

MQ 3.5 Cluster Protocol

@(#)index.html 1.2 08/24/06

The cluster protocol version is determined by the version of the master broker. If the cluster master broker is 3.0.1 or older, then the old cluster protocol is used.

If the master broker is an MQ 3.5 broker, then the new cluster protocol described in this document is used.

The new cluster protocol uses the GPacket packet format. This packet format resembles the MQ client-broker protocol. It will also allow us to make changes in future releases without breaking compatibility with old clusters.

The following is a summary of the packets exchanged between clustered 3.5 S1MQ brokers.

Message Delivery Protocol	
1	G_MESSAGE_DATA
2	G_MESSAGE_DATA_REPLY
3	G_MESSAGE_ACK
4	G_MESSAGE_ACK_REPLY

Interest Management	
5	G_NEW_INTEREST
6	G_NEW_INTEREST_REPLY
7	G_REM_DURABLE_INTEREST
8	G_REM_DURABLE_INTEREST_REPLY
9	G_INTEREST_UPDATE
10	G_INTEREST_UPDATE_REPLY

Resource Locking	
11	G_LOCK_REQUEST
12	G_LOCK_RESPONSE

Destination Forwarding	
13	G_DESTINATION_UPDATE
14	G_DESTINATION_UPDATE_REPLY

Master Broker Change Log Protocol	
15	G_CONFIG_CHANGE_EVENT
16	G_CONFIG_CHANGE_EVENT_ACK
17	G_GET_CONFIG_CHANGES_REQUEST
18	G_GET_CONFIG_CHANGES_RESPONSE

Miscellaneous	
19	G_CLIENT_CLOSED
20	G_CLIENT_CLOSED_REPLY

Cluster Flow Control Protocol	
21	G_STOP_MESSAGE_FLOW
22	G_STOP_MESSAGE_FLOW_REPLY
23	G_RESUME_MESSAGE_FLOW
24	G_RESUME_MESSAGE_FLOW_REPLY
25	G_RELOAD_CLUSTER
26	G_RELOAD_CLUSTER_REPLY

Link Initialization Protocol (Old packet format is used for backward compatibility)	
-	BROKER_INFO
-	LINK_INIT

G_CLIENT_CLOSED

Client closed notification.

This message is sent by a broker when any of its client connections goes down, so that other brokers in the cluster can remove any temporary resources (e.g. temporary queue destinations) owned by that client.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast

Reply: G_CLIENT_CLOSED_REPLY

Fixed Header Fields	
Field	Value
packet type	19 (G_CLIENT_CLOSED)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
I	Long	Client connection UID.	MQ3.5

Packet Body
None.

G_CLIENT_CLOSED_REPLY

Reply to a G_CLIENT_CLOSED packet

Fixed Header Fields	
Field	Value
packet type	20 (G_CLIENT_CLOSED_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_GET_CONFIG_CHANGES_REQUEST

Every broker sends this message during its initialization phase to the master broker. This message serves as a request for information regarding any configuration changes that might have happened while the requesting broker was down.

Brokers maintain a persistent timestamp to remember the last time they had synchronized with the cluster master broker, and only request changes since then.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast

Reply: G_GET_CONFIG_CHANGES_RESPONSE

Fixed Header Fields	
Field	Value
packet type	17 (G_GET_CONFIG_CHANGES_REQUEST)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
TS	Long	Last synchronization time. Master broker will return all the changes since this time.	MQ3.5

Packet Body
None.

G_GET_CONFIG_CHANGES_RESPONSE

Fixed Header Fields	
Field	Value
packet type	18 (G_GET_CONFIG_CHANGES_RESPONSE)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
TS	Long	Master broker's current time.	MQ3.5
C	Int	Number of records.	
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

Packet Body
EventData GPacket (byte[])
EventData GPacket (byte[])
...
EventData GPacket (byte[])

G_CONFIG_CHANGE_EVENT

Whenever a broker encounters a configuration change event, it uses this message to log the event to the master broker configuration change log. The configuration change log is maintained by the broker designated as the "master broker".

Following configuration change events are recorded -

G_DESTINATION_UPDATE
G_INTEREST_UPDATE (For durable subscriptions only)

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast

Reply: G_CONFIG_CHANGE_EVENT_REPLY

Fixed Header Fields	
Field	Value
packet type	15 (G_CONFIG_CHANGE_EVENT)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
X	Long	Unique correlation ID to match the request and reply packets.	MQ3.5
T	Int	G_DESTINATION_UPDATE This is a destination update. G_INTEREST_UPDATE This is a durable subscription update.	MQ2.0

Packet Body
Event Data.

