Ask the Experts
WebSphere MQ Failover on Distributed Platforms

30 March 2011

WebSphere® Support Technical Exchange

ON DEMAND BUSINESS™

Agenda

- Introduce the panel of experts

- Introduction

- Answer questions submitted by email (14 questions)

- Open telephone lines for questions

- Summarize highlights

Panel of Experts

Panelist	Role at IBM®
Barry Robbins	Team Lead, WebSphere MQ Distributed Support
Snezhana Johnson	WebSphere MQ Distributed Support
Tameka Woody	WebSphere MQ System i – Windows® Support
Andrew Schofield	Senior Software Engineer, WebSphere MQ
Jason Edmeades	Service Architect, WebSphere MQ
Jonathan Rumsey	Lead System i Developer, WebSphere MQ

Introduction

- We will address questions covering the following items as they apply to WebSphere MQ:

  - ▶ High availability (HA) clustering

  - ▶ Multi-instance queue managers

  - ▶ Relative advantages of each

  - ▶ Hardware and software requirements

  - ▶ Support and management

- Platforms covered will include AIX, HP-UX, Linux®, Solaris, Windows and System i.

Question 1

- What failover options are available for WebSphere MQ?

Answer to Question 1

▶ WebSphere MQ V6 and V7

- WebSphere MQ Clusters
- Backup Queue Managers
- Independent Auxiliary Storage Pools (IASP) on IBM i
- PowerHA for AIX
- Third Party High Availability Products

▶ WebSphere MQ V7

- Multi-Instance Queue Managers
- Channel and Client Reconnection

Question 2

- How should I choose between multi-instance queue managers and HA clusters?

Answer to Question 2

- Environmental complexity
  - ‣ Connected applications
  - ‣ MQ configuration
  - ‣ ip address

- Failover events

- Multiple failover ability

- Cost

Question 3

- Why was MC91 withdrawn and what replaces it?

Answer to Question 3

- Customers don't like using unsupported code to use MQ with their HA clusters
  - ▶ MC91 was provided 'as-is'

- MQ 7.0.1 can separate node-specific and shared data without needing environment variables and shell scripts
  - ▶ New DataPath attribute controlled by `crtmqm -md`
  - ▶ Much of what MC91 does is now redundant

- Each version of MQ means a new version of MC91
  - ▶ Gives customers an extra job when upgrading MQ

- Support integrated into the product would be preferable

Answer to Question 3

- MC91 has been replaced by documentation in the WebSphere MQ Information Center
  - ▶ http://publib.boulder.ibm.com/infocenter/wmqv7/v7r0/topic/com.ibm.mq.amqzag.doc/fa70230_.htm

- The scripts have been cleaned up, but are essentially equivalent

Answer to Question 3

MC91	Using MQ 7.0.1
hacrtmqm to create queue manager on shared disk and point symbolic links back to node's /var/mqm	New crtmqm –md option
halinkmqm	New addmqinf command
hadltmqm	New rmvmqinf command to remove queue manager from a node, dltmqm to delete the queue manager
hamqm_start	hamqm_start script in InfoCenter
hamqm_stop	hamqm_stop script in InfoCenter
rc.local script	No equivalent
hamqm_applmon	hamqm_applmon script in InfoCenter

Question 4

- What should I consider when setting up a multi-instance queue manager within a cluster in a mixed MQ version 6 & 7 environment?

Answer to Question 4 - Considerations

- System IP address is not carried over during a failover, which will effect communication between 7.0.1 and older queue managers

- Avoid making a multi-instance queue manager a full repository

Answer to Question 4 - Recommendations

- Configure listener as a service to start during queue manager startup

- Define only two full repositories

- Best option - upgrade all cluster members to 7.0.1
- If upgrade to 7.0.1 is not possible, use blank CONNAME and default port on a 7.0.1 multi instance queue manager cluster receiver channel

http://www.ibm.com/support/docview.wss?uid=swg27018127

Question 5

- How do you migrate or update MQ in a multi-instance or HA environment?

Answer to Question 5

- Ensure a backup is available
  - ▶ tar – zip – savf of the queue manager data
  - ▶ Use MS03 to save queue manager definitions

- Failover to standby node/system

- Update primary node/system

- Fail back to primary node/system

- Update backup node

Question 6

- If we are looking for to share/run WMQ (HUB) on multiple platforms using the new capabilities added in MQV7 (i.e. MultiInstances), is it necessary to recommend/use an external NFS storage device, in order to meet the customer expectations in terms of HA?

- And if yes, how do we integrate them?

Answer to Question 6

- Queue managers not shareable on different platforms

- External storage is best option, but not required

- Integration depends on OS platforms involved
 - ‣ UNIX® – NFS v4
  - ‣ Windows/System i - CIFS

Question 7

- ■ What needs to be changed within Windows registry to allow for the proper startup of WMQ Services on a "backup server" when manually failing over outside the control of something like MSCS?

