



IBM Software Group

WebSphere MQ Queue Manager Clustering

Beverly Brown, Angel Rivera
WebSphere MQ Level 2 Customer Support



WebSphere® Support Technical Exchange



Agenda

- What is clustering
- Setting up a basic cluster
- DISPLAY commands
- Workload balancing
- References to intermediate and advanced configuration and administration topics
- Troubleshooting
- References to manuals, technotes, articles and Redbooks

What is clustering?

- Clustering is a way to logically group WebSphere MQ queue managers so that you have:
 - reduced system administration due to fewer channel, remote queue, and transmission queue definitions.
 - increased availability and workload balancing due to the ability to define instances of the same queue on multiple queue managers in the cluster and distribute messages to those queues.



Setting up a basic cluster (Step 1)

- For details of these steps, reference “Chapter 3. First tasks” in the WebSphere MQ Queue Manager Clusters manual.
- Decide on the organization of the cluster and its name.

Setting up a basic cluster (Step 2)

- Determine which queue manager should hold full repositories.
- You will need at least one, preferably two. A full repository contains a complete set of information about every queue manager and object in the cluster.
- If you need more than two, refer to “Designing clusters” in Chapter 7 of the Queue Manager Clusters manual. You need to ensure connectivity since a queue manager sends out information about itself to only two full repositories.



Setting up a basic cluster (Step 3)

- Alter the queue manager definitions to add repository definitions.
- Do this for each queue manager that is to hold a full repository.

```
ALTER QMGR REPOS(cluster_name)
```

Setting up a basic cluster (Step 4)

- Define the CLUSRCVR channels.
- Every queue manager in the cluster needs a CLUSRCVR with a conname pointing to itself.

```
DEFINE CHANNEL(channel_name)  
CHLTYPE(CLUSRCVR) TRTYPE(TCP)  
CONNNAME('my_ip_name_or_address(port)')  
CLUSTER(cluster_name)
```

Setting up a basic cluster (Step 5)

- Define the CLUSSDR channels.
- Define one CLUSSDR to a full repository queue manager. The channel name must match that of the CLUSRCVR on the full repository. DO NOT define a CLUSSDR to point to a partial repository.

```
DEFINE CHANNEL(channel_name)  
CHLTYPE(CLUSSDR) TRPTYP(TCP)  
CONNNAME('remote_ip_name_or_address(port)')  
CLUSTER(cluster_name)
```


Setting up a basic cluster (Step 6)

- Define a cluster queue.
- Other queue managers in the cluster can send message to it without making remote-queue definitions for it.
- Only the local queue manager can read messages from an instance of the cluster queue.

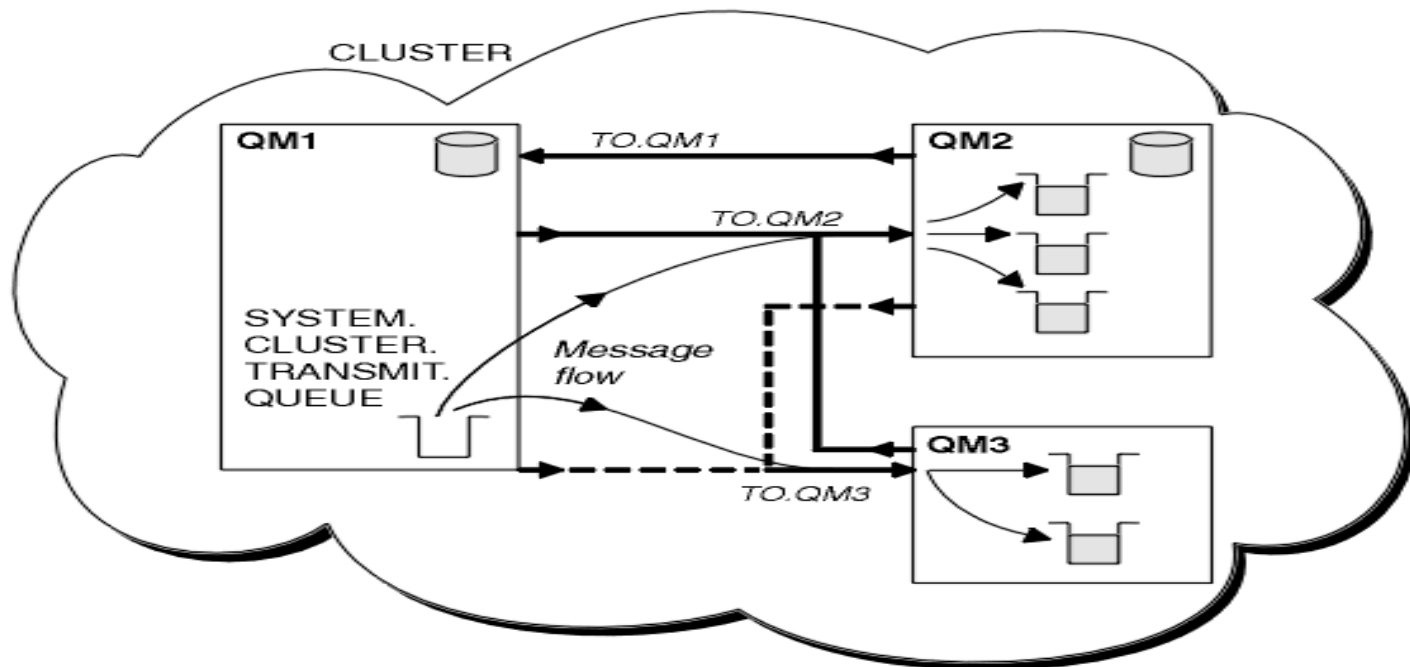
```
DEFINE QLOCAL(qname)  
CLUSTER(cluster_name)
```

