 55750100 223FD4E5 *.....MV*
000010 F4F1D4D8 F0F70001 F6F0F000 000001DC *41MQ07..600....*
000020 00240001 00000000 00000000 00000000 *.....*
000030 00000000 00000000 00000000 0000007C *.....@*
000040 00400001 000000BC 00600001 00000000 *. ....~.....*
000050 00000000 00000000 00000000 00000000 *.....*
000060 00000000 0000011C 00480001 00000000 *.....*
000070 00000000 00000164 00780001 00000000 *.....*
000080 00000000 00000000 00000000 00000000 *.....*
:
000110 00000000 00000000 00000000 003C0048 *.....*
000120 D8E2E2E3 0000004F 00000003 00000002 *QSST...|.....*

```

Figure 24. SMF record 115, subtype 1

Figure 25 on page 354 shows an example of part of the SMF record for subtype 2. Subtype 2 includes the statistics records for the message, data, lock, coupling facility, topic, and Db2 managers. The SMF record header is shown underlined; the self-defining sections are shown alternately **bold** and *italic*.

- The self-defining section at offset X'24' refers to message manager statistics. The message manager statistics record is located at offset X'00000064' from the start of the header and is X'48' bytes long. There is one set of these statistics, identified by the eye-catcher string QMST.
- The self-defining section at offset X'2C' refers to data manager statistics. The data manager statistics record is located at offset X'000000AC' from the start of the header and is X'50' bytes long. There is one set of these statistics, identified by the eye-catcher string QIST.
- In releases prior to IBM MQ 9.1.0, the self-defining section at offset X'34' referred to buffer manager statistics. As this SMF record was taken from an IBM MQ 9.1.0 queue manager, the buffer manager self-defining section is set to zeros to indicate that there are no buffer manager statistics. Instead, these statistics are in SMF 115 subtype 215 records.
- The self-defining section at offset X'3C' refers to lock manager statistics. The lock manager statistics record is located at offset X'000000FC' from the start of the header and is X'20' bytes long. There is one set of these statistics, identified by the eye-catcher string QLST.
- The self-defining section at offset X'44' refers to Db2 manager statistics. The Db2 manager statistics record is located at offset X'0000011C' from the start of the header and is X'2A0' bytes long. There is one set of these statistics, identified by the eye-catcher string Q5ST.
- The self-defining section at offset X'4C' refers to coupling facility manager statistics. The coupling facility manager statistics record is located at offset X'000003BC' from the start of the header and is X'1008' bytes long. There is one set of these statistics, identified by the eye-catcher string QEST.
- The self-defining section at offset X'54' refers to topic manager statistics. The topic manager statistics record is located at offset X'000013C4' from the start of the header and is X'64' bytes long. There is one set of these statistics, identified by the eye-catcher string QTST.
- The self-defining section at offset X'5C' is for SMDS statistics. This self defining section is set to zeros indicating that SMDS is not being used.