  - ▸ Given a configuration of two servers (A & B) both having WMQ installed and having used the hamvmqm utility to move the qmgr to a "common" drive that is only mapped on the "active" server - what needs to happen on server B within the registry to ensure that WMQ services start up automatically?

Answer to Question 7

- MQ Version 7.0.1 introduced the Multi-Instance support.

- This means both nodes can have the standard MQ service running, both 'starting' the queue manager, but one actually running it whilst the other is waiting to take-over.

- Very little administrative overhead once it is setup in standby mode, and takeover can happen really quickly.

- Does require both machines to be Domain Controllers

- Pre-7.0.1 instructions are more complex (see later foil).

Answer to Question 7 - Multi Instance

- One server creates the qmgr on the shared disk
  - crtmqm –ld ... –md ...

- The qmgr details are added to the other server
  - dspmqinf / addmqinf

- Both start 'permitting a standby'
  - strmqm –x

- One starts the queue manager

  ▶ The other sits waiting

  ▶ ... For failure or manual switch

- (http://publib.boulder.ibm.com/infocenter/wmqv7/v7r0/topic/com.ibm.mq.amqzag.doc/fa70150_.htm)
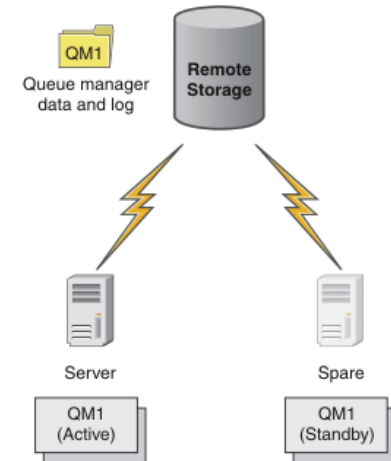
Answer to Question 7 - Pre-Multi instance

- Pre-v7.0.1 there is no supplied way to support manual failover between 2 Windows machines outside of MSCS

- Manual registry migration (with MQ (amqmsrvn) stopped)
  - HKLM\SOFTWARE\IBM\MQSeries\CurrentVersion\Configuration\QueueManager\QM
  - HKLM\SOFTWARE\IBM\MQSeries\CurrentVersion\Configuration\Services\QM
  - HKLM\SOFTWARE\IBM\MQSeries\CurrentVersion\Status\QueueManager\QM

- Giving the "local mqm" group 'full control' over all files and directories in the log/data tree on the new machine.

- Almost equivalent to backup/restore on a different node.

Question 8

- We are running SLES 10.3 on a distributed platform. Is SLES 11.x required to support multi-instance?

- Also, is it correct to assume that it doesn't work with NFS v3 due to the auto-locking capabilities enabled by NFS v4?

Answer to Question 8

- No, SLES 11 is not required to use multi-instance queue managers.

- If you are going to use NFS v4, you will need at least SLES 10.

- SLES 10 Update 3 introduced a suspected problem in the NFS v4 server which prevents correct operation of multi-instance queue managers. The problem was rectified in kernel level 2.6.16.60-0.60.1.

Answer to Question 8

- It is correct that multi-instance queue managers do not work with NFS v3.

- This is due to the locking behavior of NFS v3.

- NFS v4 has lease-based locks which mean they are automatically released if the NFS client fails.

Question 9

- I have an iSeries, MQ version 7, multi-instance queue manager wishing to join cluster, where the full repository is mainframe, MQ version 6. The current security policy prevents the use of default port (1414). Our conclusion at the moment is that this isn't possible to configure (making use of the auto switching). Our current plan is to set the cluster receiver channel CONNAME to be the DNS(port) of the system where the queue manager is running. E.g. if the iASP is active on PROD, then I will set the CONNAME of the cluster receiver to "prod(99999)". Then, when the iASP is varied onto the CONT iSeries, we will set the CONNAME to "cont(99999)", & perform a "refresh cluster" (we have found that the refresh is required, please advise if you think that this is not the case ?). Can you please provide some inputs on the scenario above?

Answer to Question 9

- Assuming you have no immediate plans to bring the full repository to V7.0.1, this is recognised as being an issue

- Manually altering the CONNAME of the cluster receiver after switchover will work, though REFRESH CLUSTER should not be required under normal circumstances (Channel retry should discover the new address)

- Using Cluster Resource Groups to handle switchover of the iASP and IP address together on IBM i is an option to consider

- Or, use a hostname and update the DNS entry

Answer to Question 9

- Another alternative...