Setting up a basic cluster (Step 7)

- Verify and test the cluster.
- See the following slides for display commands to review your definitions.
- You can use a sample program to test putting messages to clustered queues, for example
 - amqsput for distributed platforms
 - CSQ4BCK1 or CSQ4BVK1 for z/OS



Picture of manual and auto defined channels



- “Figure 8. A cluster of queue managers, showing auto-defined channels” from “WebSphere MQ Queue Manager Clusters Version 6.0” (SC34-6061-03) at <http://publib.boulder.ibm.com/infocenter/wmqv6/v6r0/index.jsp?topic=/com.ibm.mq.csqzah.doc/csq684b.htm>

Description of picture

- In the figure on the prior page, QM1 and QM2 are full repositories.
- The solid lines represent manually defined CLUSRCVR and CLUSSDR channels.
- The two dashed lines represent cluster-sender channels that have been created automatically.
- QM2 and QM3 host some clustered queues that can be put to by any other queue manager in the cluster.
- All queue managers have a `SYSTEM.CLUSTER.TRANSMIT.QUEUE` from which they can send messages to any other queue manager in the same cluster.

DISPLAY QMGR

- **DISPLAY QMGR REPOS REPOSNL QMID**

AMQ8408: Display Queue Manager details.

```
QMNAME(QM1) QMID(QM1_2005-07-12_17.14.38)  
REPOS(QMCLUS)REPOSNL( )
```

- QMID is an internally generated unique name that consists of the queue manager name plus the time the queue manager was created. For z/OS, the time is expressed as 16 hex digits, for example RTPH.BDA2D50C17934846.

DISPLAY CLUSQMGR

- **DISPLAY CLUSQMGR(*) ALL**

[The following output is partial—it is only for one cluster queue manager of several]

AMQ8441: Display Cluster Queue Manager details.

CLUSQMGR(QM2)	ALTDATE()
ALTTIME()	BATCHHB(0)
BATCHINT(0)	BATCHSZ(50)
CHANNEL(TO.QM2)	CLUSDATE(2007-01-09)
CLUSTER(QMCLUS)	CLUSTIME(13.59.18)
CLWLPRTY(0)	CLWLRANK(0)
CLWLWGHT(50)	COMPHDR(NONE)
COMPMSG(NONE)	
CONNNAME(craigbev.dyn.webahead.ibm.com(1516))	
CONVERT(NO)	DEFTYPE(CLUSDRB)
DESCR()	DISCINT(6000)
HBINT(300)	KAINT(AUTO)
LOCLADDR()	LONGRTY(999999999)
LONGTMR(1200)	MAXMSGL(4194304)
MCANAME()	MCTYPE(THREAD)

... (continued on next slide)

DISPLAY CLUSQMGR (continuation)

- (continuation of previous slide)

MCAUSER()

MRDATA()

MRRTY(10)

MSGDATA()

NETPRTY(0)

PASSWORD()

QMID(QM2_2005-07-12_17.22.29)

RCVDATA()

SCYDATA()

SENDDATA()

SEQWRAP(999999999)

SHORTTMR(60)

SSLCIPH()

STATUS(RUNNING)

TPNAME()

USERID()

MODENAME()

MREXIT()

MRTMR(1000)

MSGEXIT()

NPMSPEED(FAST)

PUTAUT(DEF)

QMTYPE(REPOS)

RCVEXIT()

SCYEXIT()

SENDEXIT()

SHORTRTY(10)

SSLCAUTH(REQUIRED)

SSLPEER()

SUSPEND(NO)

TRPTYPE(TCP)

...