```

000000 09F40000 5E730033 4DBE0113 142FD4E5 *.4...;...(.MV*
000010 F4F1D4D8 F2F10002 F9F1F000 00001428 *41MQ21..910....*
000020 00240001 00000064 00480001 000000AC *.*****
000030 00500001 00000000 00000000 000000FC *.*****
000040 00200001 0000011C 02A00001 000003BC *.*****
000050 10080001 000013C4 00640001 00000000 *......D.....*
000060 00000000 D40F0048 D8D4E2E3 00000000 *....M...QMST...*
000080 00000000 00000000 00000000 00000000 *.*****
000090 00000000 00000000 00000000 00000000 *.*****
0000A0 00000000 00000000 00000000 C90F0050 *......I..&*
0000B0 D8C9E2E3 00000000 00000000 00000000 *QIST.....*
0000C0 00000000 00000000 00000000 00000000 *.*****
0000D0 00000000 00000000 00000000 00000000 *.*****
0000E0 00000000 00000000 00000000 00000000 *.*****
0000F0 00000000 00000000 00000000 D30F0020 *......L...*
000100 D8D3E2E3 00000000 00000000 00000000 *QLST.....*
000110 00000000 00000000 00000000 F50F02A0 *......5...*
000120 D8F5E2E3 00000000 00000000 00000000 *Q5ST.....*
.
.

```

Figure 25. SMF record 115, subtype 2

Processing type 115 SMF records

Use this topic as a reference for processing type 115 SMF records.

You must process any data you collect from SMF to extract useful information. When you process the data, verify that the records are from IBM MQ and that they are the records you are expecting.

Validate the values of the following fields:

- SM115RTY, the SMF record number, must be X'73' (115)
- SM115STF, the record subtype, must be 0001, 0002, 0005, 0006, 0007, 0201, 0215, or 0231

Reading from the active SMF data sets (or SMF logstreams) is not supported. You must use the SMF program IFASMFDP (or IFASMF DL if logstreams are being used) to dump SMF records to a sequential data set so that they can be processed. For more information see [“Using System Management Facility” on page 343](#).

Details of the structures and fields can be found in IBM MQ SupportPac MP1B.

There is a C sample program called CSQ4SMFD which prints the contents of SMF type 115 and 116 records from the sequential data set. The program is provided as source in thlqual.SCSQC37S and in executable format in thlqual.SCSQLOAD. Sample JCL is provided in thlqual.SCSQPROC(CSQ4SMFJ).

Storage manager data records

Use this topic as a reference for storage manager data records.

The format of the storage manager statistics record is described in assembler macro thlqual.SCSQMACS(CSQDQSST).

The data contains information about the number of fixed and variable storage pools that the queue manager has allocated, expanded, contracted, and deleted during the statistics interval, plus the number of GETMAIN, FREEMAIN, and STORAGE requests to z/OS, including a count of those requests that were unsuccessful. Additional information includes a count of the number of times the short-on-storage condition was detected and a count of the number of abends that occurred as a result of that condition.

Additional data about storage usage in the queue manager is produced by class 2 and class 3 statistics trace. While you can view this information, some of it is intended only for IBM use.

- The format of the storage manager pool header statistics record, which is present in subtype 5 records, is described in assembler macro thlqual.SCSQMACS(CSQDQSPH).
- The format of the storage manager getmain statistics record, which is present in subtype 6 records, is described in assembler macro thlqual.SCSQMACS(CSQDQSGM).

- The format of the storage manager region summary record, which is present in subtype 7 records, is described in assembler macro `thlqual.SCSQMACS(CSQDQSRS)`.

Log manager data records

Use this topic as a reference for format of log manager data records.

The format of the log manager statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQJST)`.

In the statistics, these counts are important:

1. The total number of log write requests:

$$N_{\text{logwrite}} = \text{QJSTWRNW} + \text{QJSTWRF}$$

2. The total number of log read requests:

$$N_{\text{logread}} = \text{QJSTRBUF} + \text{QJSTRACT} + \text{QJSTRARH}$$

The problem symptoms that can be examined using log manager statistics are described in the following table.

Symptom 1 QJSTWTB is nonzero. Reason Tasks are being suspended while the in-storage buffer is being written to the active log. There might be problems writing to the active log. The OUTBUFF parameter within CSQ6LOGP is too small. Action Investigate the problems writing to the active log. Increase the value of the OUTBUFF parameter within CSQ6LOGP.
Symptom 2 The ratio: $\text{QJSTTV C} / N_{\text{logread}}$ is greater than 1%. Reason Log reads were initiated that had to read from an archive log, but IBM MQ could not allocate a data set because MAXRTU data sets were already allocated. Action Increase MAXRTU.
Symptom 3 The ratio: $\text{QJSTRARH} / N_{\text{logread}}$ is larger than normal. Reason Most log read requests should come from the output buffer or the active log. To satisfy requests for back out, unit-of-recovery records are read from the in-storage buffer, the active log, and the archived logs. A long-running unit of recovery, extending over a period of many minutes, might have log records spread across many different logs. This degrades performance because extra work has to be done to recover the log records. Action Change the application to reduce the length of a unit of recovery. Also, consider increasing the size of the active log to reduce the possibility of a single unit of recovery being spread out over more than one log. Other pointers The ratio $N_{\text{logread}} / N_{\text{logwrite}}$ gives an indication of how much work has to be backed out.

Symptom 4

QJSTLLCP is more than 10 an hour.

Reason

On a busy system, you would expect to see typically 10 checkpoints an hour. If the QJSTLLCP value is larger than this, it indicates a problem in the setup of the queue manager.

The most likely reason for this is that the LOGLOAD parameter in CSQ6SYSP is too small. The other event that causes a checkpoint is when an active log fills up and switches to the next active log data set. If your logs are too small, this can cause frequent checkpoints.

QJSTLLCP is a count of the total number of checkpoints.

Action

Increase the LOGLOAD parameter, or increase the size of your log data sets as required.

Symptom 5

QJSTCmpFail > 0 or QJSTCmpComp not much less than QJSTCmpUncmp

Reason

The queue manager is unable to significantly compress log records.

QJSTCmpFail is the number of times the queue manager was unable to achieve any reduction in record length. You should compare the number to QJSTCmpReq (number of compression requests) to see if the number of failures is significant.

QJSTCmpComp is the total of compressed bytes written to the log and QJSTCmpUncmp is the total bytes before compression. Neither total contains bytes written for log records that were not eligible for compression. If the numbers are similar then compression has achieved little benefit.

Action

Turn off log compression. Issue the SET LOG COMPLOG(NONE) command. See the [SET LOG](#) command for details.

Note: In the first set of statistics produced after system startup, there might be significant log activity due to the resolution of in-flight units of recovery.

Changes to zHyperWrite

From IBM MQ 9.3.5 onwards, zHyperWrite processing alters the behavior of the QJSTHWC and QJSTHWE statistics.

Prior to IBM MQ 9.3.5:

- QJSTHWC is the number of log data sets used in the SMF interval that are on zHyperWrite capable volumes. The zHyperWrite capability of a volume is obtained at queue manager start up and can change over time, so this information might be stale.
- QJSTHWE is the number of log data sets used in the SMF interval that are written to with zHyperWrite enabled. This happens if the log data set is on a zHyperWrite capable volume and if the queue manager has been enabled for zHyperWrite by setting ZHYWRITE(YES).

From IBM MQ 9.3.5:

- QJSTHWC is the number of log data sets used in the SMF interval that are on zHyperWrite capable volumes. The zHyperWrite capability of a volume is obtained at queue manager start up and can change over time, so this information might be stale.
- QJSTHWE is the number of log data sets used in the SMF interval that are written to with zHyperWrite enabled. This happens if the queue manager has been enabled for zHyperWrite by setting ZHYWRITE(YES).



Attention: The zHyperWrite capability of a volume can change over time. From IBM MQ 9.3.5 this can result in scenarios where QJSTHWE is greater than QJSTHWC.

zHyperLink statistics

From IBM MQ 10.0.0 onwards, log manager statistics have been added to track the performance of zHyperLink.

The following statistics have been added to the QJST:

Name	Type	Length bytes	Description
QJSTHLSCIW	Unsigned integer	4	Number of single page writes which requested zHyperLink.
QJSTHLMCIW	Unsigned integer	4	Number of multiple page writes which requested zHyperLink.
QJSTHLCICNTW	Unsigned integer	4	Number of pages written where the write requested zHyperLink.
QJSTHLSCIS	Unsigned integer	4	Number of page writes which successfully used zHyperLink. That is, the write was made synchronously.
QJSTHLMCIS	Unsigned integer	4	Number of multiple page writes which successfully used zHyperLink. That is, the write was made synchronously.
QJSTHLCICNTS	Unsigned integer	4	Number of pages written successfully using zHyperLink. That is, the write was made synchronously.
QJSTHLSCIF	Unsigned integer	4	Number of single page writes which tried to use zHyperLink, but the write could not be made synchronously. Instead the write was successfully made asynchronously.
QJSTHLMCIF	Unsigned integer	4	Number of multiple page writes which tried to use zHyperLink, but the write could not be made synchronously. Instead the write was successfully made asynchronously.
QJSTHLCICNTF	Unsigned integer	4	Number of pages written where the write tried to use zHyperLink, but the write could not be made synchronously. Instead the write was successfully made asynchronously.
QJSTHLSCICON CNT	Unsigned integer	4	Number of single page writes before a zHyperLink write session was established. See Note “1” on page 358 .
QJSTHLMCICON CNT	Unsigned integer	4	Number of multiple page writes before a zHyperLink write session was established. See Note “1” on page 358 .
QJSTHLCONCIC NT	Unsigned integer	4	Number of pages written before a zHyperLink connection was established. See Note “1” on page 358 .
QJSTHLSCIWTM AX	Unsigned integer	8	Longest zHyperLink write time in interval, single page.
QJSTHLSCIWTM IN	Unsigned integer	8	Shortest zHyperLink write time in interval, single page.

Name	Type	Length bytes	Description
QJSTHLSCIWTTOT	Unsigned integer	8	Total zHyperLink write time in interval, single page.
QJSTHLMCIWTMAX	Unsigned integer	8	Longest zHyperLink write time in interval, multiple pages.
QJSTHLMCIWTMIN	Unsigned integer	8	Shortest zHyperLink write time in interval, multiple pages.
QJSTHLMCIWTTOT	Unsigned integer	8	Total zHyperLink write time in interval, multiple pages.
QJSTHLIOSQU	Unsigned integer	16	Sum of squares of zHyperLink write times, for single page writes.
QJSTHLC	Unsigned integer	4	Number of new logs, used in this SMF interval, which are zHyperLink capable.
QJSTHLE	Unsigned integer	4	Number of new logs, used in this SMF interval, which are zHyperLink enabled.

Note:

1. When switching to a new active log copy, there might be a period of time before a zHyperLink write session is established with the DASD.

Message manager data records

Use this topic as a reference for message manager data records.

The format of the message manager statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQMST)`.

The data gives you counts of different IBM MQ API requests.

 Additionally, the following OpenTelemetry trace related statistics are captured:

QMSTOCPA – Context propagated from environment on MQPUT

Number of times, in the interval, that trace context was automatically propagated from the application environment to the message on MQPUT.

QMSTOCGA – Context propagated to environment on MQGET

Number of times, in the interval, that trace context was automatically propagated from the message to the application environment on destructive MQGET.

QMSTOAHC – Context discarded on MQPUT, message already had context

Number of times, in the interval, that trace context from the environment was discarded on MQPUT because the message already had context.

QMSTOCPD – Incorrect traceparent on MQPUT

Number of times, in the interval, that an incorrectly formatted **traceparent** property was removed from the message and discarded during MQPUT. Producer spans are not generated.

QMSTOCMP – Context discarded on MQPUT, would exceed MAXPROPL

Number of times, in the interval, that trace context from the environment was discarded during MQPUT because adding it to the message would exceed the **MAXPROPL** value for the queue manager.

QMSTOCDG – Context discarded on MQGET

Number of times, in the interval, that invalid trace context was discarded from the message during MQGET.

QMSTOCPS – tracestate discarded on MQPUT

Number of times, in the interval that, **tracestate** was truncated or discarded during automatic propagation from the environment to the message on MQPUT. IBM MQ supports a maximum of 512 bytes of **tracestate**.

QMSTOCGS – tracestate discarded on MQGET

Number of times, in the interval that, **tracestate** was truncated or discarded during automatic propagation from the message to the environment on MQGET. IBM MQ supports a maximum of 512 bytes of **tracestate**.

QMSTOSCS – consumer spans scheduled for emission

Number of consumer spans scheduled for emission in the interval.

QMSTOSPS – producer spans scheduled for emission

Number of producer spans scheduled for emission in the interval.

QMSTOSNS – sampled bit off

Number of times, in the interval, that a message had trace context on MQPUT but sampled bit was off so a producer span was not generated.

QMSTOCDA – context discarded on MQGET, application cannot accept properties

Number of times, in the interval, that a message had context which was discarded on MQGET because of **OTELPCTL (AUTO)** processing or because application cannot accept properties.

QMSTOSPL – long running transaction limit reached on MQPUT

Number of messages, in the interval, which had trace context, but a producer span was not emitted due to the span limit for long running transactions being reached.

QMSTOSCL – long running transaction limit reached on MQGET

Number of messages, in the interval, which were got from the queue which had trace context, but no consumer span was emitted due to the span limit for long running transactions being reached.

QMSTOSFS – span storage limit reached

Number of spans, in the interval, which could not be emitted because the span storage limit was reached.

QMSTOSOK - number of spans emitted successfully

Number of spans, in the interval, emitted to SMF.

QMSTOSFR - number of spans not emitted as SMF not enabled

The number of spans that could not be emitted in the interval because SMF was not enabled or was not collecting type 1158 records.

QMSTOSFC - number of spans not emitted as SMF did not have capacity

The number of spans that could not be emitted in the interval because SMF did not have capacity for them.

QMSTOTSO – number of SMF records emitted successfully

The queue manager tries to batch multiple spans into a single SMF record prior to emitting it to SMF. This is the number of SMF records successfully emitted in the interval.

QMSTOTSR – number of SMF records not emitted as SMF not enabled

The queue manager tries to batch multiple spans into a single SMF record prior to emitting it to SMF. This is the number of SMF records that could not be emitted in the interval because SMF was not enabled or was not collecting type 1158 records.

QMSTOTSC – number of SMF records not emitted as SMF did not have capacity

The queue manager tries to batch multiple spans into a single SMF record prior to emitting it to SMF. This is the number of SMF records that could not be emitted in the interval because SMF did not have capacity for them.

QMSTOTPU – CPU time used by the span emission task

A single task is responsible for emitting spans via SMF. This is the amount of CPU used by the task in the interval in STCK units.

QMSTOTPW – time span emission task spent waiting for work

A single task is responsible for emitting spans via SMF. This is the amount of time the task waited for work in the interval in STCK units.

QMSTOTPD – time span emission task spent waiting for SMF

A single task is responsible for emitting spans via SMF. This is the amount of time the task spent waiting for SMF to emit records.

QMSTOTSK – number of span emission tasks

Currently this is always 1.

QMSTOBCD

The number of times that trace context was discarded because there was not sufficient space in the IMS user data prefix.

QMSTOBCP

The number of times that trace context was propagated to IMS via the IMS bridge.

QMSTOBSE

The number of IMS bridge spans scheduled for emission.

z/OS Data manager data records

Use this topic as a reference for the format of the Data Manager data records.

The format of the data manager statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQIST)`.

The data gives you counts of different object requests.

z/OS Data manager page set data records

Use this section as a reference for the format of the data manager page set data records

The format of the data manager page set statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQIS1)`.

The page set usage information helps to facilitate better management of local queues within a queue manager by recording information such as page set input-output rates and highest usage.

The data provides the same basic page set information as output by the MQSC **DISPLAY USAGE TYPE (PAGESET)** command, or the PCF Inquire Usage (**MQCMD_INQUIRE_USAGE**) command.

For example:

- The total pages
- The current used pages
- Unused persistent and nonpersistent pages
- Expansion method
- Number of extends
- Number of stripes

The data also provides some performance indicators, together with other performance information. For example:

- How many times deferred write, immediate write, and read page I/O requests happened during the SMF interval,
- The number of pages moved, elapsed time and number of read and write operations.
- How many pages written in checkpoints.
- Has expansion occurred during the SMF interval?
- How many times the page set became full.
- An indication of where new space is being allocated within the page set.

From the information displayed, you should be able to understand the general status of each page set, and consider whether you need to retune the system.

Related reference

[“The SMF header” on page 350](#)

Use this topic as a reference for the format of the SMF header.

[“Self-defining sections” on page 351](#)

Use this topic as a reference for format of the self-defining sections of the SMF record.

[“Examples of SMF statistics records” on page 352](#)

Use this topic to understand some example SMF records.

Buffer manager data records

Use this topic as a reference for the format of buffer manager data records.

The format of the buffer manager statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQPST)`.

Note: Buffer manager statistics records will only be created for buffer pools that are defined. If a buffer pool is defined but not used then no values will be set and its buffer manager statistics record will not contain any data.

For information about efficiently managing your buffer pools, see [“Managing your buffer pools” on page 362](#).

When interpreting the statistics, you are recommended to consider the following factors because the values of these fields can be used to improve the performance of your system:

1. If QPSTSOS, QPSTDMC, or QPSTIMW is greater than zero, you should either increase the size of the buffer pool or reallocate the page sets to different buffer pools.
 - QPSTSOS is the number of times that there were no buffers available for page get requests. If QPSTSOS ever becomes nonzero, it shows that IBM MQ is under severe stress. The buffer pool size should be increased. If increasing the buffer pool size does not make the value of QPSTSOS zero, there might be I/O contention on the DASD page sets.
 - QPSTDMC is the number of updates that were performed synchronously because there was either more than 95% of the pages in the buffer pool waiting for write I/O, or there was less than 5% of the buffer pool available for read requests. If this number is not zero, the buffer pool might be too small and should be enlarged. If increasing the buffer pool size does not reduce QPSTDMC to zero, there might be I/O contention on the DASD page sets.
 - QPSTIMW is a count of the number of times pages were written out synchronously. If QPSTDMC is zero, QPSTIMW is the number of times pages were found on the queue waiting for write I/O that had been there for at least two checkpoints.
2. For buffer pool zero and buffer pools that contain short-lived messages:
 - QPSTDWT should be zero, and the percentage QPSTCBSL/QPSTNBUF should be greater than 15%.

QPSTDWT is the number of times the asynchronous write processor was started because there was either more than 85% of the pages in the buffer pool waiting for write I/O, or there was less than 15% of the buffer pool available for read requests. Increasing the buffer pool size should reduce this value. If it does not, the pattern of access is one of long delays between puts and gets.
 - QPSTTPW might be greater than zero due to checkpointing activity.
 - QPSTRIO should be zero unless messages are being read from a page set after the queue manager is restarted.

The ratio of QPSTRIO to QPSTGETP shows the efficiency of page retrieval within the buffer pool. Increasing the buffer pool size should decrease this ratio and, therefore, increase the page retrieval efficiency. If this does not happen, it indicates that pages are not being frequently reaccessed. This implies a transaction pattern where there is a long delay between messages being put and then later retrieved.

The ratio of QPSTGETN to QPSTGETP indicates the number of times an empty page, as opposed to a non-empty page, has been requested. This ratio is more an indication of transaction pattern, than a value that can be used to tune the system.

- If QPSTSTL has a value greater than zero, this indicates that pages that have not been used before are now being used. This might be caused by an increased message rate, messages not being processed as fast as they were previously (leading to a buildup of messages), or larger messages being used.

QPSTSTL is a count of the number of times a page access request did not find the page already in the buffer pool. Again, the lower the ratio of QPSTSTL to (QPSTGETP + QPSTGETN) is, the higher the page retrieval efficiency. Increasing the buffer pool size should decrease this ratio but, if it does not, it is an indication that there are long delays between puts and gets.

- You are recommended to have sufficient buffers to handle your peak message rate.

3. For buffer pools with long-lived messages, where there are more messages than can fit into the buffer pool:

- $(QPSTRIO + QPSTWIO) / \text{Statistics interval}$ is the I/O rate to page sets. If this value is high, you should consider using multiple page sets on different volumes to allow I/O to be carried out in parallel.
- Over the period of time that the messages are processed (for example, if messages are written to a queue during the day and processed overnight) the number of read I/Os (QPSTRIO) should be approximately the total number of pages written (QPSTTPW). This shows that one page is read for every page written.

If QPSTRIO is much larger than QPSTTPW, this shows that pages are being read in multiple times. This might be a result of the application using MQGET by *MsgId* or *CorrelId* when the queue is not indexed, or browsing messages on the queue using get next.

The following actions might relieve this problem:

- a. Increase the size of the buffer pool so that there are enough pages to hold the queue, in addition to any changed pages.
- b. Use the INDXTYPE queue attribute, which allows a queue to be indexed by *MsgId* or *CorrelId* and eliminates the need for a sequential scan of the queue.