G_CONFIG_CHANGE_EVENT_ACK

Fixed Header Fields	
Field	Value
packet type	16 (G_CONFIG_CHANGE_EVENT_ACK)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields	
None	

Properties:			
Property	Type	Value	Since
X	Long	Unique correlation ID to match the request and reply packets.	MQ3.5
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_DESTINATION_UPDATE

Destination create, update or destroy notification. This notification is sent only if the destination is created by the administrator and not for temporary or auto-created destinations.

Only running brokers receive these broadcast notifications. If a broker joins the cluster later it is expected to talk to the master broker first and get all the updates it may have missed.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast

Reply: G_DESTINATION_UPDATE_REPLY

Fixed Header Fields	
Field	Value
packet type	13 (G_DESTINATION_UPDATE)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
N	String	Destination name string.	MQ3.5
DT	Int	Destination type.	MQ3.5
T	Int	G_NEW_DESTINATION Create a new destination. G_UPDATE_DESTINATION Update the destination attributes. G_REM_DESTINATION Remove the destination.	MQ2.0

Packet Body
Destination properties name value pairs.

G_DESTINATION_UPDATE_REPLY

Reply to a G_DESTINATION_UPDATE packet

Fixed Header Fields	
Field	Value
packet type	14 (G_DESTINATION_UPDATE_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_INTEREST_UPDATE

Interest change notification. Sent as Unicast as well as Broadcast. Unicast is used when a new broker connection is established. Existing brokers sent their current interests to the new broker using a unicast message.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Both

Reply: G_INTEREST_UPDATE_REPLY

Fixed Header Fields	
Field	Value
packet type	9 (G_INTEREST_UPDATE)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
C	Int	Number of interests.	MQ3.5
T	Int	G_REM_INTEREST Remove an interest.	MQ3.5
		G_DURABLE_DETACH Client has detached from durable topic interest.	
		G_NEW_PRIMARY_INTEREST This is the new primary interest for the queue destination.	

Packet Body
ConsumerUID = Interest(1)
ConsumerUID = Interest(2)
ConsumerUID = Interest(3)
...
ConsumerUID = Interest(n)

G_INTEREST_UPDATE_REPLY

Reply to a G_INTEREST_UPDATE packet

Fixed Header Fields	
Field	Value
packet type	10 (G_INTEREST_UPDATE_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_LOCK

Request a mutually exclusive ownership for a resource.

This message is sent when a broker needs to ensure mutually exclusive ownership of a resource on behalf of a client (i.e. obtain a cluster-wide lock on the resource). Examples of such resources are - clientID, durable subscription, queue primary receiver role etc.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast

Reply: G_LOCK_REPLY

Fixed Header Fields	
Field	Value
packet type	11 (G_LOCK)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
I	String	Unique string representing the resource to be locked.	MQ2.0
TS	Long	Resource creation time. Used for resolving lock contention - locks are granted in chronological order.	MQ2.0
X	Long	Unique correlation ID. Used for matching requests and replies.	MQ2.0

Packet Body
None.

G_LOCK_REPLY

Reply to a G_LOCK packet

Fixed Header Fields	
Field	Value
packet type	12 (G_LOCK_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields	
None	

G_MESSAGE_ACK

Message acknowledgement sent from target broker to the home broker.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Unicast

Reply: G_MESSAGE_ACK_REPLY

Fixed Header Fields	
Field	Value
packet type	3 (G_MESSAGE_ACK)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
T	Int	MB_MESSAGE_IGNORED The target interest is invalid. If it's a topic message, acknowledge it. For queue message, deliver it again to the next receiver. MB_MESSAGE_DELIVERED The message has been delivered to the client. MB_MESSAGE_CONSUMED The client has acknowledged the message.	MQ2.0

Packet Body
Number of Acknowledgements
SysMessageID (1)
ConsumerUID (1)
SysMessageID (2)
ConsumerUID (2)
...
SysMessageID (n)
ConsumerUID (n)

G_MESSAGE_ACK_REPLY

Reply to a G_MESSAGE_ACK packet

Fixed Header Fields	
Field	Value
packet type	4 (G_MESSAGE_ACK_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_MESSAGE_DATA

Deliver a message to remote targets (i.e. interests residing on a different broker).