DISPLAY CLUSQMGR notes

- DEFTYPE can be one of the following:
 - CLUSSDR** - Defined explicitly as a cluster-sender channel
 - CLUSSDRA** - Defined by auto-definition as a cluster-sender channel (A=auto)
 - CLUSSDRB** - Defined as a cluster-sender channel, both explicitly and by auto-definition (B=both)
 - CLUSRCVR** - Defined as a cluster-receiver channel

Auto-defined cluster-sender channels take their attributes from those specified in the corresponding cluster-receiver channel definition on the receiving queue manager. You cannot directly modify an auto-defined CLUSSDR. You cannot see them via DISPLAY CHANNEL.

DISPLAY CLUSQMGR notes (continuation)

- QMTYPE may be REPOS (full repository) or NORMAL (partial repository).
- You will not see the status of the CLUSRCVR—use DISPLAY CHSTATUS for that.



DISPLAY CHSTATUS

- **dis chstatus(*) all**

AMQ8417: Display Channel Status details.

CHANNEL (TO . QM2)

CONNNAME (craigbev.dyn.webahead.ibm.com (1516))

CURRENT

STATUS (RUNNING)

XMITQ (SYSTEM . CLUSTER . TRANSMIT . QUEUE)

CHLTYPE (CLUSSDR)

RQMNAME (QM2)

SUBSTATE (MQGET)

AMQ8417: Display Channel Status details.

CHANNEL (TO . QM1)

CONNNAME ()

RQMNAME (QM3)

SUBSTATE (RECEIVE)

CHLTYPE (CLUSRCVR)

CURRENT

STATUS (RUNNING)

XMITQ ()

AMQ8417: Display Channel Status details.

CHANNEL (TO . QM3)

CONNNAME (craigbev.dyn.webahead.ibm.com (1517))

CURRENT

STATUS (RUNNING)

XMITQ (SYSTEM . CLUSTER . TRANSMIT . QUEUE)

CHLTYPE (CLUSSDR)

RQMNAME (QM3)

SUBSTATE (MQGET)

AMQ8417: Display Channel Status details.

CHANNEL (TO . QM1)

CONNNAME ()

RQMNAME (QM2)

SUBSTATE (RECEIVE)

CHLTYPE (CLUSRCVR)

CURRENT

STATUS (RUNNING)

XMITQ ()

DISPLAY QCLUSTER

- **DISPLAY QCLUSTER(*) ALL**

AMQ8409: Display Queue details.

QUEUE (INQUIRY)

ALTDATE(2007-01-09)

CLUSDATE(2007-01-09)

CLUSQMGR (QM2)

CLUSTIME(15.30.25)

CLWLRANK(0)

DEFPRTY(0)

DESCR()

QMID(QM2_2005-07-12_17.22.29)

AMQ8409: Display Queue details.

QUEUE (INQUIRY)

ALTDATE(2007-01-09)

CLUSDATE(2007-01-09)

CLUSQMGR (QM1)

CLUSTIME(15.30.03)

CLWLRANK(0)

DEFPRTY(0)

DESCR()

QMID(QM1_2005-07-12_17.14.38)

TYPE(QCLUSTER)

ALTTIME(15.30.25)

CLUSTER(QMCLUS)

CLUSQT(QLOCAL)

CLWLPRTY(0)

DEFBIND(NOTFIXED)

DEFPSIST(NO)

PUT(ENABLED)

TYPE(QCLUSTER)

ALTTIME(15.30.03)

CLUSTER(QMCLUS)

CLUSQT(QLOCAL)

CLWLPRTY(0)

DEFBIND(NOTFIXED)

DEFPSIST(NO)

PUT(ENABLED)

DISPLAY QCLUSTER notes

- It displays information about clustered queues only.
- A cluster queue will not be displayed on a partial repository until an application has opened it.
- The DEFBIND(NOTFIXED|OPEN) option affects workload balance (mentioned in later screens) if the application uses the MQOPEN option MQOO_BIND_AS_Q_DEF.