- c. Change the design of the application to eliminate the use of MQGET with *MsgId* or *CorrelId*, or the get next with browse option.

Note: Applications using long-lived messages typically process the first available message and do not use MQGET with *MsgId* or *CorrelId*, and they might browse only the first available message.

- d. Move page sets to a different buffer pool to reduce contention between messages from different applications.

Managing your buffer pools

To manage your buffer pools efficiently, you must consider the factors that affect the buffer pool I/O operations and also the statistics associated with the buffer pools.

The following factors affect buffer pool I/O operations.

- If a page containing the required data is not found in the buffer pool, it is read in synchronously to an available buffer from its DASD page set.
- Whenever a page is updated, it is put on an internal queue of pages to be (potentially) written out to DASD. This means that the buffer used by that page is unavailable for use by any other page until the buffer has been written to DASD.
- If the number of pages queued to be written to DASD exceeds 85% of the total number of buffers in the pool, an asynchronous write processor is started to put the buffers to DASD.

Similarly, if the number of buffers available for page get requests become less than 15% of the total number of buffers in the pool, the asynchronous write processor is started to perform the write I/O operations.

The write processor stops when the number of pages queued to be written to DASD has fallen to 75% of the total number of buffers in the pool.

- If the number of pages queued for writing to DASD exceeds 95% of the total number of buffers in the pool, all updates result in a synchronous write of the page to DASD.

Similarly, if the number of buffers available for page get requests becomes less than 5% of the total number of buffers in the pool, all updates result in a synchronous write of the page to DASD.

- If the number of buffers available for page get requests ever reaches zero, a transaction that encounters this condition is suspended until the asynchronous write processor has finished.
- If a page is frequently updated, the page spends most of its time on the queue of pages waiting to be written to DASD. Because this queue is in least recently used order, it is possible that a frequently updated page placed on this least recently used queue is never written out to DASD. For this reason, at the time of update, if the page is found to have been waiting on the write operation to DASD queue for at least two checkpoints, it is synchronously written to DASD. Updating occurs at checkpoint time and is suspended until the asynchronous write processor has finished.

The aim of this algorithm is to maximize the time pages spend in buffer pool memory while allowing the system to function if the system load puts the buffer pool usage under stress.

► **z/OS** **Lock manager data records**

Use this topic as a reference to the format of the lock manager data records.

The format of the lock manager statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQLST)`.

The records contain data about the following information:

- The number of lock get requests and lock release requests.
- The number of times a lock get request determined that the requested lock was already held.

► **z/OS** **Db2 manager data records**

Use this topic as a reference to the format of the Db2 manager data records.

The format of the Db2 manager statistics record is described in the following table and in assembler macro `thlqual.SCSQMACS(CSQDQ5ST)` and C header file `thlqual.SCSQC370(CSQDSMFC)`. The field names in C are all in lowercase, for example `q5st`, `q5stid`.

If the queue manager was not started as a member of a queue sharing group, no data is recorded in this record.

<i>Table 40. Db2 statistics record (Q5ST)</i>					
Offset: Dec	Offset: Hex	Type	Len	Name	Description
0	0	Structure	668	Q5ST	Db2 manager statistics
0	0	Bitstring	2	Q5STID	Control block identifier
2	2	Integer	2	Q5STLL	Control block length
4	4	Character	4	Q5STEYEC	Control block eye catcher
8	8	Character	660	Q5STZERO	QMST part cleared on occasion
8	8	Integer	4	NUMTASK	Number of server tasks
12	C	Integer	4	ACTTASK	Number of active server tasks
16	10	Integer	4	CONNCNT	Number of connect requests
20	14	Integer	4	DISCCNT	Number of disconnect requests

Table 40. Db2 statistics record (Q5ST) (continued)

Offset: Dec	Offset: Hex	Type	Len	Name	Description
24	18	Integer	4	DHIGMAX	Max. request queue depth
28	1C	Integer	4	ABNDCNT	Number of Db2SRV task abends
32	20	Integer	4	REQUCNT	Number of requests queued
36	24	Integer	4	DEADCNT	Number of deadlock timeouts
40	28	Integer	4	DELECNT	Number of delete requests
44	2C	Integer	4	LISTCNT	Number of list requests
48	30	Integer	4	READCNT	Number of read requests
52	34	Integer	4	UPDTCNT	Number of update requests
56	38	Integer	4	WRITCNT	Number of write requests
60	3C	Integer	4	SCSSEL	SCST (shared-channel-status) selects
64	40	Integer	4	SCSINS	SCST inserts
68	44	Integer	4	SCSUPD	SCST updates
72	48	Integer	4	SCSDEL	SCST deletes
76	4C	Integer	4	SSKSEL	SSKT (shared-sync-key) selects
80	50	Integer	4	SSKINS	SSKT inserts
84	54	Integer	4	SSKDEL	SSKT deletes
88	58	Integer	4	SCSBFTS	SCST number of times buffer too small
92	5C	Integer	4	SCSMAXR	SCST maximum rows on query
96	60	Integer	4	* (2)	Reserved
104	68	Character	8	DELETCUW	Cumulative STCK difference - Thread delete
112	70	Character	8	DELETMXW	Maximum STCK difference - Thread delete
120	78	Character	8	DELESCUW	Cumulative STCK difference - SQL delete
128	80	Character	8	DELESMXW	Maximum STCK difference - SQL delete
136	88	Character	8	LISTTCUW	Cumulative STCK difference - Thread list
144	90	Character	8	LISTTMXW	Maximum STCK difference - Thread list
152	98	Character	8	LISTSCUW	Cumulative STCK difference - SQL list
160	A0	Character	8	LISTSMXW	Maximum STCK difference - SQL list
168	A8	Character	8	READTCUW	Cumulative STCK difference - Thread read
176	B0	Character	8	READTMXW	Maximum STCK difference - Thread read
184	B8	Character	8	READSCUW	Cumulative STCK difference - SQL read
192	C0	Character	8	READSMXW	Maximum STCK difference - SQL read
200	C8	Character	8	UPDTTCUW	Cumulative STCK difference - Thread update
208	D0	Character	8	UPDTTMXW	Maximum STCK difference - Thread update
216	D8	Character	8	UPDTSCUW	Cumulative STCK difference - SQL update

Table 40. Db2 statistics record (Q5ST) (continued)

Offset: Dec	Offset: Hex	Type	Len	Name	Description
224	E0	Character	8	UPDTSMXW	Maximum STCK difference - SQL update
232	E8	Character	8	WRITTCUW	Cumulative STCK difference - Thread write
240	F0	Character	8	WRITTMXW	Maximum STCK difference - Thread write
248	F8	Character	8	WRITSCUW	Cumulative STCK difference - SQL write
256	100	Character	8	WRITSMXW	Maximum STCK difference - SQL write
264	108	Character	8	SCSSTCUW	Cumulative STCK difference - Thread select
272	110	Character	8	SCSSTMXW	Maximum STCK difference - Thread select
280	118	Character	8	SCSSSCUW	Cumulative STCK difference - SQL select
288	120	Character	8	SCSSSMXW	Maximum STCK difference - SQL select
296	128	Character	8	SCSITCUW	Cumulative STCK difference - Thread insert
304	130	Character	8	SCSITMXW	Maximum STCK difference - Thread insert
312	138	Character	8	SCSISCUW	Cumulative STCK difference - SQL insert
320	140	Character	8	SCSISMXW	Maximum STCK difference - SQL insert
328	148	Character	8	SCSUTCUW	Cumulative STCK difference - Thread update
336	150	Character	8	SCSUTMXW	Maximum STCK difference - Thread update
344	158	Character	8	SCSUSCUW	Cumulative STCK difference - SQL update
352	160	Character	8	SCSUSMXW	Maximum STCK difference - SQL update
360	168	Character	8	SCSDTCUW	Cumulative STCK difference - Thread delete
368	170	Character	8	SCSDTMXW	Maximum STCK difference - Thread delete
376	178	Character	8	SCSDSCUW	Cumulative STCK difference - SQL delete
384	180	Character	8	SCSDSMXW	Maximum STCK difference - SQL delete
392	188	Character	8	SSKSTCUW	Cumulative STCK difference - Thread select
400	190	Character	8	SSKSTMXW	Maximum STCK difference - Thread select
408	198	Character	8	SSKSSCUW	Cumulative STCK difference - SQL select
416	1A0	Character	8	SSKSSMXW	Maximum STCK difference - SQL select
424	1A8	Character	8	SSKITCUW	Cumulative STCK difference - Thread insert
432	1B0	Character	8	SSKITMXW	Maximum STCK difference - Thread insert
440	1B8	Character	8	SSKISCUW	Cumulative STCK difference - SQL insert
448	1C0	Character	8	SSKISMXW	Maximum STCK difference - SQL insert
456	1C8	Character	8	SSKDTCUW	Cumulative STCK difference - Thread delete
464	1D0	Character	8	SSKDTMXW	Maximum STCK difference - Thread delete
472	1D8	Character	8	SSKDSCUW	Cumulative STCK difference - SQL delete
480	1E0	Character	8	SSKDSMXW	Maximum STCK difference - SQL delete
488	1E8	Integer	4	LMSSEL	Number of Db2 BLOB read requests

Table 40. Db2 statistics record (Q5ST) (continued)

Offset: Dec	Offset: Hex	Type	Len	Name	Description
492	1EC	Integer	4	LMSINS	Number of Db2 BLOB insert requests
496	1F0	Integer	4	LMSUPD	Number of Db2 BLOB update requests
500	1F4	Integer	4	LMSDEL	Number of Db2 BLOB delete requests
504	1F8	Integer	4	LMSLIS	Number of Db2 BLOB list requests
508	IFC	64 bit integer	8	LMSSTCUW	Total elapsed time for all thread read BLOB requests
516	204	64 bit integer	8	LMSSTMXW	Maximum elapsed time for a thread read BLOB request
524	20C	64 bit integer	8	LMSSSCUW	Total elapsed time for all SQL read BLOB requests
532	214	64 bit integer	8	LMSSSMXW	Maximum elapsed time for an SQL read BLOB request
540	21C	64 bit integer	8	LMSITCUW	Total elapsed time for all thread insert BLOB requests
548	224	64 bit integer	8	LMSITMXW	Maximum elapsed time for a thread insert BLOB request
556	22C	64 bit integer	8	LMSISCUW	Total elapsed time for all SQL insert BLOB requests
564	234	64 bit integer	8	LMSISMXW	Maximum elapsed time for an SQL insert BLOB request
572	23C	64 bit integer	8	LMSUTCW	Total elapsed time for all thread update BLOB requests
580	244	64 bit integer	8	LMSUTMXW	Maximum elapsed time for a thread update BLOB request
588	24C	64 bit integer	8	LMSUSCUW	Total elapsed time for all SQL update BLOB requests
596	254	64 bit integer	8	LMSUSMXW	Maximum elapsed time for an SQL update BLOB request
604	25C	64 bit integer	8	LMSDTCW	Total elapsed time for all thread delete BLOB requests
612	264	64 bit integer	8	LMSDTMXW	Maximum elapsed time for a thread delete BLOB request
620	26C	64 bit integer	8	LMSDSCW	Total elapsed time for all SQL delete BLOB requests
628	274	64 bit integer	8	LMSDSMXW	Maximum elapsed time for an SQL delete BLOB request
636	27C	64 bit integer	8	LMSLTCW	Total elapsed time for all thread list BLOB requests
644	284	64 bit integer	8	LMSLTMXW	Maximum elapsed time for a thread list BLOB request

Table 40. Db2 statistics record (Q5ST) (continued)

Offset: Dec	Offset: Hex	Type	Len	Name	Description
652	28C	64 bit integer	8	LMSLSCUW	Total elapsed time for all SQL list BLOB requests
660	294	64 bit integer	8	LMSLSMXW	Maximum elapsed time for an SQL list BLOB request

The data contains counts for each request type that the Db2 resource manager supports. For these request types, maximum and cumulative elapse times are kept for the following:

- The time spent in the Db2 resource manager as a whole (called the thread time).
- The time that was spent performing the RRSF and SQL parts of the request (a subset of the thread time called the SQL time).

Information is also provided for:

- The number of server tasks attached.
- The maximum overall request depth against any of the server tasks.
- The number of times any of the server task requests terminated abnormally.

If the abnormal termination count is not zero, a requeue count is provided indicating the number of queued requests that were requeued to other server tasks as a result of the abnormal termination.

If the average thread time is significantly greater than the average SQL time, this might indicate that thread requests are spending an excessive amount of time waiting for a server task to process the SQL part of the request. If this is the case, examine the DHIGMAX field and, if the value is greater than one, consider increasing the number of Db2 server tasks specified in the QSGDATA parameter of the CSQ6SYSP system parameter macro.

Coupling facility manager data records

Use this topic as a reference to the format of the coupling facility manager data records.

The format of the coupling facility manager statistics record is described in the following table and in assembler macro thlqual.SCSQMACS(CSQDQEST) and C header file thlqual.SCSQC370(CSQDSMFC). The field names in C are all in lowercase, for example qest, qestid.

If the queue manager was not started as a member of a queue sharing group, no data is recorded in this record.

Table 41. Coupling facility statistics record (QEST)

Offset: Dec	Offset: Hex	Type	Len	Name	Description
0	0	Structure	4104	QEST	CF manager statistics
0	0	Bitstring	2	QESTID	Control block identifier
2	2	Integer	2	QESTLL	Control block length
4	4	Character	4	QESTYEC	Control block eye catcher
8	8	Character	4096	QESTZERO	QEST part cleared on occasion
8	8	Character	64	QESTSTUC (0:63)	Array (one entry per structure)
8	8	Character	12	QESTSTR	Structure name
20	14	Integer	4	QESTSTRN	Structure number
24	18	Integer	4	QESTCSEC	Number of IXLLSTE calls

Table 41. Coupling facility statistics record (QEST) (continued)					
Offset: Dec	Offset: Hex	Type	Len	Name	Description
28	1C	Integer	4	QESTCMEC	Number of IXLLSTM calls
32	20	Character	8	QESTSSTC	Time spent doing IXLLSTE calls
40	28	Character	8	QESTMSTC	Time spent doing IXLLSTM calls
48	30	Integer	4	QESTRSEC	Number of IXLLSTE redrives
52	34	Integer	4	QESTRMEC	Number of IXLLSTM redrives
56	38	Integer	4	QESTSFUL	Number of structure fulls
60	3C	Integer	4	QESTMNUS	Maximum number of entries in use
64	40	Integer	4	QESTMLUS	Maximum number of elements in use
68	44	Character	4	*	Reserved
4104	1008	Character	0	*	End of control block

The data contains information for each coupling facility list structure, including the CSQ_ADMIN structure, that the queue manager could connect to during the statistics interval. The information for each structure includes the following:

- The number of and cumulative elapsed times for IXLLSTE and IXLLSTM requests.
- The number of times a request had to be retried because of a timeout.
- The number of times a 'structure full' condition occurred.

Topic manager data records

Use this topic as a reference to the format of the topic manager data records.

The format of the topic manager statistics record is described in the following table and in assembler macro `thlqual.SCSQMACS(CSQDQTST)` and C header file `thlqual.SCSQC370(CSQDSMFC)`. The field names in C are all in lowercase, for example `qtst`, `qtstid`.

Table 42. Topic manager statistics record (QTST)					
Offset: Dec	Offset: Hex	Type	Len	Name	Description
0	0	Structure	96	QTST	Topic manager statistics
0	0	Bitstring	2	QTSTID	Control block identifier
2	2	Integer	2	QTSTLL	Control block length
4	4	Character	4	TESTEYEC	Control block eye catcher
8	8	Character	88	QTSTZERO	QTST part cleared on occasion
8	8	Integer	4	QTSTSTOT	Total subscription requests
12	0C	Integer	4	QTSTSDUR	Durable subscription requests
16	10	Integer	4	QTSTSHIG (1:3)	Subscription high water mark array (API, ADMIN, PROXY)
28	1C	Integer	4	QTSTSLOW (1:3)	Subscription low water mark array (API, ADMIN, PROXY)

Table 42. Topic manager statistics record (QTST) (continued)

Offset: Dec	Offset: Hex	Type	Len	Name	Description
40	28	Integer	4	QTSTSEXP	Subscriptions expired
44	2C	Integer	4	QTSTTMSG	Total messages put to Sub queue
48	30	Integer	4	QTSTSPHW	Single publish subscriber high water mark
52	34	Integer	4	QTSTPTOT (1:3)	Total Publication requests (API, ADMIN, PROXY)
64	40	Integer	4	QTSTPTHI	Total publish high water mark
68	44	Integer	4	QTSTPTLO	Total publish low water mark
72	48	Integer	4	QTSTPNOS	Count of publishes to no subscriber
76	4C	Integer	4	*	Reserved
80	50	Bitstring	8	QTSTETHW	Elapse time HW on publish
88	58	Bitstring	8	QTSTETTO	Elapse time total on publish

z/OS Coupling facility manager SMDS data records

Use this topic as a reference to the format of the coupling facility manager shared message data set (SMDS) data records.

The format of the coupling facility manager shared message data set (SMDS) statistics record is described in assembler macro `thlqual.SCSQMACS(CSQDQESD)` and in C header file `thlqual.SCSQC370(CSQDSMFC)`.

The statistics provide information about the utilization of the owned shared message data set, I/O activity for the group of shared message data sets, and SMDS buffer utilization.

If the queue manager was not started as a member of a queue sharing group, no data is recorded in this record.

z/OS Layout of channel initiator SMF type 115 records

The layout of channel initiator statistics data (SMF type 115, subtype 231) records is described in this topic.

Self-defining section

The self-defining section for the channel initiator statistics data follows the standard SMF header. It is structured in the standard triplet format. The format of the triplets is described in structure `qwsx` in the C programming language header file `thlqual.SCSQC370(CSQDSMFC)`, and in assembler macro `thlqual.SCSQMACS(CSQDQWSX)`.

Table 43 on page 369 shows the format of the self-defining section.

Table 43. Structure of the channel initiator statistics self-defining section

Offset: Dec	Offset : Hex	Type	Length	Name	Description
0	0	Integer	4	QWSX0PSO	Offset from the start of the SMF record to the first instrumentation standard header (QWHS)

Table 43. Structure of the channel initiator statistics self-defining section (continued)

Offset: Dec	Offset : Hex	Type	Length	Name	Description
4	4	Integer	2	QWSX0PSL	Length of the QWHS
6	6	Integer	2	QWSX0PSN	Number of instances of QWHS
8	8	Integer	4	QWSX0R1O	Offset from the start of the SMF record to the first channel initiator control information block (QCCT)
12	C	Integer	2	QWSX0R1L	Length of the QCCT
14	E	Integer	2	QWSX0R1N	Number of instances of QCCT
16	10	Integer	4	QWSX0R2O	Offset from the start of the SMF record to the first dispatcher task block (QCT_DSP)
20	14	Integer	2	QWSX0R2L	Length of the QCT_DSP
22	16	Integer	2	QWSX0R2N	Number of instances of QCT_DSP
24	18	Integer	4	QWSX0R3O	Offset from the start of the SMF record to the first adapter task block (QCT_ADP)
28	1C	Integer	2	QWSX0R3L	Length of the QCT_ADP
30	1E	Integer	2	QWSX0R3N	Number of instances of QCT_ADP
32	20	Integer	4	QWSX0R4O	Offset from the start of the SMF record to the first SSL task block (QCT_SSL)
36	24	Integer	2	QWSX0R4L	Length of the QCT_SSL
38	26	Integer	2	QWSX0R4N	Number of instances of QCT_SSL
40	28	Integer	4	QWSX0R5O	Offset from the start of the SMF record to the first DNS task block (QCT_DNS)
44	2C	Integer	2	QWSX0R5L	Length of the QCT_DNS
46	2E	Integer	2	QWSX0R5N	Number of instances of QCT_DNS

Typically one record contains all the data. If there are a large number of dispatchers, adapters, or SSL tasks, the data is split over more than one record.

If this happens, the count of instances of some type of tasks can be zero, and information about a group of tasks can be spread across multiple records. The channel initiator control information block (QCCT) is only present in the first record. For example the data could be split between two SMF records like this:

Table 44. Example data

Count	First record	Last record
QWHS	1	1
QCCT	1	0
QCT_DSP	50	5
QCT_ADP	0	10
QCT_SSL	0	3
QCT_DNS	0	1

This example shows that there were 55 dispatcher TCBs running during the SMF interval.

Instrumentation standard header (QWHS)

The format of the QWHS is described in structure `qwhs` in the C programming language header file `thlqual.SCSQC370(CSQDSMFC)`, and in assembler macro `thlqual.SCSQMACS(CSQDQWHS)`. It contains the following key fields that are relevant to channel initiator SMF 115 records:

Table 45. Key fields in the QWHS		
Name	Length	Description
QWHSNDA	1 byte	Number of self-defining sections
QWHSSSID	4 bytes	Subsystem name
QWHSSMFC	1 bit	Indicates whether there are multiple SMF records containing information for this interval. If this bit is on, information for this interval is continued in further SMF records. If this bit is off, this is the last or only record. The subsystem ID in QWHSSSID, and the SMF interval start time in QWHSTIME, can be used to group multiple records for the same interval.
QWHSTIME	8 bytes	Local time of the start of the interval in STCK format
QWHSDURN	8 bytes	Duration from the start of the interval to the end of the interval in STCK format
QWHSSTCK	8 bytes	End of the interval in UTC in STCK format

Channel initiator statistics data records

Use this topic as a reference for channel initiator statistics data records.

The format of the channel initiator statistics data record contains two parts:

- The first part is the channel initiator control information block, described in assembler macro `thlqual.SCSQMACS(CSQDQCCT)`. For further information, see [“Channel initiator control information block” on page 372](#).
- The second part is the channel initiator task block, described in assembler macro `thlqual.SCSQMACS(CSQDQCTA)`.

The channel initiator task block contains information about the four types of task within the CHINIT. For further information, see:

- [“Dispatcher tasks” on page 373](#)
- [“Adapter tasks” on page 374](#)
- [“Domain Name Server \(DNS\) task” on page 375](#)
- [“SSL tasks” on page 376](#)

Each task includes:

- The elapsed time that the task spent processing requests in the interval (*qcteltm*)
- The CPU time used by the task in the interval, which is made up of CPU used while processing requests and CPU used between requests (*qctcptm*)
- The total wait time of this task in the interval (*qctwtm*)
- The number of requests in the interval (*qctreqn*)

You can use this information to see how busy the task was, and determine whether you need to add more tasks based on the analysis.

For TLS and DNS tasks, the duration of the longest request (*qctlgdu*, *qctlsdu*) and the time of day when this occurred (*qctlgdm*, *qctlsdm*) are also included.

These can be useful to identify when channel requests took a long time. For example, a DNS lookup request going to a server outside of your enterprise taking seconds rather than milliseconds.

The CPU time (*qctcptm*) value includes all CPU consumed by the task, both processing requests and between processing requests. The elapsed time (*qcteltm*) value only includes time while processing requests. This means that the CPU time may be greater than the elapsed time.

The example accounting data in the following tasks has been formatted using IBM MQ SupportPac MP1B.

Both of the parts are also described in the C programming language header file `thlqual.SCSQC370` (CSQDSMFC). Note that the field names in C are all in lowercase, for example, *qcct*, *qct_adp*.

Channel initiator control information block

Use this topic as a reference for the channel initiator control information block.

The channel initiator control information block contains basic information for this CHINIT, including:

- CHINIT job name (*qcctjobn*)
- QSG name if it is in a queue sharing group (*qcctqsgn*)
- Peak number used of current channels (*qcctnocc*)
- Peak number used of active channels (*qcctnoac*)
- MAXCHL - maximum permitted current channels (*qcctmxcc*)
- ACTCHL - maximum permitted active channels (*qcctmxac*)
- TCPCHL - maximum permitted TCP/IP channels (*qcctmxtp*)
- LU62CHL - maximum permitted LU62 channels (*qcctmxlu*)
- 31-bit storage used by CHINIT in the extended private region (*qcctstus*). This information is also provided by the `CSQX004I` message in the CHINIT job log.
- 64-bit storage limit available to the CHINIT (*qcctslim*)
- 64-bit storage used by CHINIT (*qcctstab*). This information is also provided by the `CSQX004I` message in the CHINIT job log.

The format of the channel initiator control information block is described in structure *qcct* in the C programming language header file `thlqual.SCSQC370` (CSQDSMFC), and in assembler macro `thlqual.SCSQMACS` (CSQDQCCT).

You can use this information to see if the number of active channels is approaching the configured maximum value. Note that the number of current and active channels are the values when the record was created. So, between the two intervals there might have been more than this number of channels active.

Channel information from SMF data

Here is an example of channel information from SMF data:

```
MV4A,MQ27,2023/10/02,11:53:02,VRM:934,
From 2023/10/02,11:52:52 to 2023/10/02,11:53:02, duration 10 seconds.
Peak number used of current channels..... 1
Peak number used of active channels ..... 1
MAXCHL. Max allowed current channels..... 9999
ACTCHL. Max allowed active channels..... 9999
TCPCHL. Max allowed TCP/IP channels..... 9999
LU62CHL. Max allowed LU62 channels..... 200
31-bit storage used..... 436 MB
64-bit storage limit.....16384 PB
```

```
64-bit storage used..... 187 MB
64-bit storage free.....16384 PB
```

You can monitor the storage usage and see whether the value is trending upwards. If the total used is approaching the total storage available, you might be running out of storage, and so might not be able to support many more channels.

If the numbers of active current channels are trending towards the maximum number of channels, you might need to increase the maximum number of channels.

Dispatcher tasks

This topic contains example data for the dispatcher tasks statistics, and information about how to interpret the data.

The format of the dispatcher task block is described in structure `qct_dsp` in the C programming language header file `thlqual.SCSQC370(CSQDSMFC)`, and in assembler macro `thlqual.SCSQMACS(CSQDQCTA)`.

Example data

Task	Type	Requests	Busy %	CPU used, Seconds	CPU %	"avg CPU", uSeconds	"avg ET", uSeconds
0	DISP	26587	0.4	0.592463	0.1	22	127
1	DISP	26963	0.3	0.588092	0.1	22	112
2	DISP	864329	2.7	2.545668	0.3	3	28
3	DISP	26875	0.4	0.590825	0.1	22	120
4	DISP	26874	0.4	0.603285	0.1	22	123
Summ	DISP	971628	0.8	4.920332	0.1	5	38

The example data shows that there were five dispatchers. A channel is associated with a dispatcher when it starts. The channel initiator tries to distribute work across all the dispatchers when allocating a channel to a dispatcher. This example shows that one dispatcher is processing more requests than other dispatchers. This is normal, as some channels might stop, so the dispatcher is processing fewer channels, and some channels can be busier than others.

- 4.9 seconds of CPU were used by the dispatchers.
- The average request used 5 microseconds of CPU and took 38 microseconds elapsed time.
- A dispatcher is used to send and receive data over a communications network, and this is not usually dependent on external events. The average elapsed time should, therefore, be close to the average CPU time used. The CPU time (*qctcptm*) value includes all CPU consumed by the task, both processing requests and between processing requests.

The elapsed time (*qcteltm*) value only includes time while processing requests. This means that the CPU time may be greater than the elapsed time. If the CHINIT is delayed due to lack of CPU, then the ratio of average elapsed time to average CPU time is much larger, compared to when the CHINIT is not delayed for CPU.

- The average CPU used per request depends on the message traffic. For example, bigger messages use more CPU than smaller messages.

The fields are calculated from:

- Duration: *qwhs.qwhsdurn*
- Requests : *qctreqn*
- Busy %: *qcteltm* and duration
- CPU used: *qctcptm*
- CPU %: *qctcptm* and duration
- Average CPU: *qctcptm* and *qctreqn*
- Average ET: *qcteltm* and *qctreqn*

Usually, the number of dispatchers should be less than, or equal to, the number of processors in the LPAR. If you have more dispatchers than processors in the LPAR they might compete for CPU resources. For more information about tuning your system, see [SupportPac MP16](#).

Channels have an affinity to a dispatcher, so you might find that some dispatchers process many more requests than another dispatcher.

You can use the ALTER QMGR CHIDISPS() command to change the number of dispatchers used. Any change comes into effect the next time the channel initiator is started.

Adapter tasks

This topic contains example data for the adapter tasks statistics, and information about how to interpret the data.

The format of the adapter task block is described in structure `qct_adp` in the C programming language header file `thlqual1.SCSQC370(CSQDSMFC)`, and in assembler macro `thlqual1.SCSQMACS(CSQDQCTA)`.

Example data

Task	Type	Requests	Busy %	CPU used, Seconds	CPU %	"avg CPU", uSeconds	"avg ET", uSeconds
0	ADAP	470297	10.2	41.290670	4.6	88	194
1	ADAP	13907	0.6	1.589428	0.2	114	365
2	ADAP	2517	0.2	0.185325	0.0	74	746
3	ADAP	1095	0.1	0.085774	0.0	78	907
4	ADAP	535	0.1	0.040743	0.0	76	947
5	ADAP	220	0.0	0.016228	0.0	74	1175
6	ADAP	82	0.0	0.005521	0.0	67	1786
7	ADAP	80	0.0	0.004248	0.0	53	1160
Summ	ADAP	488733	1.4	43.217938	0.6	88	205

The fields are calculated from:

- Duration: *qwhs.qwhsdurn*
- Requests: *qctreqn*
- Busy %: *qcteltm* and duration
- CPU used: *qctcptm*
- CPU %: *qctcptm* and duration
- Average CPU: *qctcptm* and *qctreqn* average
- ET: *qcteltm* and *qctreqn*

This example shows that there were eight adapter tasks.

Adapter number 0

- Processed the majority of the requests (470297 out of 488733)
- Was busy 10.2% of the interval
- Used 41.3 seconds of CPU

Overall

The average CPU per request was 88 microseconds of CPU and took 205 microseconds

The adapters process IBM MQ requests. Some of these requests might wait, for example, for log I/O during a commit, so the average Elapsed Time per request has little meaning.

The CPU time (*qctcptm*) value includes all CPU consumed by the task, both processing requests and between processing requests. The elapsed time (*qcteltm*) value only includes time while processing requests. This means that the CPU time may be greater than the elapsed time.

When an IBM MQ request is made the first free adapter task is used.

- If there is at least one adapter that has been little used (less than 1%) busy, you have enough adapters.

- If at least one adapter was not used, you have enough adapters defined.
- If all the adapters were used, you might need to allocate more adapters.
- If all of the adapters were used, and they were all busy for most of the interval, you need to allocate more adapters.

You can use the ALTER QMGR CHIADAPS() command to change the number of adapters used. Any changes come into effect the next time the channel initiator is started.



Attention: If there are too many adapters acting on a small set of queues, you might get contention within the queue manager.

Related reference

[ALTER QMGR](#)



Domain Name Server (DNS) task

This topic contains example data for the DNS tasks statistics, and information about how to interpret the data.

The format of the DNS task block is described in structure `qct_dns` in the C programming language header file `thlqual.SCSQC370(CSQDSMFC)`, and in assembler macro `thlqual.SCSQMACS(CSQDQCTA)`.

Example data

Task, date,	Type,	Requests, time	Busy %,	CPU used,	CPU %,	"avg CPU",	"avg ET",	longest,
				Seconds,		uSeconds,	uSeconds,	uSeconds,
0',	DNS',	14002',	0.0',	0.122578,	0.0',	9,	11,	463,
12:56:33.987671								2014/03/18,
Summ,	DNS,	14002,	0.0,	0.122578,	0.0,	9,	11,	463,
12:56:33.987671								2014/03/18,

The channel initiator uses a single DNS task. The example shows that the task processed 14002 requests and on average the request used 9 microseconds of CPU and took 11 microseconds of elapsed time.

The longest DNS request took 463 microseconds elapsed time, and this occurred at 12:56:33 local time.

The fields are calculated from:

- Duration: *qwhs.qwhsdurn*
- Requests : *qctreqn*
- Busy %: *qcteltn* and duration
- CPU used: *qctcptm*
- CPU %: *qctcptm* and duration
- Average CPU: *qctcptm* and *qctreqn*
- Average ET: *qcteltn* and *qctreqn*
- Longest: *qctlgdu*
- Longest at: *qctlgtm*

The DNS task can go out of your enterprise to look up the IP address associated with a name. If the average Elapsed time is significantly more than the average CPU time used, you might have some long requests.

If the value of the longest request time is unacceptable you should work with your network team to investigate why you are having long requests. It might be that you have an invalid name in your connections.

If the DNS task is busy for 25% of the duration, consider investigating the cause further.

The CPU time (*qctcptm*) value includes all CPU consumed by the task, both processing requests and between processing requests. The elapsed time (*qcteltm*) value only includes time while processing requests. This means that the CPU time might be greater than the elapsed time.

Note: There are requests to the DNS task that are not DNS lookups, so you might have the number of requests being greater than zero - but no longest request information.

z/OS SSL tasks

This topic contains example data for the SSL tasks statistics, and information about how to interpret the data.

The format of the SSL task block is described in structure `qct_ssl` in the C programming language header file `thlqual.SCSQC370(CSQDSMFC)`, and in assembler macro `thlqual.SCSQMACS(CSQDQCTA)`.

Example data

Task, date,	Type,	Requests, time	Busy %,	CPU used,	CPU %,	"avg CPU",	"avg ET",	longest,	
				Seconds,		uSeconds,	uSeconds,	uSeconds,	
0,	SSL,	3112,	1.2,	0.248538,	0.3,	80,	362,	8864,	2014/03/18,
12:46:40.237697									
1,	SSL,	3070,	1.2,	0.245433,	0.3,	80,	359,	4714,	2014/03/18,
12:46:18.938022									
2,	SSL,	3170,	1.2,	0.255557,	0.3,	81,	362,	7273,	2014/03/18,
12:46:35.358145									
3,	SSL,	3060,	1.2,	0.246542,	0.3,	81,	365,	13164,	2014/03/18,
12:46:44.514045									
4,	SSL,	3120,	1.3,	0.251927,	0.3,	81,	373,	22438,	2014/03/18,
12:46:22.134123									
Summ,	SSL,	15532,	1.2,	1.247998,	0.3,	80,	364,	22438,	2014/03/18,
12:46:22.134123									

This example data shows that the average request took 364 microseconds. The longest request was for SSL task 4, took 22,438 microseconds, and occurred at 12:46:22.134123 local time.

The fields are calculated from:

- Duration: *qwhs.qwhsdurn*
- Requests : *qctreqn*
- Busy %: *qcteltm* and duration
- CPU used: *qctcptm*
- CPU %: *qctcptm* and duration
- Average CPU: *qctcptm* and *qctreqn*
- Average ET: *qcteltm* and *qctreqn*
- Longest: *qctlsdu* longest at: *qctlstm*

A running channel is associated with an SSL task, in a similar way that a channel is associated with a dispatcher. The SSL tasks can use the cryptographic coprocessors available to the LPAR. So, the elapsed time can include time spent on a coprocessor. You should monitor the average elapsed time throughout the day. If this time increases significantly during peak periods you should work with your z/OS systems programmers, as your coprocessors might be over-used.

If the SSL tasks are busy for a significant proportion of the interval, increasing the number of SSL tasks might help. If the SSL tasks are waiting for external resources such as a coprocessor, increasing the number of SSL tasks has little effect.

You can use the `ALTER QMGR SSLTASKS()` command to change the number of SSL tasks used. Any changes come into effect the next time the channel initiator is started.

The CPU time (*qctcptm*) value includes all CPU consumed by the task, both processing requests and between processing requests. The elapsed time (*qcteltm*) value only includes time while processing requests. This means that the CPU time might be greater than the elapsed time.

Related reference

[ALTER QMGR](#)

Queue data records

Use this topic as a reference for queue (SMF type 115, subtype 216) data records. The statistics are designed to make it easier for you to monitor usage and performance of your queue over time, and give an insight into what happened with your queue during the last SMF interval. This includes all the **DISPLAY QSTATUS** information and information on message flow, expiry, high and low watermarks and more.

The format of the queue statistics data record is described in assembler macro `thlqual.SCSQMACS(CSQDQST)`.

The queue statistics record contains information on the performance of selected queues and includes the following fields:

QQSTID – Control block identifier

The identifier for the queue statistics control block; is always x' D80F '.

QQSTLL – Length of control block

The length of a queue statistics record.

QQSTEYEC – Control block eyecatcher

The eyecatcher used to make identification of the control block easier; is always ' QQST '.

QQSTQNAM – Queue name

The name of the queue.

QQSTFLAG

An array of bits containing the following information about the queue:

QQSTDISP – Queue disposition

This bit identifies whether the queue is of private or shared disposition. If the bit is on, then it is a shared queue.

QQSTPART – Partial record identifier

This bit identifies whether the record is a full or partial record. If the bit is on, then it is a partial record. When this flag is set there was an issue accessing the information on the queue, for example if there is a CF structure failure.

In a shared queue partial record, the accuracy of **qqstdpth**, **qqstmage**, and **qqstuncm** cannot be guaranteed. Therefore, the fields **qqstdpth** and **qqstmage** are populated with x' 00 ' and the **qqstuncm** flag is not set.

In a private queue partial record, the accuracy of **qqstmage** cannot be guaranteed, therefore the field is populated with x' 00 '.

QQSTUNCM – Uncommitted changes pending

This bit indicates whether there are any uncommitted changes (puts and gets) pending for the queue. If the bit is on, there are uncommitted changes.

This is checked and set at the time of the SMF data collection and provides the same result as a [DISPLAY QSTATUS](#) command would if run at the time the SMF record was generated.

If the queue is a shared queue and QQSTPART is set, this bit is always off as there might have been an issue obtaining the correct value.

QQSTPSID – Page set ID

The page set ID where the queue is located, if allocated and a private queue. If the queue is a shared queue, or a private queue that does not have a page set assigned to it, this field is set to -1 (x' FFFF ').

This value is correct at the time the SMF record was generated. It is possible that the page set changed during the SMF interval, in which case, the value reflected in the next SMF record will be the new page set.

QQSTBPID – Buffer pool ID

The buffer pool ID used by the queue, if allocated and a private queue. If the queue is a shared queue, or a private queue that does not have a buffer pool assigned to it, this field is set to -1 (x' FFFF ').

This value is correct at the time the SMF record was generated. It is possible that the buffer pool changed during the SMF interval. In this case, the value reflected in the SMF record is the new buffer pool.

QQSTQSGN – QSG name

The Queue Sharing Group name that the queue manager is a member of, if it is a shared queue. If the queue is a private queue this field is blank.

QQSTCFST – CF Structure name

The coupling facility (CF) structure name the queue uses if it is a shared queue. If the queue is a private queue this field is blank.

QQSTDPTH – Current depth

The depth of the queue at the time the SMF data was captured.

If the queue is a shared queue and QQSTPART is set, this value is always zero as there might have been an issue obtaining the correct value.

QQSTOPCT – Current open for output count

The number of handles that are currently open for output for the queue at the time when the SMF data was captured. For shared queues, the number returned applies only to the queue manager generating the record. The number is not the total for all the queue managers in the queue sharing group.

This is the same as OPPROCS from a DISPLAY QSTATUS command.

This is checked and set at the time of the SMF data collection and provides the same result as a DISPLAY QSTATUS command would, if run at the time the SMF record was generated.

QQSTIPCT – Current open for input count

The number of handles that are currently open for input for the queue at the time when the SMF data was captured. For shared queues, the number returned applies only to the queue manager generating the record. The number is not the total for all the queue managers in the queue sharing group.

This is the same as IPPROCS from a DISPLAY QSTATUS command.

This is checked and set at the time of the SMF data collection and provides the same result as a DISPLAY QSTATUS command would, if run at the time the SMF record was generated.

QQSTMAGE – Oldest message age

The age, in seconds, of the oldest message on the queue.

This is checked and set at the time of the SMF data collection and provides the same result as a DISPLAY QSTATUS command would, if run at the time the SMF record was generated.

If QQSTPART is set, this value is always zero as there might have been an issue obtaining the correct value.

QQSTQTST – Short term QTIME

The interval, in microseconds, between messages being put on the queue and then being destructively read. Value based on the last few messages processed. For shared queues, the values shown are for measurements collected on this queue manager only.

This is the same as the first value in QTIME from a DISPLAY QSTATUS command.

This is checked and set at the time of the SMF data collection and provides the same result as a DISPLAY QSTATUS command would, if run at the time the SMF record was generated.

QQSTQTLT – Long term QTIME

The interval, in microseconds, between messages being put on the queue and then being destructively read. The value is based on a larger sample of the recently processed messages. For shared queues, the values shown are for measurements collected on this queue manager only.

This is the same as the second value in QTIME from a DISPLAY QSTATUS command.

This is checked and set at the time of the SMF data collection and provides the same result as a DISPLAY QSTATUS command would, if run at the time the SMF record was generated.

QQSTLPUT – Last put date/time

The time, in store clock format, at which the last message was put to the queue since the queue manager started. For shared queues, the value shown is for messages put by this queue manager only.

This is the same as LPUTDATE and LPUTTIME from a DISPLAY QSTATUS command.

This is checked and set at the time of the SMF data collection and provides the same result as a `DISPLAY QSTATUS` command would, if run at the time the SMF record was generated.

QQSTLGET – Last get date/time

The time, in store clock format, at which the last message was retrieved from the queue since the queue manager started. For shared queues, the value shown is for messages put by this queue manager only.

A message being browsed does not count as a message being retrieved.

This is the same as `LGETDATE` and `LGETTIME` from a `DISPLAY QSTATUS` command.

This is checked and set at the time of the SMF data collection and provides the same result as a `DISPLAY QSTATUS` command would, if run at the time the SMF record was generated.

QQSTDPHI – Highest depth

The highest depth reached by the queue during the SMF interval.

For shared queues, queue managers only have partial information about the change in depth of the queue over time. The `QQSTDPHI` value is based off this partial information as follows:

- At the start of the interval the value of `QQSTDPHI` is set to zero.
- When an application puts a message to the queue in the interval the queue manager checks the depth of the queue, including the message just being put. If this value is higher than the current value of `QQSTDPHI`, then it is used as the new value of `QQSTDPHI`.
- When SMF data for the queue is collected, the queue manager will check if the current queue depth is higher than `QQSTDPHI`, if so the current queue depth is used as the new value of `QQSTDPHI`.

This approach means that the value of `QQSTDPHI` does not take into account messages put by other queue managers in the queue sharing group, unless those messages contributed to the queue depth at the point where SMF data is collected.

QQSTDPL0 – Lowest depth

The lowest depth reached by the queue during the SMF interval.

For shared queues, queue managers only have partial information about the change in depth of the queue over time. The `QQSTDPL0` value is based off this partial information as follows:

- At the start of the interval the value of `QQSTDPL0` is set to a special value.
- The first time during the interval the queue manager obtains the depth of the queue, `QQSTDPL0` to that value.
- When an application puts a message to the queue in the interval the queue manager checks the depth of the queue, including the message just being put. If this value is lower than the current value of `QQSTDPL0`, then it is used as the new value of `QQSTDPL0`.
- When SMF data for the queue is collected, the queue manager will check if the current queue depth is lower than `QQSTDPL0`, if so the current queue depth is used as the new value of `QQSTDPL0`.

This approach means that the value of `QQSTDPL0` does not take into account messages got by other queue managers in the queue sharing group, unless those messages contributed to the queue depth at the point where SMF data is collected.

QQSTPUTS – MQPUT count

The number of messages put to the queue using `MQPUT` during the SMF interval.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTPUT1 – MQPUT1 count

The number of messages put to the queue using `MQPUT1` during the SMF interval.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTNPPT – Non-persistent MQPUT count

The number of non-persistent messages put to the queue using `MQPUT` during the SMF interval.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTPPT – Persistent MQPUT count

The number of persistent messages put to the queue using MQPUT during the SMF interval.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTNPP1 – Non-persistent MQPUT1 count