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Unicast

Reply: G_MESSAGE_DATA_REPLY

Fixed Header Fields	
Field	Value
packet type	1 (G_MESSAGE_DATA)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
D	Boolean	If this property is set, the target broker must send a message delivery notification when the message is delivered to the client.	MQ2.0

Packet Body
Number of target ConsumerUIDs sendMsgDeliveredAck (boolean) <ul style="list-style-type: none"> ConsumerUID = Target(1) ConsumerUID = Target(2) ... ConsumerUID = Target(n) ReadOnlyPacket

G_MESSAGE_DATA_REPLY

Reply to a G_MESSAGE_DATA packet

Fixed Header Fields	
Field	Value
packet type	2 (G_MESSAGE_DATA_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Consumer deleted ERROR(500) Internal server error	MQ3.5

G_NEW_INTEREST

New interest creation notification. Sent as unicast as well as broadcast. Unicast is used when a new broker connection is established. Existing brokers send their current interests to the new broker using a unicast message.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Both

Reply: G_NEW_INTEREST_REPLY

Fixed Header Fields	
Field	Value
packet type	5 (G_NEW_INTEREST)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
N	String	Durable name. (Optional. Only used when the packet is stored as a master broker change log entry.	MQ3.5
I	String	Client ID. (Optional. Only used when the packet is stored as a master broker change log entry.	MQ3.5
M	Boolean	True if this packet is received as a config change event from the master broker. False if this is a "real time" event - i.e. the packet was broadcast by the broker where the interest was created.	
C	Integer	Number of interests.	MQ3.5

Packet Body
Serialized Interest object = Interest(1) Serialized Interest object = Interest(2) Serialized Interest object = Interest(3) ... Serialized Interest object = Interest(n)

G_NEW_INTEREST_REPLY

Reply to a G_NEW_INTEREST packet

Fixed Header Fields	
Field	Value
packet type	6 (G_NEW_INTEREST_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_RELOAD_CLUSTER

Reload the cluster configuration and initialize the connections.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast

Reply: G_RELOAD_CLUSTER_REPLY

Fixed Header Fields	
Field	Value
packet type	25 (G_RELOAD_CLUSTER)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:
None.

Packet Body
None.

G_RELOAD_CLUSTER_REPLY

Reply to a G_RELOAD_CLUSTER packet

Fixed Header Fields	
Field	Value
packet type	26 (G_RELOAD_CLUSTER_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_REM_DURABLE_INTEREST

Durable interest unsubscribe notification. This packet is also stored in the master broker change log as a record of the durable unsubscription event.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast.

Reply: G_NEW_INTEREST_REPLY

Fixed Header Fields	
Field	Value
packet type	7 (G_REM_DURABLE_INTEREST)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:			
Property	Type	Value	Since
N	String	Durable name. (Optional. Only used when the packet is stored as a master broker change log entry.	MQ3.5
I	String	Client ID. (Optional. Only used when the packet is stored as a master broker change log entry.	MQ3.5
C	Integer	Number of interests.	MQ3.5

Packet Body
DurableName(1) = String ClientID(1) = String DurableName(2) = String ClientID(2) = String ... DurableName(n) = String ClientID(n) = String

G_REM_DURABLE_INTEREST_REPLY

Reply to a G_REM_DURABLE_INTEREST packet

Fixed Header Fields	
Field	Value
packet type	8 (G_REM_DURABLE_INTEREST_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_RESUME_MESSAGE_FLOW

This message is sent by a broker when it is ready to accept messages from other brokers.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast

Reply: G_RESUME_MESSAGE_FLOW_REPLY

Fixed Header Fields	
Field	Value
packet type	23 (G_RESUME_MESSAGE_FLOW)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:
None.

Packet Body
None.

G_RESUME_MESSAGE_FLOW_REPLY

Reply to a G_RESUME_MESSAGE_FLOW packet

Fixed Header Fields	
Field	Value
packet type	24 (G_RESUME_MESSAGE_FLOW_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5

G_STOP_MESSAGE_FLOW

Ask other brokers to stop message delivery.

This message may be sent by a broker when it is experiencing resource shortage.

Since: MQ2.0

Modified: MQ3.5

Unicast / Broadcast: Broadcast

Reply: G_STOP_MESSAGE_FLOW_REPLY

Fixed Header Fields	
Field	Value
packet type	21 (G_STOP_MESSAGE_FLOW)
version, size, timestamp, sequence number, property offset, property size.	As defined in the GPacket specification.
A bit	1 (remote broker must send a reply)
All other bits	Not Applicable
priority	Not applicable

Properties:
None.

Packet Body
None.

G_STOP_MESSAGE_FLOW_REPLY

Reply to a G_STOP_MESSAGE_FLOW packet

Fixed Header Fields	
Field	Value
packet type	22 (G_STOP_MESSAGE_FLOW_REPLY)
All bit fields	Not Applicable
priority	Not Applicable

Variable Header Fields
None

Properties:			
Property	Type	Value	Since
S	Int	OK(200) Okay! ERROR(500) Internal server error	MQ3.5