DISPLAY QUEUE(*) CLUSINFO

- **DISPLAY QUEUE(*) CLUSINFO**
[The following output is partial]

```
AMQ8409: Display Queue details.  
  QUEUE(SYSTEM.CLUSTER.COMMAND.QUEUE)      TYPE(QLOCAL)  
AMQ8409: Display Queue details.  
  QUEUE(SYSTEM.CLUSTER.REPOSITORY.QUEUE)    TYPE(QLOCAL)  
AMQ8409: Display Queue details.  
  QUEUE(SYSTEM.CLUSTER.TRANSMIT.QUEUE)      TYPE(QLOCAL)  
AMQ8409: Display Queue details.  
  QUEUE(TESTCLUS)                            TYPE(QLOCAL)  
AMQ8409: Display Queue details.  
  QUEUE(INQUIRY)                             TYPE(QCLUSTER)  
AMQ8409: Display Queue details.  
  QUEUE(INQUIRY)                             TYPE(QCLUSTER)
```

DISPLAY QUEUE(*) CLUSINFO

- This command displays information about queues with TYPE(QCLUSTER), i.e. clustered queues, in addition to the usual QLOCAL, QALIAS, QREMOTE, and QMODEL types.
- For clustered queues, the information is obtained from this queue manager's repository. A partial repository will not know about a clustered queue until it has been opened by that queue manager.
- On z/OS, you cannot issue this from CSQINP2.



Workload balancing

- When a cluster contains more than one instance of the same queue, workload balancing determines the best queue manager to route a message to. It takes into consideration:
 - the availability channels, the queue manager and queues
 - how to treat a local copy of the queue (CLWLUSEQ in V6).
- At its simplest, workload management results in a round-robin effect.



Workload balancing (continuation 1)

- MQ V6 has additional parameters that can be used to influence the results of the algorithm.
 - Queues: CLWLPRTY, CLWLRANK, CLWLUSEQ
 - Queue Managers: CLWLUSEQ, CLWLMRUC
 - Channels: CLWLPRTY, CLWLRANK, CLWLWGHT, NETPRTY
- You can write your own cluster workload user exit

Workload balancing (continuation 2)

- For workload balancing to occur:
 - open the queue with the MQOO_BIND_NOT_FIXED open option, or
 - open with the default MQOO_BIND_AS_Q_DEF and with DEFBIND(NOTFIXED) set in the queue definition. DEFBIND(OPEN) is the default.

Workload balancing (continuation 3)

- Leave MQMD.ObjectQMgrName blank to allow the queue manager to choose the queue instance. To force the message to a specific instance of the clustered queue, specify that queue manager's name in ObjectQmgrName.



Intermediate configuration

- A queue manager may be a member of more than one cluster. List those clusters in a NAMELIST. You can have more than one namelist to achieve overlapping clusters.

You can alter a full repository QMGR to use REPOSNL(namelist) rather than REPOS.

For channels and queues, you can specify CLUSNL(namelist) rather than specifying the CLUSTER parameter.

References to advanced tasks

- Chapter 4 of the Queue Manager Clusters manual describes how you can use queue-manager aliases, reply-to queue aliases, and queue aliases to allow more complex operations.

You can put from a queue manager outside a cluster or put to a queue manager outside the cluster using these configuration options. The manual offers example definitions.

References to advanced tasks (continuation)

- Chapter 9 of the Queue Manager Clusters manual describes advanced tasks, including:
 - removing a cluster queue from a queue manager
 - moving a full repository to another queue manager
 - removing a queue manager from a cluster
 - removing a cluster network
- Follow the steps in the manual to properly execute these tasks! Otherwise, you may end up with problems in your cluster.



Troubleshooting tips

- **Is the repository manager still running?**
Check the AMQERRxx.log or CHIN joblog.
- **Are channels able to run in both directions?**
Display CLUSQMGR and CHSTATUS information.
- **Are the SYSTEM.CLUSTER.* queues enabled?**
Issue DISPLAY QUEUE(SYSTEM.C*) ALL
- **Are there messages on
SYSTEM.CLUSTER.COMMAND.QUEUE or
SYSTEM.CLUSTER.TRANSMIT.QUEUE?**
- **Are there duplicate QMIDs for a given QMGR?**
Issue DISPLAY CLUSQMG(*) QMID.

Troubleshooting scenarios

- **DISPLAY CLUSQMGR** may show **CLUSQMGR** names starting with **SYSTEM.TEMP**.
The queue manager has not received all necessary information from the full repository.
- **Cluster cache on z/OS fills.**
 - See technote 1240731 for a problem caused by distributed queue managers in the cluster
 - In MQ V6, you can set **CLCACHE(DYNAMIC)** to dynamically increase cache storage
 - For z/OS, is module **CSQXRRMF** at a current level?