The number of non-persistent messages put to the queue using MQPUT1 during the SMF interval.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTPP1 – Persistent MQPUT1 count

The number of persistent messages put to the queue using MQPUT1 during the SMF interval.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTPUTB – MQPUT bytes

The number of bytes of message data, including any message properties, put to the queue using MQPUT during the SMF interval. This does not include message headers in the calculation of the size.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTPT1B – MQPUT1 bytes

The number of bytes of message data, including any message properties, put to the queue using MQPUT1 during the SMF interval. This does not include message headers in the calculation of the size.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTNPPB – Non-persistent MQPUT bytes

The number of bytes of non-persistent message data, including any message properties, put to the queue using MQPUT during the SMF interval. This does not include message headers in the calculation of the size.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTPPB – Persistent MQPUT bytes

The number of bytes of persistent message data, including any message properties, put to the queue using MQPUT during the SMF interval. This does not include message headers in the calculation of the size.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTNP1B – Non-persistent MQPUT1 bytes

The number of bytes of non-persistent message data, including any message properties, put to the queue using MQPUT1 during the SMF interval. This does not include message headers in the calculation of the size.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTP1B – Persistent MQPUT1 bytes

The number of bytes of persistent message data, including any message properties, put to the queue using MQPUT1 during the SMF interval. This does not include message headers in the calculation of the size.

For shared queues, the count only includes messages put through the queue manager that generated the SMF record.

QQSTFLPT – Failed MQPUT count

The number of MQPUT calls targeting the queue, which failed with a completion code of MQCC_FAILED, during the SMF interval.

For shared queues, the count only includes failed puts attempted through the queue manager that generated the SMF record.

QQSTFLP1 – Failed MQPUT1 count

The number of MQPUT1 calls targeting the queue, which failed with a completion code of MQCC_FAILED, during the SMF interval.

For shared queues, the count only includes failed puts attempted through the queue manager that generated the SMF record.

QQSTFPTC – Fast puts to a waiting getter count

The number of MQPUT and MQPUT1 calls targeting the queue, that were fast put to a waiting getter during the SMF interval.

Note: *Put to a waiting getter* is a technique whereby a message might not actually be put onto a queue if there is an application already waiting to get the message. Certain conditions must be satisfied for this to occur, in particular the message must be non-persistent and the putting and getting application must be processing the message outside syncpoint control.

If these conditions are met, then the message is transferred from the putting application's buffer into the getting application's buffer without actually touching the IBM MQ queue. This removes a lot of processing involved in putting the message on the queue and therefore leads to increased throughput and lower CPU costs.

QQSTFPTB – Fast puts to a waiting getter bytes

The number of message and properties bytes from MQPUT and MQPUT1 calls targeting the queue, that were fast put to a waiting getter during the SMF interval. This does not include message headers in the calculation of the size.

QQSTSTRM – Streamed message count

The number of messages that were successfully streamed from the queue during the interval. This is always zero if STREAMQ is not set for the queue.

QQSTMSMI – Minimum message size put

The minimum message size, in bytes, put to the queue during the SMF interval.

This includes message and properties bytes, and does not include message headers such as the MQMD.

QQSTMSMA – Maximum message size put

The maximum message size, in bytes, put to the queue during the SMF interval.

This includes message and properties bytes, and does not include message headers such as the MQMD.

QQSTMSAV – Average message size put

The average message size, in bytes, put to the queue during the SMF interval.

This includes message and properties bytes, and does not include message headers such as the MQMD.

QQSTGETS – Destructive MQGET count

The number of messages got from the queue using destructive MQGET during the SMF interval.

For shared queues, the count only includes messages got through the queue manager that generated the SMF record.

QQSTNPDG – Non-persistent destructive MQGET count

The number of non-persistent messages got from the queue using destructive MQGET during the SMF interval.

For shared queues, the count only includes messages got through the queue manager that generated the SMF record.

QQSTPDG – Persistent destructive MQGET count

The number of persistent messages got from the queue using destructive MQGET during the SMF interval.

For shared queues, the count only includes messages got through the queue manager that generated the SMF record.

QQSTGETB – Destructive MQGET byte count

The number of message and properties bytes got from the queue using destructive MQGET during the SMF interval.

For shared queues, the count only includes messages got through the queue manager that generated the SMF record.

QQSTNPDB – Non-persistent destructive MQGET bytes

The number of non-persistent message and properties bytes got from the queue using destructive MQGET during the SMF interval.

For shared queues, the count only includes messages got through the queue manager that generated the SMF record.

QQSTPDB – Persistent destructive MQGET bytes

The number of persistent message and properties bytes got from the queue using destructive MQGET during the SMF interval.

For shared queues, the count only includes messages got through the queue manager that generated the SMF record.

QQSTBRWS – Non-destructive MQGET count

The number of messages browsed from the queue during the SMF interval.

For shared queues, the count only includes messages browsed through the queue manager that generated the SMF record.

QQSTNPBR – Non-persistent non-destructive MQGET count

The number of non-persistent messages browsed from the queue during the SMF interval.

For shared queues, the count only includes messages browsed through the queue manager that generated the SMF record.

QQSTPBR – Persistent non-destructive MQGET count

The number of persistent messages browsed from the queue during the SMF interval.

For shared queues, the count only includes messages browsed through the queue manager that generated the SMF record.

QQSTBRWB – Non-destructive MQGET bytes

The number of message and properties bytes browsed from the queue during the SMF interval.

For shared queues, the count only includes messages browsed through the queue manager that generated the SMF record.

QQSTNPBB – Non-persistent non-destructive MQGET bytes

The number of non-persistent message and properties bytes browsed from the queue during the SMF interval.

For shared queues, the count only includes messages browsed through the queue manager that generated the SMF record.

QQSTPBB – Persistent non-destructive MQGET bytes

The number of persistent message and properties bytes browsed from the queue during the SMF interval.

For shared queues, the count only includes messages browsed through the queue manager that generated the SMF record.

QQSTFLGT – Failed destructive MQGET count

The number of destructive MQGET calls targeting the queue, that failed with a completion code of MQCC_FAILED, during the SMF interval.

Also included in the count is MQCC_WARNING when accompanied by a return code of MQRC_TRUNCATED_MSG_FAILED. However, not included in this count is any MQGET with a wait that receives MQRC_NO_MSG_AVAILABLE.

For shared queues, the count only includes failed MQGET attempts through the queue manager that generated the SMF record.

QQSTNMAG – Failed destructive MQGET with MQRC_NO_MSG_AVAILABLE count

The number of destructive MQGET calls, without wait, targeting the queue, that failed with both a completion code of MQCC_FAILED and a return code of MQRC_NO_MSG_AVAILABLE, during the SMF interval.

This value is a subset of QQSTFLGT.

For shared queues, the count only includes failed MQGET, without wait, attempts through the queue manager that generated the SMF record.

QQSTTMFB – Failed destructive MQGET with MQRC_TRUNCATED_MSG_FAILED count

The number of destructive MQGET calls targeting the queue, that failed with both a completion code of MQCC_WARNING and a return code of MQRC_TRUNCATED_MSG_FAILED, during the SMF interval.

This value is a subset of QQSTFLGT.

For shared queues, the count only includes failed MQGET attempts through the queue manager that generated the SMF record.

QQSTFLGW – No message available for destructive MQGET with a wait count

The number of times that there is no message available for destructive MQGET calls with a wait, targeting the queue, causing the MQGET to continue waiting, during the SMF interval.

QQSTRDGW – Re-driven destructive MQGET with a wait count

The number of times that destructive MQGET calls with a wait are re-driven to check if there is a message on the queue that matches their criteria, during the SMF interval.

When a new message arrives on the queue, all eligible waiting MQGET calls are woken up to attempt to get the message. Each MQGET with a wait that is woken up to check for a valid message increments this count by one. If any of these MQGET with wait fails to get the message, QQSTFLGW is incremented by one and the MQGET goes back into a waiting state.

QQSTFLBW – No message available for non-destructive MQGET with a wait count

The number of times that there is no message available for non-destructive MQGET calls with a wait, targeting the queue, causing the MQGET to continue waiting, during the SMF interval.

QQSTRDBW – Re-driven non-destructive MQGET with a wait count

The number of times that non-destructive MQGET calls with a wait are re-driven to check if there is a message on the queue that matches their criteria, during the SMF interval.

When a new message arrives on the queue, all eligible waiting MQGET calls are woken up to attempt to browse the message. Each MQGET with a wait that is woken up to check for a valid message increments this count by one. If any of these MQGET with wait fails to browse the message, QQSTFLBW is incremented by one and the MQGET goes back into a waiting state.

QQSTSAGT – Destructive MQGET with MQRC_SIGNAL_REQUEST_ACCEPTED count

The number of destructive MQGET calls targeting the queue, that complete with both a completion code of MQCC_WARNING and a return code of MQRC_SIGNAL_REQUEST_ACCEPTED, during the SMF interval.

QQSTSABR – Non-destructive MQGET with MQRC_SIGNAL_REQUEST_ACCEPTED count

The number of non-destructive MQGET calls targeting the queue, that complete with both a completion code of MQCC_WARNING and a return code of MQRC_SIGNAL_REQUEST_ACCEPTED, during the SMF interval.

QQSTIPHI – High watermark for IPPROC

The highest number of concurrent input handles open on the queue during the SMF interval.

For shared queues, the watermark only includes handles owned through the queue manager that generated the SMF record.

QQSTIPLO – Low watermark for IPPROC

The lowest number of concurrent input handles open on the queue during the SMF interval.

For shared queues, the watermark only includes handles owned through the queue manager that generated the SMF record.

QQSTOPHI – High watermark for OPPROC

The highest number of concurrent output handles open on the queue during the SMF interval.

For shared queues, the watermark only includes handles owned through the queue manager that generated the SMF record.

QQSTOPLO – Low watermark for OPPROC

The lowest number of concurrent output handles open on the queue during the SMF interval.

For shared queues, the watermark only includes handles owned through the queue manager that generated the SMF record.

QQSTOPEN – Successful MQOPEN count

The number of times the queue was successfully opened during the SMF interval. This does not include opens performed as part of an MQPUT1 call.

For shared queues, the count only includes when the queue is opened through the queue manager that generated the SMF record.

QQSTCLOS – MQCLOSE count

The number of times the queue was successfully closed using MQCLOSE, during the SMF interval.

For shared queues, the count only includes when the queue is closed through the queue manager that generated the SMF record.

QQSTINQR – MQINQ count

The number of MQINQ calls that completed with a completion code of MQCC_OK or MQCC_WARNING, during the SMF interval.

QQSTSET – MQSET count

The number of MQSET calls that completed with a completion code of MQCC_OK during the SMF interval.

QQSTEXPR – Expired messages count

The number of expired messages cleared from the queue during the SMF interval.

This includes messages expired by an application issuing an MQGET, by the expired message scanning task or by a REFRESH QMGR TYPE(EXPIRY) command.

QQSTRBPT – Rolled back MQPUT counts

The number of messages that were put to the queue, that have been rolled back off the queue, during the SMF interval.

QQSTRBGT – Rolled back MQGET counts

The number of messages destructively read from the queue, that have been rolled back onto the queue, during the SMF interval.

V 10.0.0 QQSTOCPA – Context propagated from environment on MQPUT

The number of times, in the interval, that trace context was automatically propagated from the application environment to the message on MQPUT.

V 10.0.0 QQSTOCGA – Context propagated to environment on MQGET

The number of times, in the interval, that trace context was automatically propagated from the message to the application environment on destructive MQGET.

V 10.0.0 QQSTOAHG – Context discarded on MQPUT, message already had context

The number of times, in the interval, that trace context from the environment was discarded on MQPUT because the message already had context.

V 10.0.0 QQSTOCPD – Incorrect traceparent on MQPUT

The number of times, in the interval, that an incorrectly formatted **traceparent** property was removed from the message and discarded during MQPUT. Producer spans are not generated.

V 10.0.0 QQSTOCMP – Context discarded on MQPUT, would exceed MAXPROPL

The number of times, in the interval, that trace context from the environment was discarded during MQPUT because adding it to the message would exceed the **MAXPROPL** value for the queue manager.

V 10.0.0 QQSTOCGD – Context discarded on MQGET

The number of times, in the interval, that invalid trace context was discarded from the message during MQGET.

V 10.0.0 QQSTOCPS – tracestate discarded on MQPUT

The number of times, in the interval, that **tracestate** was truncated or discarded during automatic propagation from the environment to the message on MQPUT. IBM MQ supports a maximum of 512 bytes of **tracestate**.

V 10.0.0 QQSTOCGS – tracestate discarded on MQGET

The number of times, in the interval, that **tracestate** was truncated or discarded during automatic propagation from the message to the environment on MQGET. IBM MQ supports a maximum of 512 bytes of **tracestate**.

V 10.0.0 QQSTOSCS – consumer spans scheduled for emission

The number of consumer spans scheduled for emission in the interval.

V 10.0.0 QQSTOSPS – producer spans scheduled for emission

The number of producer spans scheduled for emission in the interval.

V 10.0.0 QQSTOSNS – sampled bit off

The number of times, in the interval, that a message had trace context on MQPUT but sampled bit was off so a producer span was not generated.

V 10.0.0 QQSTOCDA – context discarded on MQGET, application cannot accept properties

The number of times, in the interval, that a message had context that was discarded on MQGET because of **OTELPCTL (AUTO)** processing or because application cannot accept properties.

V 10.0.0 QQSTOSPL – long running transaction limit reached on MQPUT

The number of messages, in the interval, that had trace context, but a producer span was not emitted due to the span limit for long running transactions being reached.

V 10.0.0 QQSTOSCL – long running transaction limit reached on MQGET

The number of messages, in the interval, that were got from the queue that had trace context, but no consumer span was emitted due to the span limit for long running transactions being reached.

V 10.0.0 QQSTOSFS – span storage limit reached

The number of spans, in the interval, that could not be emitted because span storage limit was reached.

Interpreting IBM MQ for z/OS accounting data

IBM MQ for z/OS accounting data is written as SMF type 116 records. Use this topic as a reference to the different types of accounting data records.

IBM MQ accounting information can be collected for the following subtypes:

0

Message manager accounting records (how much processor time was spent processing IBM MQ API calls and the number of MQPUT and MQGET calls). This information is produced when a named task disconnects from IBM MQ, and so the information contained within the record might cover many hours.

1

Accounting data for each task, at thread and queue level.

2

Additional queue-level accounting data (if the task used more queues than could fit in the subtype 1 record).

10

Accounting data for channels.

Note: Accounting information for specific channels can be enabled or suppressed by the **STATCHL** channel attribute, and the **STATACLS** queue manager attribute.

Note that:

- Subtype 0 records are produced with accounting trace class 1.
- Subtype 1 and 2 records are produced with accounting trace class 3.
- Subtype 10 records are produced with accounting trace class 4.

Layout of an SMF type 116 record

Use this topic as a reference to the format of an SMF type record.

The standard layout for SMF records involves three parts:

SMF header

Provides format, identification, and time and date information about the record itself.

Self-defining section

Defines the location and size of the individual data records within the SMF record.

Data records

The actual data from IBM MQ that you want to analyze.

For more information about SMF record formats, see [z/OS MVS System Management Facilities \(SMF\)](#).

The SMF header

Table 46 on page 386 shows the format of SMF record header (SM116).

Table 46. SMF record header description						
Offset : Dec	Offset: Hex	Type	Len	Name	Description	Example
0	0	Structure	28	SM116	SMF record header.	
0	0	Integer	2	SM116LEN	SMF record length.	01A4
2	2		2		Reserved.	
4	4	Integer	1	SM116FLG	System indicator.	5E
5	5	Integer	1	SM116RTY	Record type. The SMF record type, for IBM MQ accounting records this is always 116 (X'74').	74
6	6	Integer	4	SM116TME	Time when SMF moved record.	00356124
10	A	Integer	4	SM116DTE	Date when SMF moved record.	0100223F
14	E	Character	4	SM116SID	z/OS subsystem ID. Defines the z/OS subsystem on which the records were collected.	D4E5F4F1 (MV41)
18	12	Character	4	SM116SSI	IBM MQ subsystem ID.	D4D8F0F7 (MQ07)
22	16	Integer	2	SM116STF	Record subtype.	0000
24	18	Character	3	SM116REL ²	Internal version number.	F9F4F6 (946)
27	1B		1		Reserved.	
28	1C	Character	0	SM116END	End of SMF header and start of self-defining section.	

Notes:

1. The (hexadecimal) values in the right-hand column relate to Figure 26 on page 388.
2.  If SMF116REL is 946 or greater, locate the QWHX header and take the IBM MQ version from the 6-character QWHXREL field. See [Common IBM MQ SMF header](#) and [Layout of channel initiator SMF type 116 records](#).

Self-defining sections

A self-defining section of an SMF record tells you where to find an accounting record, how long it is, and how many times that type of record is repeated (with different values). The self-defining sections follow the header, at a fixed offset from the start of the SMF record.

Each self-defining section points to accounting related data. [Table 47 on page 387](#) summarizes the offsets from the start of the SMF record header.

<i>Table 47. Offsets to self-defining sections</i>				
Record subtype (SMF116STF)	Source of accounting data	Offset of self-defining section		See...
		Dec	Hex	
All	Common header	28	X'1C'	“Common IBM MQ SMF header” on page 389
0	Message manager	44	X'2C'	“Message manager data records” on page 391
1	Thread identification record	36	X'24'	“Thread-level and queue-level data records” on page 392
1	Thread-level accounting	44	X'2C'	“Thread-level and queue-level data records” on page 392
1	Queue-level accounting	52	X'34'	“Thread-level and queue-level data records” on page 392 . This section is present only if the WTASWQCT field in the task-related information (WTAS) structure is non-zero.
2	Thread identification record	36	X'24'	“Thread-level and queue-level data records” on page 392
2	Queue-level accounting	44	X'2C'	“Thread-level and queue-level data records” on page 392
10	Channel accounting			“Channel accounting data records” on page 395

Note: Other self-defining sections refer to data for IBM use only.

Each self-defining section is two fullwords long and has this format:

```
sssssssl111nnnn
```

where:

ssssssss

Fullword containing the offset from start of the SMF record.

llll

Halfword giving the length of this data record.

nnnn

Halfword giving the number of data records in this SMF record.

Figure 26 on page 388 shows an example of part of an SMF type 116 record. The numbers in the left-hand column represent the offset, in hexadecimal, from the start of the record. Each line corresponds to sixteen bytes of data, where each byte is two hexadecimal characters, for example 0C. The characters in the right-hand column represent the printable characters for each byte. Non-printable characters are shown by a period (.) character.

```
000000 01A40000 5E740035 61240100 223FD4E5 *...;.../....MV*
000000 F4F1D4D8 F0F70000 F6F0F000 00000134 *41MQ07..600....*
000000 00700001 00000054 00B00001 00000104 *.....*
000000 00300001 00000000 00000000 00000000 *.....*
000000 00000000 00000000 00000000 00000000 *.....*
```

Figure 26. Part of an SMF record 116 showing the header and self-defining sections

The self-defining section for the type of message manager accounting data is located at offset X'2C' from the start of the SMF record and contains this information:

- The offset of the message manager accounting data is located X'00000104' bytes from the start of the SMF record.
- This message manager record is X'0030' bytes long.
- There is one record (X'0001').

Note: Always use offsets in the self-defining sections to locate the accounting records.

Processing type 116 SMF records

Use this topic as a reference to the format of the processing type accounting record.

Any accounting data you collect from SMF must be processed to extract useful information. When you process the data, verify that the records are from IBM MQ and that they are the records you are expecting.

Validate the value of the following fields:

- SM116RTY, the SMF record number = X'74' (116)
- SM116STF, the record subtype, must be 0000, 0001, 0002, or 0010

Reading from the active SMF data sets (or SMF logstreams) is not supported. You must use the SMF program IFASMFDP (or IFASMFDP if logstreams are being used) to dump SMF records to a sequential data set so that they can be processed. For more information see [“Using System Management Facility” on page 343](#).

Details of the structures and fields can be found in IBM MQ SupportPac [MP1B](#).

There is a C sample program called CSQ4SMFD which prints the contents of SMF type 115 and 116 records from the sequential data set. The program is provided as source in thlqual.SCSQC37S and in executable format in thlqual.SCSQLOAD. Sample JCL is provided in thlqual.SCSQPROC(CSQ4SMFJ).

You need to update the SMFIN DD card with the name of the SMF data set. Use the z/OS command '/D SMF' to show the name of the data set, and you need to update the DUMPOUT DD card with the name for the output data set.

You also need to specify the START and END times that you require.

The following sample JCL extracts SMF records from SMF data sets and dumps them to the SMFOUT data set:

```
//SMFDUMP EXEC PGM=IFASMFDP,REGION=0M
//SYSPRINT DD SYSOUT=
```

```
//SMFIN DD DSN=xxxxxx.MANA,DISP=SHR
//SMFOUT DD DSN=xxxxxx.SMFOUT,SPACE=(CYL,(1,1)),DISP=(NEW,CATLG)
//SYSIN DD *
INDD(SMFIN,OPTIONS(DUMP))
OUTDD(SMFOUT,TYPE(116))
OUTDD(SMFOUT,TYPE(115))
START(1159) END(1210)
/*
```

The following sample JCL extracts SMF records from the SMF log stream named in LSNAME and dumps them to the SMFOUT data set:

```
//SMFDUMP EXEC PGM=IFASMF DL,REGION=0M
//SYSPRINT DD SYSOUT=*
//SMFOUT DD DSN=xxxxxx.SMFOUT,SPACE=(CYL,(1,1)),DISP=(NEW,CATLG)
//SYSIN DD *
LSNAME(IFASMF.MQ,OPTIONS(DUMP))
OUTDD(SMFOUT,TYPE(116))
OUTDD(SMFOUT,TYPE(115))
START(1159) END(1210)
/*
```

Common IBM MQ SMF header

Use this topic as a reference to the common IBM MQ SMF header type accounting record.

The format of this record is described in Table 48 on page 389 and in assembler macros `thlqual.SCSQMACS(CSQDQWHS)` and `thlqual.SCSQMACS(CSQDQWHC)`, and C header file `thlqual.SCSQC370(CSQDSMFC)`. The field names in C are all in lowercase, for example `qwhs`, `qwhsnsda`.

The QWHS data includes the subsystem name. For subtype 1 records, it also shows whether there are queue-level accounting records present. If the QWHSNSDA field is 3 or less, there are not, and the corresponding self-defining section (at offset X'34') is not set.

The QWHC data gives you information about the user (for example, the user ID (QWHCAID) and the type of application (QWHCATYP)). The QWHC section is completed only for subtype 0 records. The equivalent information is present in the thread identification record for subtype 1 and 2 records.

Table 48. Structure of the common IBM MQ SMF header record QWHS

Offset: Dec	Offset : Hex	Type	Length	Name	Description
0	0	Structure	128	QWHS	
0	0		6		Reserved
6	6	Character	1	QWHSNSDA	Number of self defining sections in the SMF records
7	7		5		Reserved
12	C	Character	4	QWHSSSID	Subsystem name
16	10		24		Reserved
40	28	Character	8	QWHCAID	User ID associated with the z/OS job
48	30	Character	12	QWHCCV	Thread cross-reference
60	3C	Character	8	QWHCCN	Connection name
68	44		8		Reserved
76	4C	Character	8	QWHCOPID	User ID associated with the transaction
84	54	Integer	4	QWHCATYP	Type of connecting system (1=CICS, 2=Batch or TSO, 3=IMS control region, 4=IMS MPP or BMP, 5=Command server, 6=Channel initiator, 7=RRS Batch)

Table 48. Structure of the common IBM MQ SMF header record QWHS (continued)

Offset: Dec	Offset : Hex	Type	Length	Name	Description
88	58	Character	22	QWHCTOKN	Accounting token set to the z/OS accounting information for the user
110	6E	Character	16	QWHCNID	Network identifier
126	7E		2		Reserved

QWHX header

► V1000 - V1006

The QWHX header contains the queue manager release in the form *VVRRMM*, and the name of the queue sharing group that the queue manager is a member of. The QWHX is described in assembler macro `thlqual.SCSQMACS(CSQDQWHX)`.

The QWHX header directly follows the QWHS/QWHC headers described in Table 1, starting at offset 128 (X'80') from the start of the QWHS header.

► z/OS Combining CICS and IBM MQ performance data

Use this topic as a reference to the combination of IBM MQ and CICS performance data.

The common IBM MQ SMF header type accounting record section, QWHCTOKN, is used to correlate CICS type 110 SMF records with IBM MQ type 116 SMF records.

CICS generates an LU6.2 unit-of-work token, for each CICS task. The token is used to generate an accounting token that is written to QWHCTOKN in the correlation header of subtype zero records.

Details are also written to the WTIDACCT section in subtype 1 and 2 records. The accounting token enables correlation between CICS and IBM MQ performance data for a transaction.

► z/OS Thread cross-reference data

Use this topic as a reference to the format of the thread cross-reference type accounting record.

The interpretation of the data in the thread cross-reference (QWHCCV) field varies. This depends on what the data relates to:

- CICS connections (QWHCATYP=1) - see [Table 49 on page 390](#)
- IMS connections (QWHCATYP=3 or 4) - see [Table 50 on page 391](#)
- Batch connections (QWHCATYP=2 or 7) - this field consists of binary zeros
- Others - no meaningful data

Table 49. Structure of the thread cross-reference for a CICS system

Offset: Dec	Offset: Hex	Type	Length	Description
48	30	Character	4	CICS thread number.
52	34	Character	4	CICS transaction name.
56	38	Integer	4	CICS task number.

Some entries contain blank characters. These apply to the task, rather than to a specific transaction.

Table 50. Structure of the thread cross-reference for an IMS system

Offset: Dec	Offset: Hex	Type	Length	Description
48	30	Character	4	IMS partition specification table (PST) region identifier.
52	34	Character	8	IMS program specification block (PSB) name.

Message manager data records

Use this topic as a reference to the format of the message manager accounting records.

The message manager is the component of IBM MQ that processes all API requests. The format of the message manager accounting records is described in assembler macro `thlqual.SCSQMACS(CSQDQMAC)`.

The QMAC data gives you information about the processor time spent processing IBM MQ calls, and counts of the number of MQPUT and MQGET requests for messages of different sizes.

Note: A single IMS application might write two SMF records. In this case, add the figures from both records to provide the correct totals for the IMS application.

Records containing zero processor time

Records are sometimes produced that contain zero processor time in the QMACCPUT field. These records occur when long running tasks identified to IBM MQ either terminate or are prompted to output accounting records by accounting trace being stopped. Such tasks exist in the CICS adapter and in the channel initiator (for distributed queuing). The number of these tasks with zero processor time depends upon how much activity there has been in the system:

- For the CICS adapter, this can result in up to nine records with zero processor time.
- For the channel initiator, the number of records with zero processor time can be up to the sum of `Adapters + Dispatchers + 6`, as defined in the queue manager attributes.

These records reflect the amount of work done under the task, and can be ignored.

Sample subtype zero accounting record

Use this topic as a reference to the format of the subtype zero accounting records.

Figure 27 on page 392 shows a type 116, subtype zero SMF record. In this figure, the SMF record header and the QMAC accounting data record are underlined. The self-defining sections are in bold.

```

000000 01A40000 5E740035 61240100 223FD4E5 *. . . . ; . . . / . . . . MV *
000010 F4F1D4D8 F0F70000 F6F0F000 00000134 * 41MQ07 . . 600 . . . . *
000020 00700001 00000054 00B00001 00000104 * . . . . . . . . . . *
000030 00300001 00000000 00000000 00000000 * . . . . . . . . . . *
000040 00000000 00000000 00000000 00000000 * . . . . . . . . . . *
000050 00000000 B478AB43 9C6C2280 B478AB47 * . . . . . . % . . . . *
000060 9DB47E02 00000000 04C0F631 00000001 * . . = . . { 6 . . . . *
000070 9880E72D 00000000 014D9540 00000000 * . . X . . . . ( . . . . *
000080 08480C80 00000010 40404040 40404040 * . . . . . . . . . . *
000090 00000000 00000000 00000051 00000000 * . . . . . . . . . . *
0000A0 00000000 00000000 00000000 00000000 * . . . . . . . . . . *
0000B0 00000000 00000000 00000000 00000000 * . . . . . . . . . . *
0000C0 00000000 00000000 00000000 00000000 * . . . . . . . . . . *
0000D0 00000000 00000000 00000000 00000000 * . . . . . . . . . . *
0000E0 00000000 00000000 00000000 00000000 * . . . . . . . . . . *
0000F0 00000000 00000000 00000000 00000000 * . . . . . . . . . . *
000100 00000000 D4140030 D8D4C1C3 00000000 * . . . . M . . . QMAC . . . *
000110 689C738D 00000050 00000050 00000050 * . . . . . & . . . . & *
000120 0000000A 00000000 00000000 00000000 * . . . . . . . . . . *
000130 00000000 0024011A 00030710 02DAACF0 * . . . . . . . . . . 0 *

```

Figure 27. Example SMF type 116, subtype zero record

Thread-level and queue-level data records

Use this topic as a reference to the format of the thread-level and queue-level accounting records.

Thread level accounting records are collected for each task using IBM MQ. In addition, queue-level accounting records are gathered about each queue that the task opens. A queue-level accounting record is written for each queue that the task has used since the thread-level accounting record was last written.

If the task uses a queue that is configured with a streaming queue, there is no queue-level accounting record for the streaming queue. Instead, the accounting record for the original queue accumulates data for the data points that would have been associated with the streaming queue.

The only exception to this is that the PUTN/PUT1N value shows the number of MQPUT/MQPUT1 requests made by the application, and excludes the extra MQPUT requests made to the streaming queue.

So, for example, if an application issues a single MQPUT request, the:

- PUTN value is 1
- Elapsed time (PUTET) and CPU time (PUTCT) for the MQPUT include the time taken to put to both the primary and streaming queue
- Number of page set requests (PUTPSN) includes those for both the primary and secondary queue, and so on

For each task, data is written to SMF when the task finishes.

From IBM MQ 9.3.0 onwards, for long running tasks, data is also written at the interval specified by either the **ACCTIME**, or **STATIME** parameter of the CSQ6SYSP system parameter macro, or by the system SMF statistics broadcast, provided that the task was running the previous time data was gathered.

Thread-level and queue-level accounting records are produced if you specify class 3 when you start the accounting trace. For example, use the following command:

```
START TRACE(ACCTG) DEST(SMF) CLASS(3)
```

The thread level accounting information is written to an SMF type 116, subtype 1 record, and is followed by queue-level records. If the task opened many queues, further queue information is written to one or more SMF type 116 subtype 2 records. A thread identification control block is included in each subtype 1 and 2 record to enable you to relate each record to the correct task. Typically, the maximum number of queue-level records in each SMF record is about 45.

The format of the thread-level accounting record is described in assembler macro `thlqual.SCSQMACS(CSQDWTAS)`. The format of the queue-level accounting record is described in

assembler macro `thlqual.SCSQMACS(CSQDWQ)`. The format of the thread identification record is described in assembler macro `thlqual.SCSQMACS(CSQDWTID)`. All these records are also described in C header file `thlqual.SCSQC370(CSQDSMFC)`. The field names in C are all in lowercase, for example `wtas`, `wtasshex`.

 From IBM MQ 10.0.0, if an outbound channel has a **CONNAME** greater than 48 characters, **WTIDCHLC** will contain that **CONNAME** truncated to 48 characters.

Meaning of the channel names

Use this topic as a reference to the meaning of channel names.

The channel name in the WTID is constructed as shown in the following example. In this example a sender channel exists from queue manager QM1 to queue manager QM2.

The meaning of channel names are described in the following table.

Table 51. Meaning of channel names		
Field name	Meaning	Example
For queue manager QM1 the sender channel has the following fields set:		
WTIDCCN	The job name	QM1CHIN
WTIDCHL	The channel name	QM1.QM2
WTIDCHLC	This is defined in the CONNAME of the channel	WINMVS2B(2162)
For queue manager QM2 the receiver channel has the following fields set:		
WTIDCCN	The job name	QM2CHIN
WTIDCHL	The channel name	QM1.QM2
WTIDCHLC	Where the channel came from	9.20.101.14

Sample subtype 1 and subtype 2 records

Use this topic as a reference to the format of the subtype 1 and subtype 2 accounting records.

Figure 28 on page 393 and Figure 29 on page 394 show examples of SMF type 116, subtype 1 and subtype 2 records. These two accounting records were created for a batch job that opened 80 queues. Because many queues were opened, a subtype 2 record was required to contain all the information produced.

000000	703C0000	5E74002D	983B0100	229FD4E5	*....;.....MV*
000010	F4F1D4D8	F0F70001	F6F0F000	00006FCC	*41MQ07..600...?.*
000020	00700001	0000003C	00D00001	<i>0000010C</i>	*.....}.....*
000030	02C00001	000003CC	02400030	F70000D0	*.i.....7..{*
000040	E6E3C9C4	00000000	00000000	00000040	*WTID.....*
.					
.					
000100	00000000	00000000	7F4A4BB8	F70102C0	*....."...7..{*
000110	E6E3C1E2	B4802373	0BF07885	7F4AE718	*WTAS.....0.."X.*

Figure 28. Example SMF type 116, subtype 1 record

The first self-defining section starts at X'24' and is **bold** in the example; X'0000003C' is the offset to the WTID data record, X'00D0' is the length of the WTID record, and X'0001' is the number of WTID records.

The second self-defining section starts at X'2C' and is in *italic*; X'0000010C' is the offset to the WTAS data record, X'02C0' is the length of the WTAS record, and X'0001' is the number of WTAS records.

The third self-defining section starts at X'34' and is **bold** in the example; X'000003CC' is the offset to the first WQST data record, X'0240' is the length of the WQST record, and X'0030' is the number of WQST records.

Figure 29 on page 394 shows an example of an SMF type 116, subtype 2 record.

```

000000 49740000 5E74002D 983B0100 229FD4E5 *....;.....MV*
000010 F4F1D4D8 F0F70002 F6F0F000 00004904 *41MQ07..600....*
000020 00700001 00000034 00D00001 00000104 *.....}.....*
000030 02400020 F70000D0 E6E3C9C4 00000002 *. ..7..}WTID....*
.
.
.
000100 7F4A4BB8 F7020240 E6D8E2E3 00000001 *"...7.. WQST....*

```

Figure 29. Example SMF type 116, subtype 2 record

The first self-defining section starts at X'24' and is **bold** in the example; X'00000034' is the offset to the WTID data record, X'00D0' is the length of the WTID record, and X'0001' is the number of WTID records.

The second self-defining section starts at X'2C' and is in *italic*; X'00000104' is the offset to the first WQST data record, X'0240' is the length of the WQST record, and X'0020' is the number of WQST records.

Figure 30 on page 394 shows an example of an SMF type 116, subtype 1 record where no queues have been opened and there are consequently no self-defining sections for WQST records..