- Use of port 1414 can be avoided by setting the "Port" attribute under the "TCP" stanza in qm.ini to a value other than 1414, e.g. 99999
- This would need to be changed for each queue manager in the cluster, leaving the CONNAME blank

Answer to Question 9

- Or...

- Writing a message exit to alter the CONNAME might be another option. The MR01 SupportPac could be used as a basis for writing such an exit

- The MR01 Support Pac is also useful as-is connecting pre V7.0.1 queue managers to multi-instance queue managers
  - ▸ For example WebSphere MQ for HP OVMS and HP NonStop Server

Question 10

- Is using multiple IP addresses in CONNAME to connect to a remote queue manager supported when they are not configured as multi-instance queue managers?

Answer to Question 10

- ■ Yes, it is possible to supply multiple connection names in the CONNAME parameter on channels when the queue managers are not configured as multi-instance queue managers

- ■ Remember to separate connections using the comma as the delimiter

- ■ The CONNAME field can hold up to 264 characters

Question 11

- Does CONNAME support multiple addresses in WMQ for z/OS?

Answer to Question 11

- CONNAMES with multiple IP addresses are supported on z/OS

  ▸ CONNAME attribute has a 48 character limit

    • Long hostnames may not fit

  ▸ Workaround

    • Set up DNS servers to use hostname "myserver" instead of   "myserver.location.company.com

    • Use IP addresses

Question 12

- What kind of events cause a client connection switch from an active IP address to another.

Answer to Question 12

- Long term connection issues

- Primary queue manager unavailable

- Collect MQ documentation for review if root cause analysis is desired

http://www.ibm.com/support/docview.wss?uid=swg21229861

Question 13

- I am running V6 queue managers on HP-UX. HA configuration was done using the MC91 support pack. We use the Veritas (VCS) agent that is provided in MC91.

- I understand that MC91 is fully embedded in MQ V7 and that it has much more than MC91, but how do I migrate my existing V6 queue managers to V7? Also, optionally, is it possible to continue using the existing configuration of MC91 but still upgrade to V7?

Answer to Question 13

- MC91 was upgraded when MQ v7 was released.

- Follow the migration instructions for v6 to v7 in the MC91 documentation.

- You can continue to use MC91 with current versions of MQ, but since MQ v7.0.1 there are less convoluted alternatives for configuration using an HA cluster.

Question 14

- In the scenario where I create brand new queue managers in V7 – let us say I create them as multi-instance queue managers using the new flags as documented. My data disk is different than /var. I understand MQ will monitor and offline the queue managers as necessary. In case of failover, how will MQ bring the virtual IP and disk online on the second node? Will it have enough permissions to do that? Or is it necessary to give root authority to the MQM id?

Answer to Question 14

- If you're using multi-instance queue managers, the queue manager data is held on networked storage which is attached to both servers hosting instances of the queue manager. There is no need to move the disks from one server to another as required with an HA cluster.

- A multi-instance queue manager changes IP address when it fails over. Consequently, the issue about permissions to move a virtual IP address does not arise.

Open Lines for Questions

We Want to Hear From You!

Tell us about what you want to learn

Suggestions for future topics
Improvements and comments about our webcasts
We want to hear everything you have to say!

Please send your suggestions and comments to:
wsehelp@us.ibm.com

Summary

Additional WebSphere MQ HA links

- WebSphere MQ - Understanding High Availability Options
  http://www.ibm.com/support/docview.wss?uid=swg27019872

- WebSphere MQ V7.0 Client Enhancements
  http://www.ibm.com/support/docview.wss?uid=swg27016801

- Introduction to WebSphere MQ Multi-Instance Queue Mangers
  http://www.ibm.com/support/docview.wss?uid=swg27017382

- Implementation Considerations for Multi-Instance Queue Managers in
  WebSphere MQ Cluster Environment
  http://www.ibm.com/support/docview.wss?uid=swg27018127

- Checklist for Implementing High Availability using Multi-Instances in Message
  Broker v7
  http://www.ibm.com/support/docview.wss?uid=swg27017847

Additional WebSphere Product Resources

- Learn about upcoming WebSphere Support Technical Exchange webcasts, and access previously recorded presentations at:
http://www.ibm.com/software/websphere/support/supp_tech.html

- Discover the latest trends in WebSphere Technology and implementation, participate in technically-focused briefings, webcasts and podcasts at:
http://www.ibm.com/developerworks/websphere/community/

- Join the Global WebSphere Community:
http://www.websphereusergroup.org

- Access key product show-me demos and tutorials by visiting IBM Education Assistant:
http://www.ibm.com/software/info/education/assistant

- View a webcast replay with step-by-step instructions for using the Service Request (SR) tool for submitting problems electronically:
http://www.ibm.com/software/websphere/support/d2w.html

- Sign up to receive weekly technical My Notifications emails:
http://www.ibm.com/software/support/einfo.html