Troubleshooting - REFRESH/RESET

- REFRESH CLUSTER and RESET CLUSTER do not need to be issued under normal circumstances.

They will not fix a cluster that has a bad configuration or lacks connectivity and can make matters worse in such a situation.



Troubleshooting - REFRESH

- REFRESH CLUSTER removes and rebuilds locally held information about a cluster.
- REFRESH CLUSTER(clustername) REPOS(NO)
- REFRESH CLUSTER(clustername) REPOS(YES) also refreshes information about full repository queue managers and may not be issued from a full repository.
- REFRESH CLUSTER(*)

Troubleshooting - RESET

- RESET CLUSTER is issued from a full repository queue manager. It forcibly removes a queue manager or specific QMID from a cluster.
 - RESET CLUSTER(*clustername*)
QMNAME(*qmname*) ACTION(FORCEREMOVE)
QUEUES(NO)
 - RESET CLUSTER(*clustername*) QMID(*qmid*)
ACTION(FORCEREMOVE) QUEUES(NO)

References – Support Website and Manuals

- The WebSphere MQ product support page:
<http://www-306.ibm.com/software/integration/wmq/support/>.
You can search for fixes and non-defect solutions.
- The MQ Library:
<http://www-306.ibm.com/software/integration/wmq/library/>
- The WebSphere MQ Queue Manager Clusters manual:
 - SC34-6061-02 for MQ 5.3.x
 - SC34-6589-00 for MQ 6.0
 - “Appendix A. Troubleshooting” lists common error symptoms and causes plus some problem scenarios with their resolutions

MustGather technote for clustering

- Technote 1178839 includes:
 - How to avoid cluster problems
 - Clustering terminology
 - Documentation required and diagnostic tips
 - Best reference material for WebSphere MQ problem determination and for this component



MustGather – data from amqrfdm

- The MustGather for Clustering provides instructions to run the utility “amqrfdm” in Unix and Windows to collect all the details for:
 - ▶ Subscriptions
 - ▶ Cluster Queues
 - ▶ Cluster Queue Managers
 - ▶ Browse Transmission Queue
 - ▶ Miscellaneous data



MustGather – data from runmqsc - cluscmd

- Also, this MustGather tells how to use “runmqsc” with a cluscmd.txt file to collect:
 - ▶ Display Cluster Queue Manager details.
 - ▶ Display Queue Manager details.
 - ▶ Display Channel details - Cluster Sender
 - ▶ Display Channel details - Cluster Receiver
 - ▶ Display Channel Status details.
 - ▶ Display Namelist details.
 - ▶ Display Queue details - cluster



Additional technotes

- 1229905: WebSphere MQ Clustering - Cluster hints and tips
- 4006367: SupportPac MD05: MQSeries - Design considerations for large Clusters
- 4006424: SupportPac MP7A: WebSphere MQ for Windows V5.3 - Performance tuning for large clusters



Additional References – Articles/Redbook

- DeveloperWorks: Enabling SSL in an existing WebSphere MQ cluster
http://www-128.ibm.com/developerworks/websphere/library/techarticles/0608_vanstone/0608_vanstone.html
- DeveloperWorks: Migrating WebSphere MQ queue manager clusters to WebSphere MQ V6
http://www-128.ibm.com/developerworks/websphere/library/techarticles/0605_vanstone/0605_vanstone.html
- Redbook: WebSphere MQ V6 Fundamentals
<http://www.redbooks.ibm.com/abstracts/SG247128.html>

Summary

- Use the WebSphere MQ Queue Manager Clusters manual as guidance for configuration, administration, and troubleshooting
- Search the support database for fixes and solutions

Additional WebSphere Product Resources

- Discover the latest trends in WebSphere Technology and implementation, participate in technically-focused briefings, webcasts and podcasts at: www.ibm.com/developerworks/websphere/community/
- Learn about other upcoming webcasts, conferences and events: www.ibm.com/software/websphere/events_1.html
- Join the Global WebSphere User Group Community: www.websphere.org
- Access key product show-me demos and tutorials by visiting IBM Education Assistant: www.ibm.com/software/info/education/assistant
- Learn about the Electronic Service Request (ESR) tool for submitting problems electronically: www.ibm.com/software/support/viewlet/ESR_Overview_viewlet_swf.html
- Sign up to receive weekly technical My support emails: www.ibm.com/software/support/einfo.html

Questions and Answers