```

000000          5E740039 4E9B0104 344FD4E5 *.....|MV*
000010 F4F1D4D8 F0F70001 F6F0F000 000003DC *41MQ07..600....*
000020 00800001 00000034 00D00001 00000104 *.....*
000030 02D80001 F70000D0 E6E3C9C4 00000002 *.Q..7..WTID....*
000040 C1F8C5C1 C4C5D740 C1F8C5C1 C4C54040 *A8EADEP A8EADE *
000050 40404040 40404040 00000000 00000000 *.....*
000060 40404040 40404040 4040          *          *

```

Figure 30. Example SMF type 116, subtype 1 record with no WQST data records

The first self-defining section starts at X'24' and is **bold** in the example; X'00000034' is the offset to the WTID data record, X'00D0' is the length of the WTID record, and X'0001' is the number of WTID records.

The second self-defining section starts at X'2C' and is in *italic*; X'0000010C' is the offset to the WTAS data record, X'02D8' is the length of the WTAS record, and X'0001' is the number of WTAS records.

There is no self-defining section describing a WQST data record, equivalent to the third self-defining section in Figure 28 on page 393.

Layout of channel initiator SMF type 116 records

The layout of channel accounting data (SMF type 116, subtype 10) records is described in this topic.

Self-defining section

The self-defining section for the channel accounting data follows the standard SMF header. It is structured in the standard triplet format. The format of the triplets is described in structure qws5 in the C programming language header file thlqual.SCSQC370 (CSQDSMFC), and in assembler macro thlqual.SCSQMACS (CSQDQWS5).

Table 52 on page 395 shows the format of the self-defining section.

Table 52. Structure of the channel accounting self-defining section

Offset: Dec	Offset : Hex	Type	Length	Name	Description
0	0	Integer	4	QWS50PSO	Offset from the start of the SMF record to the first instrumentation standard header (QWHS)
4	4	Integer	2	QWS50PSL	Length of the QWHS
6	6	Integer	2	QWS50PSN	Number of instances of QWHS
8	8	Integer	4	QWS50R1O	Offset from the start of the SMF record to the first channel accounting data record (QCST)
12	C	Integer	2	QWS50R1L	Length of the QCST
14	E	Integer	2	QWS50R1N	Number of instances of QCST

Instrumentation standard header (QWHS)

The format of the QWHS is described in structure `qwhs` in the C programming language header file `thlqual.SCSQC370(CSQDSMFC)`, and in assembler macro `thlqual.SCSQMACS(CSQDQWHS)`. It contains the following key fields that are relevant to channel initiator SMF 116 records:

Table 53. Key fields in the QWHS

Name	Length	Description
QWHSNDA	1 byte	Number of self-defining sections
QWHSSSID	4 bytes	Subsystem name
QWHSSMFC	1 bit	Indicates whether there are multiple SMF records containing information for this interval. If this bit is on, information for this interval is continued in further SMF records. If this bit is off, this is the last or only record.
QWHSTIME	8 bytes	Local time of the start of the interval in STCK format
QWHSDURN	8 bytes	Duration from the start of the interval to the end of the interval in STCK format
QWHSSTCK	8 bytes	End of the interval in UTC in STCK format

QWHX header

V1000 V1000

The QWHX header contains the queue manager release in the form `VVRRMM`, and the name of the queue sharing group that the queue manager is a member of. The QWHX is described in assembler macro `thlqual.SCSQMACS(CSQDQWHX)`.

The QWHX header directly follows the QWHS header described in Table 2, starting at offset 54 (X'34') from the start of the QWHS header.

Channel accounting data records

Use this topic as a reference for channel accounting data records.

The format of the channel accounting data record is described in assembler macro `thlqual.SCSQMACS(CSQDQCST)`. The format is also described in the C programming language header

file `th1qua1.SCSQC370(CSQDSMFC)`. Note that the field names in C are all in lowercase, for example, `qcst`.

The channel accounting data gives you information about the status and statistics of each channel instance, including:

- Average network time (`qcstntav`)
- Average time on exit (`qcstetav`)
- Channel batch data limit (`qcstcbdl`)
- Channel batch interval (`qcstcbi`)
- Channel batch size (`qcstcbz`)
- Channel dispatcher number (`qcstdspn`)
- Channel disposition (`qcstchdp`)
- Channel name (`qcstchnm`)
- Channel state (`qcstchst`)
- Channel started time (`qcststrt`)
- Channel status collected time (`qcstcltm`)
- Channel stopped time (`qcstludt`)
- Channel type (`qcstchty`)
- Common name (CN) from SSLCERTI (`qcstslcn`)
- Compression rate (`qcstcpa`)
- Connection name (`qcstcnm`) From IBM MQ 10.0.0, connection names greater than 48 characters will be truncated to 45 characters and have a suffix of ... added at the end.
- Current shared conversations (`qcstcscv`)
- DNS resolution time (`qcstdnrt`)
- Effective value of STATCHL parameter (`qcststcl`)
- Last message time (`qcstlmst`)
- Maximum network time (`qcstntmx`)
- Maximum time on exit (`qcstetmx`)
- Minimum network time (`qcstntmn`)
- Minimum time on exit (`qcstetmn`)
- Name of the remote queue manager or application (`qcstrqmn`)
- Number of batches (`qcstbatc`)
- Number of bytes for message data (`qcstnbyt`)
- Number of bytes for persistent message data (`qcstnpby`)
- Number of bytes received for both message data and control information (`qcstbyrc`)
- Number of bytes sent for both message data and control information (`qcstbyst`)
- Number of full batches (`qcstfuba`)
- Number of messages, or number of MQI calls (`qcstnmsg`)
- Number of persistent messages (`qcstnpg`)
- Number of put retries (`qcstptrc`)
- Number of transmission queue becoming empty (`qcstqetc`)
- Number of transmission buffers received (**qcstbfrc**)
- Number of transmission buffers sent (`qcstbfst`)
- Remote partner product id (`qcstrpid`)
- Remote partner version (`qcstrver`)

- Serial number from SSLPEER (*qcstslsn*)
- SSL CipherSpec (zero means TLS not used) (*qcstslcs*)
- The date and time of maximum network time (*qcstntdt*)
- The date and time of maximum time on exit (*qcstetdt*)

You can use this information to see the throughput of a channel, if the actual batches are approaching the limit, the latency of the network, information about the remote end, performance of user exit, and so on.

Note:

- Both the fields *qcstetmn* (Minimum time on exit) and *qcstntmn* (Minimum network time) are initialized to the hexadecimal value of 8FFFFFFF when unused.
-  The field *qcstrpid* (Remote partner product id) gives the product identifier of the IBM MQ code running at the remote end of the channel. (For possible values, see [MQCMD_INQUIRE_CHANNEL_STATUS \(Inquire Channel Status\) Response](#).)
-  The field *qcstrver* (Remote partner version) gives the version of the IBM MQ code running at the remote end of the channel, for example, 09040300 for version 9.4.3.0.

Here is an example of the channel accounting data which has been formatted with IBM MQ SupportPac MP1B.

The fields available are based on the display channel status command (DIS CHS) and channel statistics by IBM MQ on platforms except z/OS, with some additional fields.

The data and time of the start and end of the record in local time, and the duration

SMF interval start	2014/03/26,02:30:00
SMF interval end	2014/03/26,02:45:00
SMF interval duration	899.997759 seconds

Information about the channel

Connection name	9.20.4.159
Channel disp	PRIVATE
Channel type	RECEIVER
Channel status	CLOSING
Channel STATCHL	HIGH

Start date & time	2014/03/26,02:44:58
Channel status collect time	2014/03/26,02:45:00
Last status changed	1900/01/01,00:00:00
Last msg time	2014/03/26,02:44:59

Batch size	50
Messages/batch	3.3
Number of messages	1,102
Number of persistent messages	1,102
Number of batches	335
Number of full batches	0
Number of partial batches	335
Buffers sent	337
Buffers received	1,272
Message data	5,038,344 4 MB
Persistent message data	5,038,344 4 MB
Non persistent message data	0 0 B
Total bytes sent	9,852 9 KB
Total bytes received	5,043,520 4 MB
Bytes received/Batch	15,055 14 KB
Bytes sent/Batch	29 29 B
Batches/Second	1
Bytes received/message	4,576 4 KB
Bytes sent/message	8 8 B
Bytes received/second	28,019 27 KB/sec
Bytes sent/second	54 54 B/sec
Compression rate	0

The name of the queue manager at the remote end of the connection

Remote qmgr/app	MQPH
Put retry count	0

Tuning your IBM MQ network

Use the tuning tips in this section to help improve the performance of your queue manager network.

Tuning client and server connection channels

The default setting for **SHARECNV** is 10, which allows up to 10 client conversations for each channel instance. However, using a different number of shared conversations can be better for performance. If you do not need shared conversations, or are using a distributed server, set **SHARECNV** to 1. If you have existing client applications that do not run correctly when you set **SHARECNV** to 1 or greater, set **SHARECNV** to 0.

About this task

For some configurations, using shared conversations brings significant benefits. However, for distributed servers, processing messages on channels that use the default configuration of 10 shared conversations is on average 15% slower than on channels that do not use shared conversations. On an MQI channel instance that is sharing conversations, all of the conversations on a socket are received by the same thread. If the conversations sharing a socket are all busy, the conversational threads contend with one another to use the receiving thread. The contention causes delays, and in this situation using a smaller number of shared conversations is better.

You use the **SHARECNV** parameter to specify the maximum number of conversations to be shared over a particular TCP/IP client channel instance. For details of all possible values, see [Supported IBM MQ client: Default behavior of client-connection and server-connection channels](#).

If you set **SHARECNV** to 1 or greater, you enable the following performance enhancements:

- Bi-directional heartbeats
- Administrator stop-quiesce
- Read-ahead
- Asynchronous-consume by client applications

If you do not need shared conversations, these two settings give best performance:

- **SHARECNV(1)**.
- **SHARECNV(0)**.

Notes:

- If the client-connection **SHARECNV** value does not match the server-connection **SHARECNV** value, then the lowest value is used.
- When applications are linked or compiled against a non-reentrant library, the [CURSHCNV\(0\)](#) value is negotiated even if a higher value is set in [CLNTCONN](#) and [SVRCONN](#).

To optimize performance for a given channel instance, complete any of the following steps.

Procedure

- Monitor channels that use the default **SHARECNV** value of 10.

The default setting of **SHARECNV(10)** works well in many scenarios, but might not be the optimum setting for a given channel instance. For example, for distributed servers, processing messages on channels that use this setting is on average 15% slower than on channels that do not use shared conversations.

To ensure that the default setting is appropriate for a given channel instance, monitor how the channel performs with this setting.

- Set a **SHARECNV** value of 2 or more.

You can set SHARECNV (2) to SHARECNV (999999999). To ensure that the setting you choose is appropriate for a given channel instance, monitor how the channel performs with the new setting.

- Set a SHARECNV value of 1.

If you do not need shared conversations, use this setting whenever possible. It eliminates contention to use the receiving thread, and your client applications can take advantage of the performance enhancements described in the "about this task" section.

With this setting, distributed server performance is significantly improved. The performance improvements apply to client applications that issue non read ahead synchronous get wait calls; for example C client MQGET wait calls. When these client applications are connected, the distributed server uses less threads and less memory and the throughput is increased.

If a server has clients connected to it that are sharing conversations over a socket, and you decrease the shared conversations setting from SHARECNV (10) to SHARECNV (1), this has the following effects:

- Increased socket usage on the server.
- Increased channel instances on the server.

In this case, you might also choose to increase the settings for **MaxChannels** and **MaxActiveChannels**.

Note: You can also set the MQCONN option, MQCNO_NO_CONV_SHARING and connect the application to a channel with **SHARECNV** set to a value greater than 1. The result is the same as connecting the application to a channel with **SHARECNV** set to 1.

- Set a SHARECNV value of 0.

The channel instance behaves exactly as if it was an IBM WebSphere MQ 6.0 server or client connection channel. You do not get shared conversations, or the performance enhancements that are available when you set **SHARECNV** to 1 or greater. Use a value of 0 only if you have existing client applications that do not run correctly when you set **SHARECNV** to 1 or greater.

Related concepts

[Supported IBM MQ client: Default behavior of client-connection and server-connection channels](#)

Tuning distributed publish/subscribe networks

Use the tuning tips in this section to help improve the performance of your IBM MQ distributed publish/subscribe clusters and hierarchies.

Related concepts

[“Monitoring clusters” on page 333](#)

Within a cluster you can monitor application messages, control messages, and logs. There are special monitoring considerations when the cluster load balances between two or more instances of a queue.

Direct routed publish/subscribe cluster performance

In direct routed publish/subscribe clusters, information such as clustered topics and proxy subscriptions is pushed to all members of the cluster, irrespective of whether all cluster queue managers are actively participating in publish/subscribe messaging. This process can create a significant additional load on the system. To reduce the effect of cluster management on performance you can perform updates at off-peak times, define a much smaller subset of queue managers involved in publish/subscribe and make that an "overlapping" cluster, or switch to using topic host routing.

There are two sources of workload on a queue manager in a publish/subscribe cluster:

- Directly handling messages for application programs.
- Handling messages and channels needed to manage the cluster.

In a typical point-to-point cluster, the cluster system workload is largely limited to information explicitly requested by members of the cluster as required. Therefore in anything other than a very large point-to-point cluster, for example one which contains thousands of queue managers, you can largely discount the performance effect of managing the cluster. However, in a direct routed publish/subscribe cluster,

information such as clustered topics, queue manager membership and proxy subscriptions is pushed to all members of the cluster, irrespective of whether all cluster queue managers are actively participating in publish/subscribe messaging. This can create a significant additional load on the system. Therefore you need to consider the effect of cluster management on queue manager performance, both in its timing, and its size.

Performance characteristics of direct routed clusters

Compare a point-to-point cluster with a direct routed publish/subscribe cluster in respect of the core management tasks.

First, a point to point cluster:

1. When a new cluster queue is defined, the destination information is pushed to the full repository queue managers, and only sent to other cluster members when they first reference a cluster queue (for example, when an application attempts to open it). This information is then cached locally by the queue manager to remove the need to remotely retrieve the information each time the queue is accessed.
2. Adding a queue manager to a cluster does not directly affect the load on other queue managers. Information about the new queue manager is pushed to the full repositories, but channels to the new queue manager from other queue managers in the cluster are only created and started when traffic begins to flow to or from the new queue manager.

In summary, the load on a queue manager in a point-to-point cluster is related to the message traffic it handles for application programs and is not directly related to the size of the cluster.

Second, a direct routed publish/subscribe cluster:

1. When a new cluster topic is defined, the information is pushed to the full repository queue managers, and from there directly to all members of the cluster, causing channels to be started to each member of the cluster from the full repositories if not already started. If this is the first direct clustered topic, each queue manager member is sent information about all other queue manager members in the cluster.
2. When a subscription is created to a cluster topic on a new topic string, the information is pushed directly from that queue manager to all other members of the cluster immediately, causing channels to be started to each member of the cluster from that queue manager if not already started.
3. When a new queue manager joins an existing cluster, information about all clustered topics (and all queue manager members if a direct cluster topic is defined) is pushed to the new queue manager from the full repository queue managers. The new queue manager then synchronizes knowledge of all subscriptions to cluster topics in the cluster with all members of the cluster.

In summary, cluster management load at any queue manager in a direct routed publish/subscribe cluster grows with the number of queue managers, clustered topics, and changes to subscriptions on different topic strings within the cluster, irrespective of the local use of those cluster topics on each queue manager.

In a large cluster, or one where the rate of change of subscriptions is high, this level of cluster management can be a significant overhead across all queue managers.

Reducing the effect of direct routed publish/subscribe on performance

To reduce the effect of cluster management on the performance of a direct routed publish/subscribe cluster, consider the following options:

- Perform cluster, topic, and subscription updates at off-peak times of the day.
- Define a much smaller subset of queue managers involved in publish/subscribe, and make that an "overlapping" cluster. This cluster is then the cluster where cluster topics are defined. Although some queue managers are now in two clusters, the overall effect of publish/subscribe is reduced:
 - The size of the publish/subscribe cluster is smaller.

- Queue managers not in the publish/subscribe cluster are much less affected by cluster management traffic.

If the previous options do not adequately resolve your performance issues, consider using a *topic host routed* publish/subscribe cluster instead. For a detailed comparison of direct routing and topic host routing in publish/subscribe clusters, see [Designing publish/subscribe clusters](#).

Related concepts

Topic host routed publish/subscribe cluster performance

A topic host routed publish/subscribe cluster gives you precise control over which queue managers host each topic. These topic hosts become the *routing* queue managers for that branch of the topic tree. Moreover, queue managers without subscriptions or publishers have no need to connect with the topic hosts. This configuration can significantly reduce the number of connections between queue managers in the cluster, and the amount of information that is passed between queue managers.

Balancing producers and consumers in publish/subscribe networks

An important concept in asynchronous messaging performance is *balance*. Unless message consumers are balanced with message producers, there is the danger that a backlog of unconsumed messages might build up and seriously affect the performance of multiple applications.

Subscription performance in publish/subscribe networks

Distributed publish/subscribe in IBM MQ works by propagating knowledge of where subscriptions to different topic strings have been created in the queue manager network. This enables the queue manager on which a message is published to identify which other queue managers require a copy of the published message, to match their subscriptions.

Topic host routed publish/subscribe cluster performance

A topic host routed publish/subscribe cluster gives you precise control over which queue managers host each topic. These topic hosts become the *routing* queue managers for that branch of the topic tree. Moreover, queue managers without subscriptions or publishers have no need to connect with the topic hosts. This configuration can significantly reduce the number of connections between queue managers in the cluster, and the amount of information that is passed between queue managers.

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

- Topics are manually defined on individual *topic host* queue managers in the cluster.
- When a subscription is made on a cluster queue manager, proxy subscriptions are created only on the topic hosts.
- When an application publishes information to a topic, the receiving queue manager forwards the publication to a queue manager that hosts the topic. The topic host then sends the publication to all queue managers in the cluster that have valid subscriptions to the topic.

For a more detailed introduction to topic host routing, see [Topic host routing in clusters](#).

For many configurations, topic host routing is a more appropriate topology than *direct routing* because it provides the following benefits:

- Improved scalability of larger clusters. Only the topic host queue managers need to be able to connect to all other queue managers in the cluster. Therefore, there are fewer channels between queue managers, and there is less inter-queue manager publish/subscribe administrative traffic than for direct routing. When subscriptions change on a queue manager, only the topic host queue managers need to be informed.
- More control over the physical configuration. With direct routing, all queue managers assume all roles, and therefore all need to be equally capable. With topic host routing, you explicitly choose the topic host queue managers. Therefore, you can ensure that those queue managers are running on adequate equipment, and you can use less powerful systems for the other queue managers.

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

- System configuration and maintenance require more planning than for direct routing. You need to decide which points to cluster in the topic tree, and the location of the topic definitions in the cluster.

- Just as for direct routed topics, when a new topic host routed topic is defined, the information is pushed to the full repository queue managers, and from there direct to all members of the cluster. This event causes channels to be started to each member of the cluster from the full repositories if not already started.
- Publications are always sent to a host queue manager from a non-host queue manager, even if there are no subscriptions in the cluster. Therefore, you should use routed topics when subscriptions are typically expected to exist, or when the overhead of global connectivity and knowledge is greater than the risk of extra publication traffic.
- Messages that are published on non-host queue managers do not go direct to the queue manager that hosts the subscription, they are always routed through a topic host queue manager. This approach can increase the total overhead to the cluster, and increase message latency and reduce performance.

Note: For certain configurations, you can usefully remove this constraint as described in [Topic host routing using centralized publishers or subscribers](#).

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

Topic host routing with centralized publishers or subscribers

To remove the extra "hop" incurred when publications are always routed to subscriptions through a topic host queue manager, configure the publishers or the subscriptions on the same queue manager that hosts the topic. This approach brings maximum performance benefits in the following two cases:

- Topics with many publishers and few subscriptions. In this case, host the subscriptions on the topic host queue manager.
- Topics with few publishers and many subscriptions. In this case, host the publishers on the topic host queue manager.

The following figure shows a topic host queue manager that also hosts the subscriptions. This approach removes the extra "hop" between the publisher and the subscriber, and reduces unnecessary sharing of subscription knowledge across all members of the cluster:




Figure 32. Hosting publications on a topic host queue manager

Related concepts

Direct routed publish/subscribe cluster performance

In direct routed publish/subscribe clusters, information such as clustered topics and proxy subscriptions is pushed to all members of the cluster, irrespective of whether all cluster queue managers are actively participating in publish/subscribe messaging. This process can create a significant additional load on the system. To reduce the effect of cluster management on performance you can perform updates at off-peak times, define a much smaller subset of queue managers involved in publish/subscribe and make that an "overlapping" cluster, or switch to using topic host routing.

Balancing producers and consumers in publish/subscribe networks

An important concept in asynchronous messaging performance is *balance*. Unless message consumers are balanced with message producers, there is the danger that a backlog of unconsumed messages might build up and seriously affect the performance of multiple applications.

Subscription performance in publish/subscribe networks

Distributed publish/subscribe in IBM MQ works by propagating knowledge of where subscriptions to different topic strings have been created in the queue manager network. This enables the queue manager

on which a message is published to identify which other queue managers require a copy of the published message, to match their subscriptions.

Balancing producers and consumers in publish/subscribe networks

An important concept in asynchronous messaging performance is *balance*. Unless message consumers are balanced with message producers, there is the danger that a backlog of unconsumed messages might build up and seriously affect the performance of multiple applications.

In a point-to-point messaging topology, the relationship between message consumers and message producers is readily understood. You can get estimates of message production and consumption, queue by queue, channel by channel. If there is a lack of balance, the bottlenecks are readily identified and then remedied.

It is harder to work out whether publishers and subscribers are balanced in a publish/subscribe topology. Start from each subscription, and work back to the queue managers having publishers on the topic. Calculate the number of publications flowing to each subscriber from each queue manager.

Each publication that matches a subscription on a remote queue manager (based on proxy subscriptions) is put to a transmission queue. If multiple remote queue managers have proxy subscriptions for that publication, multiple copies of the message are put to a transmission queue, each targeted for a different sender channel.

In a publish/subscribe cluster, those publications are targeted at the `SYSTEM.INTER.QMGR.PUBS` queue on the remote queue managers that host the subscriptions. In a hierarchy, each publication is targeted at the `SYSTEM.BROKER.DEFAULT.STREAM` queue, or any other stream queues listed in the `SYSTEM.QPUBSUB.QUEUE.NAMELIST` on the remote queue managers. Each queue manager processes messages arriving on that queue and delivers them to the correct subscriptions on that queue manager.

For this reason, monitor the load at the following points where bottlenecks might arise:

- Monitor the load at the individual subscription queues.
 - This bottleneck implies that the subscribing application is not consuming the publications as quick as they are being published.
- Monitor the load at the `SYSTEM.INTER.QMGR.PUBS` queue or the stream queues.
 - This bottleneck implies that the queue manager is receiving publications from one or more remote queue managers faster than it can distribute them to the local subscriptions.
 - When seen on a topic host queue manager when using topic host routing in a cluster, consider making additional queue managers topic hosts, allowing the publication workload to be balanced across them. However, this will affect the message ordering across publications. See [Topic host routing using multiple topic hosts for a single topic](#).
- Monitor the load at the channels between the publishing queue manager and the subscribing queue managers, which are fed by the transmission queues on the publishing queue manager.
 - This bottleneck implies that either one or more channels is not running, or messages are being published to the local queue manager faster than the channels can deliver them to the remote queue manager.
 - When you use a publish/subscribe cluster, consider defining additional cluster receiver channels on the target queue manager. This allows the publication workload to be balanced across them. However, this affects the message ordering across publications. Also consider moving to a multiple cluster transmission queue configuration, because this can improve performance in certain circumstances.
- If the publishing application is using a queued publish/subscribe interface, monitor the load at (a) the `SYSTEM.BROKER.DEFAULT.STREAM` queue, and any other stream queues listed in the `SYSTEM.QPUBSUB.QUEUE.NAMELIST`; and (b) the `SYSTEM.BROKER.DEFAULT.SUBPOINT` queue, and any other subpoint queues listed in the `SYSTEM.QPUBSUB.SUBPOINT.NAMELIST`.
 - This bottleneck implies that messages are being put by local publishing applications faster than the local queue manager can process the messages.

Related concepts

Direct routed publish/subscribe cluster performance

In direct routed publish/subscribe clusters, information such as clustered topics and proxy subscriptions is pushed to all members of the cluster, irrespective of whether all cluster queue managers are actively participating in publish/subscribe messaging. This process can create a significant additional load on the system. To reduce the effect of cluster management on performance you can perform updates at off-peak times, define a much smaller subset of queue managers involved in publish/subscribe and make that an "overlapping" cluster, or switch to using topic host routing.

Topic host routed publish/subscribe cluster performance

A topic host routed publish/subscribe cluster gives you precise control over which queue managers host each topic. These topic hosts become the *routing* queue managers for that branch of the topic tree. Moreover, queue managers without subscriptions or publishers have no need to connect with the topic hosts. This configuration can significantly reduce the number of connections between queue managers in the cluster, and the amount of information that is passed between queue managers.

Subscription performance in publish/subscribe networks

Distributed publish/subscribe in IBM MQ works by propagating knowledge of where subscriptions to different topic strings have been created in the queue manager network. This enables the queue manager on which a message is published to identify which other queue managers require a copy of the published message, to match their subscriptions.

"Monitoring clusters" on page 333

Within a cluster you can monitor application messages, control messages, and logs. There are special monitoring considerations when the cluster load balances between two or more instances of a queue.

Subscription performance in publish/subscribe networks

Distributed publish/subscribe in IBM MQ works by propagating knowledge of where subscriptions to different topic strings have been created in the queue manager network. This enables the queue manager on which a message is published to identify which other queue managers require a copy of the published message, to match their subscriptions.

This approach minimizes the sending of published messages to queue managers on which no matching subscriptions exist. However, the propagation of the subscription knowledge can become a significant overhead, when the number of topic strings being subscribed to is high and constantly changing through frequent subscription creation and deletion.

You can affect performance by adjusting how publications and subscriptions are flowed around your publish/subscribe network. If your network traffic has few publications, and rapid subscription creation, deletion, or change, you can stop subscription information being flowed to all queue managers, and instead forward all publications to all queue managers in the network. You can also restrict the flow of proxy subscriptions and publications for a given topic between connected queue managers, restrict the flow of proxy subscriptions containing wildcards, and reduce the number and transient nature of topic strings.

Individual subscription propagation and *publish everywhere*

Publish everywhere is an alternative to individual subscription propagation. With individual propagation, only publications that have a matching subscription on a queue manager are forwarded to that queue manager. With *publish everywhere*, all publications are forwarded to all queue managers in the network. The receiving queue managers then deliver those publications that match local subscriptions.

Individual subscription propagation

This mechanism results in the least amount of inter-queue manager publication traffic, because only those publications that match subscriptions on a queue manager are sent.

However:

- For each individual topic string that is subscribed to, a proxy subscription is sent to other queue managers in the publish/subscribe topology. The set of queue managers depends on the routing model being used, as described in [Planning your distributed publish/subscribe network](#).

- This messaging overhead can be significant if there are many thousands of subscriptions to create or delete (for example, recreating all non-durable subscriptions after a restart of a queue manager), or if the set of subscriptions is changing rapidly, and each is to a different topic string.
- The number of queue managers to which the proxy subscription is propagated also affects the scale of the overhead.
- Proxy subscriptions are flowed to other queue managers using asynchronous messaging. This has the following effect:
 - There is a delay between the creation of a subscription, and the creation, delivery, and processing of the proxy subscription by the other queue managers.
 - Messages that are published at those queue managers in that interval are not delivered to the remote subscription.

Publish everywhere

With this mechanism there is no per topic string proxy subscription overhead on the system. This means that rapid subscription creation, deletion, or change does not result in increased network load and processing.

There is also no delay between creating a subscription and publications being flowed to a queue manager, because all publications are flowed to all queue managers. Therefore there is no window in which publications are not delivered to newly-created remote subscriptions.

However:

- Sending all publications to all queue managers in the publish/subscribe topology, can result in excessive network traffic where publications do not have matching subscriptions on each queue manager.
 - The greater the number of queue managers in the topology, the greater the overhead.

You should consider using the *publish everywhere* mechanism when you expect a publication to be subscribed to from a significant proportion of your queue managers, or where the proxy subscription overheads are too great because of the frequency of subscription changes. You should use individual proxy subscription forwarding in cases where you experience increased messaging traffic when publications are sent to all queue managers, rather than to the queue managers with matching subscriptions.

You can set *publish everywhere* behavior at any level within the topic tree. To enable *publish everywhere*, you set the **PROXYSUB** parameter to **FORCE** for a high-level topic object. This results in a single wildcard proxy subscription that matches all topics below this topic object in the topic tree. When set it on a clustered topic object, the **PROXYSUB (FORCE)** attribute is propagated to every queue manager in the network, not just the queue manager that the topic was defined on.

Note: When used in a hierarchy, you set **PROXYSUB (FORCE)** individually on each queue manager, so the topology mechanism naturally limits the number of channels. However, when used in a cluster, many additional channels might be started:

- In a topic host routed cluster, channels are started from each queue manager to each topic host queue manager.
- In a direct routed cluster, channels are started from every queue manager to every other queue manager.

The overhead of starting many channels is most pronounced in a direct routed cluster, and can cause performance issues. See [“Direct routed publish/subscribe cluster performance”](#) on page 399.

Other ways of restricting the flow of proxy subscriptions and publications between connected queue managers

Consolidate topic strings

The use of many distinct, transient, topic strings introduces some level of management overhead on each queue manager in the system where publishers or subscriptions are attached. You should periodically assess the use of topic strings to see whether they can be consolidated. Reducing the

number and transient nature of topic strings, and therefore the publishers and subscriptions to them, reduces the impact on the system.

Restrict publication and subscription scope

For a given topic, you can use the [Publication scope](#) and [Subscription scope](#) settings to keep publications and subscriptions local to the queue manager on which they are defined.

Block subscriptions made to wildcarded topics

You can restrict the flow of proxy subscriptions containing wildcards by setting the **Topic** attribute WILDCARD to BLOCK. See [Wildcards in proxy subscriptions](#).

See also [“Balancing producers and consumers in publish/subscribe networks”](#) on page 405

Monitoring proxy subscription traffic in clusters

When considering the load on the system from the proxy subscription traffic, in addition to monitoring the queues listed in [“Balancing producers and consumers in publish/subscribe networks”](#) on page 405, monitor the following cluster queues:

- The SYSTEM.INTER.QMGR.FANREQ queue on the subscriber queue manager.
- The SYSTEM.INTER.QMGR.CONTROL queue on all other queue managers in the cluster.

Any significant message backlog on these queues implies that either the rate of subscription change is too great for the system, or a queue manager is not correctly functioning in the cluster. If you suspect the problem lies with a specific queue manager, check that publish/subscribe support is not disabled for that queue manager. See **PSMODE** in [ALTER QMGR](#).

Related concepts

Direct routed publish/subscribe cluster performance

In direct routed publish/subscribe clusters, information such as clustered topics and proxy subscriptions is pushed to all members of the cluster, irrespective of whether all cluster queue managers are actively participating in publish/subscribe messaging. This process can create a significant additional load on the system. To reduce the effect of cluster management on performance you can perform updates at off-peak times, define a much smaller subset of queue managers involved in publish/subscribe and make that an "overlapping" cluster, or switch to using topic host routing.

Topic host routed publish/subscribe cluster performance

A topic host routed publish/subscribe cluster gives you precise control over which queue managers host each topic. These topic hosts become the *routing* queue managers for that branch of the topic tree. Moreover, queue managers without subscriptions or publishers have no need to connect with the topic hosts. This configuration can significantly reduce the number of connections between queue managers in the cluster, and the amount of information that is passed between queue managers.

Balancing producers and consumers in publish/subscribe networks

An important concept in asynchronous messaging performance is *balance*. Unless message consumers are balanced with message producers, there is the danger that a backlog of unconsumed messages might build up and seriously affect the performance of multiple applications.

Proxy subscriptions in a publish/subscribe network

Reducing the number of unwanted topics in the topic tree

The performance of a publish/subscribe system is improved by reducing the number of unwanted topics in the topic tree. What is an unwanted topic and how do you remove them?

You can create large numbers of topics without affecting performance adversely. However, some ways of using publish/subscribe result in continually expanding topic trees. An exceptionally large number of topics are created once and never used again. The growing number of topics might become a performance problem.

How can you avoid designs that lead to a large and growing number of unwanted topics? What can you do to help the queue manager remove unwanted topics from the topic tree?

The queue manager recognizes an unwanted topic because it has been unused for 30 minutes. The queue manager removes unused topics from the topic tree for you. The 30 minute duration can be changed by altering the queue manager attribute, **TREELIFE**. You can help the queue manager to remove unwanted topics by making sure that the topic appears to the queue manager to be unused. The section, [“What is an unused topic?”](#) on page 409 explains what an unused topic is.

A programmer, designing any application, and especially designing a long running application, considers its resource usage: how much resource the program requires, are there any unbounded demands, and any resource leaks? Topics are a resource that publish/subscribe programs use. Scrutinize the use of topics just like any other resource a program uses.

What is an unused topic?

Before defining what an unused topic is, what exactly counts as a topic?

When a topic string, such as USA/Alabama/Auburn, is converted into a topic, the topic is added to the topic tree. Additional topic nodes, and their corresponding topics, are created in the tree, if necessary. The topic string USA/Alabama/Auburn is converted into a tree with three topics.

- USA
- USA/Alabama
- USA/Alabama/Auburn

To display all the topics in the topic tree, use the **runmqsc** command `DISPLAY TPSTATUS( '# ' ) TYPE(TOPIC)`.

An unused topic in the topic tree has the following properties.

It is not associated with a topic object

An administrative topic object has a topic string that associates it with a topic. When you define the topic object Alabama, if the topic, USA/Alabama, it is to be associated with does not exist, the topic is created from the topic string. If the topic does exist, the topic object and the topic are associated together using the topic string.

It does not have a retained publication

A topic with a retained publication results from a publisher putting a message to a topic with the option `MQPMO_RETAIN`.

Use the **runmqsc** command `DISPLAY TPSTATUS( 'USA/Alabama' ) RETAINED` to check if USA/Alabama has a retained publication. The response is YES or NO.

Use the **runmqsc** command `CLEAR TOPICSTR( 'USA/Alabama' ) CLTRTYPE(RETAINED)` to remove a retained publication from USA/Alabama.

It has no child topics

USA/Alabama/Auburn is a topic with no child topics. USA/Alabama/Auburn is the direct child topic of USA/Alabama.

Display the direct children of USA/Alabama with the **runmqsc** command `DISPLAY TPSTATUS( 'USA/Alabama/+' )`.

There are no active publishers to the node

An active publisher to a node is an application that has the topic open for output.

For example, an application opens the topic object named **Alabama** with open options `MQOO_OUTPUT`.

To display active publishers to USA/Alabama and all its children, use the **runmqsc** command `DISPLAY TPSTATUS( 'USA/Alabama/#' ) TYPE(PUB) ACTCONN`.

There are no active subscribers to the node

An active subscriber can either be a durable subscription, or an application that has registered a subscription to a topic with `MQSUB`, and not closed it.

To display active subscriptions to USA/Alabama, use the **runmqsc** command `DISPLAY TPSTATUS('USA/Alabama') TYPE(SUB) ACTCONN`.

To display active subscriptions to USA/Alabama and all its children, use the **runmqsc** command `DISPLAY TPSTATUS('USA/Alabama/#') TYPE(SUB) ACTCONN`.

Managing the number of topics in a topic tree

In summary, there are a number of ways to manage the number of topics in a topic tree.

Display TPCOUNT

Use the **runmqsc** command `DISPLAY PUBSUB ALL` periodically to display the **TPCOUNT** property. This is the number of topic nodes in the topic tree. If the number is growing it might indicate that a shorter **TREELIFE** is required, or that a redesign of the topics themselves is required.

Modify TREELIFE

An unused topic has a lifetime of 30 minutes by default. You can make the lifetime of an unused topic smaller.

For example, The **runmqsc** command, `ALTER QMGR TREELIFE(900)`, reduces lifetime of an unused topic from 30 minutes to 15 minutes.

Exceptionally, restart the queue manager

When the queue manager is restarted, the topic tree is reinitialized from topic objects, nodes with retained publications, and durable subscriptions. Topics that had been created by the operation of publisher and subscriber programs are eliminated.

As a last resort, if the growth in unwanted topics has been the cause of performance problems in the past, restart the queue manager.

Related concepts

[Topic trees](#)

Aspera gateway can improve performance over high latency networks

The IBM Aspera® faspio Gateway provides a fast TCP/IP tunnel that can significantly increase network throughput for IBM MQ.

The Aspera gateway can be used to improve the performance of queue manager channels. It is especially effective if the network has high latency or tends to lose packets, and it is typically used to speed up the connection between queue managers in different data centers.

However, for a fast network that does not lose packets there is a decrease in performance when using the Aspera gateway, so it is important to check network performance before and after defining an Aspera gateway connection.

A queue manager running on any entitled platform can connect through an Aspera gateway. The gateway itself is deployed on Red Hat® or Ubuntu Linux, or Windows.

For more information, see [Defining an Aspera gateway connection on Linux or Windows](#).

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 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
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 IBM 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 (<https://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: