Oracle® Solaris Cluster Data Service for IBM WebSphere MQ Guide

Part No: F24929
November 2019

ORACLE

Oracle Solaris Cluster Data Service for IBM WebSphere MQ Guide

Part No: F24929

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` or visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` if you are hearing impaired.

Référence: F24929

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, accorder de licence, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique :

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer un risque de dommages corporels. Si vous utilisez ce logiciel ou ce matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour des applications dangereuses.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée de The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers, sauf mention contraire stipulée dans un contrat entre vous et Oracle. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` ou le site `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` si vous êtes malentendant.

Contents

Using This Documentation

- **Overview** – Describes how to install and configure the Oracle Solaris Cluster HA for WebSphere MQ data service.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Oracle Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Oracle Solaris Cluster software.

Product Documentation Library

Documentation and resources for this product and related products are available at http://www.oracle.com/pls/topic/lookup?ctx=product_intuitive_ID.

Feedback

Provide feedback about this documentation at http://www.oracle.com/goto/docfeedback.

♦ ♦ ♦ CHAPTER 1

Installing and Configuring HA for WebSphere MQ

This chapter explains how to install and configure HA for WebSphere MQ.

This chapter contains the following sections.

HA for WebSphere MQ Overview

The HA for WebSphere MQ data service provides a mechanism for the orderly startup and shutdown, fault monitoring, and automatic failover of the IBM MQ service.

The following components can be protected by the HA for WebSphere MQ data service within the global zone or a zone cluster.

- Queue Manager
- Channel Initiator

- Command Server
- Listener
- Trigger Monitor

Overview of Installing and Configuring HA for WebSphere MQ

The following table summarizes the tasks for installing and configuring HA for WebSphere MQ and provides cross-references to detailed instructions for performing these tasks. Perform the tasks in the order that they are listed in the table.

TABLE 1 Tasks for Installing and Configuring HA for WebSphere MQ

Task	Instructions
Plan the installation	"Planning the HA for WebSphere MQ Installation and Configuration" on page 10
Install and configure the IBM MQ software	"How to Install and Configure IBM MQ" on page 14
Verify the installation and configuration	"How to Verify the Installation and Configuration of IBM MQ" on page 20
Install HA for WebSphere MQ packages	"How to Install the HA for WebSphere MQ Package" on page 23
Register and configure HA for WebSphere MQ resources	"How to Register and Configure HA for WebSphere MQ" on page 24
Verify the HA for WebSphere MQ installation and configuration	"How to Verify the HA for WebSphere MQ Installation and Configuration" on page 28
Upgrade the HA for WebSphere MQ data service	"Upgrading HA for WebSphere MQ" on page 28
Tune the HA for WebSphere MQ fault monitor	"Understanding the HA for WebSphere MQ Fault Monitor" on page 29
Debug HA for WebSphere MQ	"How to Turn on Debugging for HA for WebSphere MQ" on page 31

Planning the HA for WebSphere MQ Installation and Configuration

This section contains the information you need to plan your HA for WebSphere MQ installation and configuration.

Configuration Restrictions

The configuration restrictions in the subsections that follow apply only to HA for WebSphere MQ.

Caution - Your data service configuration might not be supported if you do not observe these restrictions.

Restriction for the Supported Configurations of HA for WebSphere MQ

The HA for WebSphere MQ data service can only be configured as a failover service.

Single or multiple queue managers can be deployed in the cluster.

Multiple installations are not currently supported by HA for WebSphere MQ. Only a primary installation is supported on `/opt/mqm`. After installing IBM WebSphere v8.0 or later, you must also execute the WebSphere MQ `setmqinst` script on each of the cluster node as follows:

```
# /opt/mqm/bin/setmqinst -i -p /opt/mqm
```

IBM MQ can be deployed in either the global zone or the zone cluster. See “Restriction for Multiple IBM MQ Instances” on page 11 for more information.

Restriction for the Location of IBM MQ Files

The IBM MQ files are where the queue manager data files `/var/mqm/qmgr/`*queue-manager* and `/var/mqm/log/`*queue-manager* are stored.

These IBM MQ files needs to be placed on shared storage as either a cluster file system or a highly available local file system.

Restriction for Multiple IBM MQ Instances

The HA for WebSphere MQ data service can support multiple IBM MQ instances, potentially with different versions.

If you intend to deploy multiple IBM MQ instances with different versions you will need to consider deploying IBM MQ in separate zone clusters.

Configuration Requirements

The configuration requirements in this section apply only to HA for WebSphere MQ.

Caution - If your data service configuration does not conform to these requirements, the data service configuration might not be supported.

Determine Which Oracle Solaris Zone IBM MQ Will Use

Oracle Solaris Zones provides a means of creating virtualized operating system environments within an instance of the Oracle Solaris 11 OS. Oracle Solaris Zones allow one or more applications to run in isolation from other activity on your system. For complete information about installing and configuring Oracle Solaris Zones, see *Creating and Using Oracle Solaris Zones*.

You must determine which zone IBM MQ will run in. IBM MQ can run within a global zone or a zone cluster.

Requirements If Multiple IBM MQ Instances Are Deployed on Cluster File Systems

If a cluster file system is being used for the IBM MQ files, it is possible to manually start the queue manager on one node of the cluster and at the same time to also manually start the same queue manager on another node of the cluster.

Note - Although it is possible, you should not attempt this as doing so will cause severe damage to the IBM MQ files.

Although it is expected that no-one will manually start the same queue manager on separate nodes of the cluster at the same time the HA for WebSphere MQ provides a mechanism to prevent someone from doing so, albeit by mistake.

To prevent against this happening you must implement one of the following two solutions.

1. Use a highly available local file system for the IBM MQ files.

   This is the recommended approach as the IBM MQ files would be mounted only on one node of the cluster at a time. This then limits starting the queue manager on only one node of the cluster at a time.

2. Create a symbolic link for `/opt/mqm/bin/strmqm` and `/opt/mqm/bin/endmqm` to `/opt/SUNWscmqs/mgr/bin/check-start`.

   `/opt/SUNWscmqs/mgr/bin/check-start` provides a mechanism to prevent manually starting or stopping the queue manager, by verifying that the start or stop is being attempted by the HA for WebSphere MQ data service.

   `/opt/SUNWscmqs/mgr/bin/check-start` will report the following error if an attempt to manually start or stop the queue manager.

```
$ strmqm qmgr1
$ Request to run </usr/bin/strmqm qmgr1> within Oracle Solaris Cluster has been
 refused
```

   If a cluster file system is used for the IBM MQ files, you must create a symbolic link for `strmqm` and `endmqm` to `/opt/SUNWscmqs/mgr/bin/check-start` and inform the HA for WebSphere MQ data service of this change.

   To do this, you must perform the following on each node of the cluster.

```
# cd /opt/mqm/bin
#
# mv strmqm strmqm_sc3
# mv endmqm endmqm_sc3
#
# ln -s /opt/SUNWscmqs/mgr/bin/check-start strmqm
# ln -s /opt/SUNWscmqs/mgr/bin/check-start endmqm
```

   After renaming `strmqm` and `endmqm` you must use these new program names (`strmqm_sc3` and `endmqm_sc3`) for the `START_CMD` and `STOP_CMD` variables when you edit the `/opt/SUNWscmqs/mgr/util/mgr_config` file in Step 7 in "How to Register and Configure HA for WebSphere MQ" on page 24.

Note - If you implement this workaround, then you must back it out whenever you need to apply any maintenance to IBM MQ. Afterwards, you must again apply this workaround.

Instead the recommended approach is to use a highly available local file system for the IBM MQ files.

Installing and Configuring IBM MQ

This section contains the procedures you need to install and configure IBM MQ.

How to Install and Configure IBM MQ

This section contains the procedures you need to install and configure IBM MQ.

1. **Determine how many IBM MQ instances will be used.**

   Refer to "Restriction for Multiple IBM MQ Instances" on page 11 for more information.

2. **Determine which Oracle Solaris Zone to use.**

   Refer to "Determine Which Oracle Solaris Zone IBM MQ Will Use" on page 12 for more information.

3. **Determine how IBM MQ should be deployed in the cluster.**

   IBM MQ can be deployed onto a cluster file system or highly available file system on the cluster.

`/var/mqm`	Can be deployed on a cluster file system, highly available local file system or on local storage on each cluster node. It is recommended to deploy `/var/mqm` on local storage on each cluster node.
`/var/mqm/qmgrs/`*queue-manager* and `/var/mqm/log/`*queue-manager*	Can be deployed on a cluster file system or highly available local file system. It is recommended to deploy `/var/mqm/qmgrs/`*queue-manager* and `/var/mqm/log/`*queue-manager* on highly available local file system.

4. **Create a cluster file system or highly available local file system for the IBM MQ files.**

   Within this step you will create file systems for the IBM MQ files and `/var/mqm`. Once you have determined how IBM MQ should be deployed in the cluster, you can choose one of the sub steps below.

   - Create the IBM MQ files and `/var/mqm` on cluster file systems by using Step 4a.
   - Create the IBM MQ files on SVM highly available local file systems and `/var/mqm` on cluster file system by using Step 4b.
   - Create the IBM MQ files on ZFS highly available local file systems and `/var/mqm` on local storage by using Step 4c.

   a. **IBM MQ files and `/var/mqm` on cluster file systems.**

Within this deployment:

- The IBM MQ files are deployed on cluster file systems.
- The IBM MQ instances are `qmgr1` and `qmgr2`.
- `/var/mqm` uses a cluster file system with a symbolic link for `/var/mqm/qmgrs/@SYSTEM` to a local file (`/var/mqm_local/qmgrs/@SYSTEM`) on each node in the cluster.

Note - Refer to Step 4d for more information about setting up this symbolic link.

```
# ls -l /var/mqm
lrwxrwxrwx   1 root     other         11 Jan  8 14:17 /var/mqm ->
 /global/mqm
#
# ls -l /global/mqm/qmgrs
total 6
lrwxrwxrwx   1 root      other          512 Dec 16 09:57 @SYSTEM ->
 /var/mqm_local/qmgrs/@SYSTEM
drwxr-xr-x   4 root     root         512 Dec 18 14:20 qmgr1
drwxr-xr-x   4 root     root         512 Dec 18 14:20 qmgr2
#
# ls -l /global/mqm/log
total 4
drwxr-xr-x   4 root     root         512 Dec 18 14:20 qmgr1
drwxr-xr-x   4 root     root         512 Dec 18 14:20 qmgr2
#
# more /etc/vfstab (Subset of the output)
/dev/md/dg_d4/dsk/d40   /dev/md/dg_d4/rdsk/d40  /global/mqm
     ufs     3       yes     logging,global
/dev/md/dg_d4/dsk/d43   /dev/md/dg_d4/rdsk/d43  /global/mqm/qmgrs/qmgr1
 ufs     4       yes     logging,global
/dev/md/dg_d4/dsk/d46   /dev/md/dg_d4/rdsk/d46  /global/mqm/log/qmgr1
   ufs     4       yes     logging,global
/dev/md/dg_d5/dsk/d53   /dev/md/dg_d5/rdsk/d53  /global/mqm/qmgrs/qmgr2
 ufs     4       yes     logging,global
/dev/md/dg_d5/dsk/d56   /dev/md/dg_d5/rdsk/d56  /global/mqm/log/qmgr2
   ufs     4       yes     logging,global
```

b. IBM MQ files on SVM highly available local file systems and `/var/mqm` on cluster file system.

Within this deployment:

- The IBM MQ files are deployed on SVM highly available local file systems.
- The IBM MQ instances are `qmgr1` and `qmgr2`.

- `/var/mqm` uses a cluster file system with a symbolic link for `/var/mqm/qmgrs/@SYSTEM` to a local file (`/var/mqm_local/qmgrs/@SYSTEM`) on each node in the cluster.

Note - Refer to Step 4d for more information about setting up this symbolic link.

```
# ls -l /var/mqm
lrwxrwxrwx   1 root     other         11 Sep 17 16:53 /var/mqm ->
 /global/mqm
#
# ls -l /global/mqm/qmgrs
total 6
lrwxrwxrwx   1 root      other          512 Sep 17 09:57 @SYSTEM ->
 /var/mqm_local/qmgrs/@SYSTEM
lrwxrwxrwx   1 root     other         22 Sep 17 17:19 qmgr1 ->
 /local/mqm/qmgrs/qmgr1
lrwxrwxrwx   1 root     other         22 Sep 17 17:19 qmgr2 ->
 /local/mqm/qmgrs/qmgr2
#
# ls -l /global/mqm/log
total 4
lrwxrwxrwx   1 root     other         20 Sep 17 17:18 qmgr1 ->
 /local/mqm/log/qmgr1
lrwxrwxrwx   1 root     other         20 Sep 17 17:19 qmgr2 ->
 /local/mqm/log/qmgr2
#
# more /etc/vfstab (Subset of the output)
/dev/md/dg_d4/dsk/d40   /dev/md/dg_d4/rdsk/d40  /global/mqm
     ufs     3       yes     logging,global
/dev/md/dg_d4/dsk/d43   /dev/md/dg_d4/rdsk/d43  /local/mqm/qmgrs/qmgr1
 ufs     4       no      logging
/dev/md/dg_d4/dsk/d46   /dev/md/dg_d4/rdsk/d46  /local/mqm/log/qmgr1
   ufs     4       no      logging
/dev/md/dg_d5/dsk/d53   /dev/md/dg_d5/rdsk/d53  /local/mqm/qmgrs/qmgr2
 ufs     4       no      logging
/dev/md/dg_d5/dsk/d56   /dev/md/dg_d5/rdsk/d56  /local/mqm/log/qmgr2
   ufs     4       no      logging
```

c. IBM MQ files on ZFS highly available local file systems and `/var/mqm` on local storage.

Within this deployment:

- The IBM MQ files are deployed on ZFS highly available local file systems.
- The IBM MQ instances are `qmgr1` and `qmgr2`.
- `/var/mqm` uses local storage on each cluster node.

As `/var/mqm` is on a local file system you must copy `/var/mqm/mqs.ini` from the node where the queue managers was created to all other nodes in the cluster where the queue manager will run.

Note - Refer to Step 8 for more information about copying `/var/mqm/mqs.ini`.

```
# df -k /var/mqm
Filesystem            kbytes    used   avail capacity  Mounted on
/                   59299764 25657791 33048976    44%    /
#
# ls -l /var/mqm/qmgrs
total 6
drwxrwsr-x   2 mqm      mqm          512 Sep 11 11:42 @SYSTEM
lrwxrwxrwx   1 mqm      mqm           14 Sep 11 11:45 qmgr1 -> /ZFSwmq1/qmgrs
lrwxrwxrwx   1 mqm      mqm           14 Sep 11 11:50 qmgr2 -> /ZFSwmq2/qmgrs
#
# ls -l /var/mqm/log
total 4
lrwxrwxrwx   1 mqm      mqm           12 Sep 11 11:44 qmgr1 -> /ZFSwmq1/log
lrwxrwxrwx   1 mqm      mqm           12 Sep 11 11:54 qmgr2 -> /ZFSwmq2/log
#
# df -k /ZFSwmq1
Filesystem            kbytes    used   avail capacity  Mounted on
HAZpool1             4096453   13180 4083273     1%    /ZFSwmq1
#
# df -k /ZFSwmq2
Filesystem            kbytes    used   avail capacity  Mounted on
HAZpool2             4096453   13133 4083320     1%    /ZFSwmq2
```

d. Cluster file system is used for `/var/mqm`.

Within this deployment:

- If `/var/mqm` is placed on shared storage as a cluster file system, a symbolic link is made from `/var/mqm/qmgrs/@SYSTEM` to local file `/var/mqm_local/qmgrs/@SYSTEM`.
- You must perform this step on all nodes in the cluster *only* if `/var/mqm` is a cluster file system.

```
# mkdir -p /var/mqm_local/qmgrs/@SYSTEM
# mkdir -p /var/mqm/qmgrs
# ln -s /var/mqm_local/qmgrs/@SYSTEM /var/mqm/qmgrs/@SYSTEM
```

This restriction is required because IBM MQ uses keys to build internal control structures. Mounting `/var/mqm` as a cluster file system with a symbolic link for `/var/mqm/qmgrs/`

`@SYSTEM` to a local file ensures that any derived shared memory keys are unique on each node.

If multiple queue managers are required and your queue manager was created before you setup a symbolic link for `/var/mqm/qmgrs/@SYSTEM`, you must copy the contents, with permissions, of `/var/mqm/qmgrs/@SYSTEM` to `/var/mqm_local/qmgrs/@SYSTEM` before creating the symbolic link.

You must stop all queue managers before doing this and perform this on each node of the cluster.

```
# mkdir -p /var/mqm_local/qmgrs/@SYSTEM
# cd /var/mqm/qmgrs
# cp -rp @SYSTEM/* /var/mqm_local/qmgrs/@SYSTEM
# rm -r @SYSTEM
# ln -s /var/mqm_local/qmgrs/@SYSTEM @SYSTEM
```

5. Mount the highly available local file system.

Perform this step on one node of the cluster.

a. If a non-ZFS highly available file system is being used for the IBM MQ files.

Ensure the node has ownership of the disk set or disk group.

For Solaris Volume Manager:

```
# metaset -s disk-set -t
```

i. If the global zone is being used for IBM MQ.

```
# mount websphere-mq-highly-available-local-file-system
```

ii. If a zone cluster is being used for IBM MQ.

Create the mount point on all zones of the cluster that are being used for IBM MQ.

Mount the highly available local file system on one of the zones being used .

```
# zlogin zonename mkdir websphere-mq-highly-available-local-file-system
#
# mount -F lofs websphere-mq-highly-available-local-file-system \
/zonepath/root/websphere-mq-highly-available-local-file-system
```

b. If a ZFS highly available file system is being used for IBM MQ.

```
# zpool export -f HAZpool
# zpool import -R /zonepath/root HAZpool
```

6. **Install IBM MQ on all nodes of the cluster.**

   After you have created and mounted the appropriate file systems for the IBM MQ files and `/var/mqm`, you must install IBM MQ on *each* node of the cluster, either in the global zone or zone cluster, as required.

   Follow the *IBM MQ* documentation to install IBM MQ.

   Note - You may choose to locate the mqm userid and group within `/etc/passwd` and `/etc/group` or within a name service such as NIS or NIS+. However, as the HA for WebSphere MQ uses the `su user` command to start, stop and probe IBM MQ, it is recommend that the mqm userid/group is located within `/etc/passwd` and `/etc/group` in the cluster. This is to ensure that the `su(1M)`command is not impacted if a name service such as NIS or NIS+ is unavailable.

   If you choose to locate the mqm userid/group within a network information name service such as NIS or NIS+, IBM MQ maybe affected if the network information name service is unavailable.

   ---

7. **Create the IBM MQ queue manager.**

   Follow the *IBM MQ* documentation to create a queue manager.

8. **If a local file system is used for `/var/mqm`, copy `/var/mqm/mqs.ini` to all nodes of the cluster.**

   Within this deployment:

   - If `/var/mqm/mqs.ini` is placed on local storage as a local file system, you must copy `/var/mqm/mqs.ini` from the node where the queue manager was created to all other nodes in the cluster where the queue manager will run.
   - You must perform this step on all nodes in the cluster *only* if `/var/mqm` is a local file system.

   a. **If the global zone is being used for IBM MQ.**

      ```
      # rcp /var/mqm/mqs.ini remote-node:/var/mqm/mqs.ini
      ```

   b. **If a zone cluster is being used for IBM MQ.**

      ```
      # rcp /zonepath/root/var/mqm/mqs.ini \
      remote-node:/zonepath/root/var/mqm/mqs.ini
      ```

Verifying the Installation and Configuration of IBM MQ

This section contains the procedure you need to verify the installation and configuration.

▼ How to Verify the Installation and Configuration of IBM MQ

This procedure does not verify that your application is highly available because you have not yet installed your data service.

Perform this procedure on one node of the cluster unless a specific steps indicates otherwise.

1. **If using a zone cluster, ensure that the zone is booted and running.**

   ```
   # clzonecluster status
   ```

   If required, boot the zone cluster.

   ```
   # clzonecluster boot zone-cluster-name
   ```

2. **Login to the zone, if a zone cluster is being used.**

   ```
   # zlogin zonename
   ```

3. **Start the queue manager, create a persistent queue and put a test message to that queue.**

   ```
   # su - mqm
   $ strmqm queue-manager
   $ runmqsc queue-manager
   def ql(sc3test) defpsist(yes)
   end
   $
   $ /opt/mqm/samp/bin/amqsput SC3TEST queue-manager
   test test test test test
   ^C
   ```

4. **Stop the queue manager.**

   ```
   $ endmqm -i queue-manager
   $ exit
   ```

5. **Logout from the zone, if a zone cluster is being used.**

```
# exit
```

6. **Unmount the highly available local file system.**

   Perform this step in the global zone only.

   You should unmount the highly available file system you mounted in Step 5 in "How to Install and Configure IBM MQ" on page 14.

   a. **If a non-ZFS highly available local file system is being used for IBM MQ.**

      i. **If the global zone is being used for IBM MQ.**

      ```
      # umount websphere-mq-highly-available-local-file-system
      ```

      ii. **If a zone cluster is being used for IBM MQ.**

      Unmount the highly available local file system from the zone.

      ```
      # umount /zonepath/root/websphere-mq-highly-available-local-file-system
      ```

   b. **If a ZFS highly available file system is being used for IBM MQ.**

   ```
   # zpool export -f HAZpool
   ```

7. **Relocate the shared storage to other node.**

   Perform this step on another node of the cluster.

   a. **If a non-ZFS highly available local file system is being used for the IBM MQ files.**

      Ensure the node has ownership of the disk set or disk group.

      For Solaris Volume Manager.

      ```
      # metaset -s disk-set -t
      ```

      i. **If the global zone is being used for IBM MQ.**

      ```
      # mount websphere-mq-highly-available-local-file-system
      ```

      ii. **If a zone cluster is being used for IBM MQ.**

      Create the mount point on all zones of the cluster that are being used for IBM MQ.

      Mount the highly available local file system on one of the zones being used .

      ```
      # zlogin zonename mkdir websphere-mq-highly-available-local-file-system
      ```

```
#
# mount -F lofs websphere-mq-highly-available-local-file-system \
/zonepath/root/websphere-mq-highly-available-local-file-system
```

b. If a ZFS highly available file system is being used for IBM MQ.

```
# zpool import -R /zonepath/root HAZpool
```

8. Login to the zone, if a zone cluster is being used.

Perform this step on the other node of the cluster.

```
# zlogin zonename
```

9. Start the queue manager, get the test message and delete the queue.

Perform this step on the other node of the cluster.

```
# su - mqm
$ strmqm queue-manager
$ /opt/mqm/samp/bin/amqsget SC3TEST queue-manager
^C
$ runmqsc queue-manager
delete ql(sc3test)
end
```

10. Stop the queue manager.

Perform this step on the other node of the cluster.

```
$ endmqm -i queue-manager
$ exit
```

11. Logout from the zone, if a zone cluster is being used.

```
# exit
```

12. Unmount the highly available local file system.

Perform this step in the global zone only.

You should unmount the highly available file system you mounted in Step 5 in "How to Install and Configure IBM MQ" on page 14.

a. If a non-ZFS highly available local file system is being used for IBM MQ.

i. If the global zone is being used for IBM MQ.

```
# umount websphere-mq-highly-available-local-file-system
```

ii. **If a zone cluster is being used for IBM MQ.**

Unmount the highly available local file system from the zone.

```
# umount /zonepath/root/websphere-mq-highly-available-local-file-system
```

b. **If a ZFS highly available file system is being used for IBM MQ.**

```
# zpool export -f HAZpool
```

Installing the HA for WebSphere MQ Packages

If you did not install the HA for WebSphere MQ packages during your initial Oracle Solaris Cluster installation, perform this procedure to install the packages. To install the packages, use the `installer` program.

Note - You need to install the HA for WebSphere MQ packages in the global cluster and not in the zone cluster.

▼ How to Install the HA for WebSphere MQ Package

Perform this procedure on each cluster node where you want the HA for WebSphere MQ software to run.

1. **On the cluster node where you are installing the data service package, assume the `root` role.**

2. **Ensure that the data service package is available from the configured publisher and that the `solaris` and `ha-cluster` publishers are valid.**

```
# pkg list -a ha-cluster/data-service/websphere-message-queue
# pkg publisher
PUBLISHER                             TYPE     STATUS   P  LOCATION
solaris                               origin   online   F  solaris-repository
ha-cluster                            origin   online   F  ha-cluster-repository
```

For information about setting the `solaris` publisher, see "Adding, Modifying, or Removing Package Publishers" in *Updating Systems and Adding Software in Oracle Solaris 11.4*.

Tip - Use the `-nv` options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the `-nv` options, run the command again without the `-n` option to actually perform the installation or update. If you do get error messages, run the command again with more `-v` options (for example, `-nvv`) or more of the package FMRI pattern to get more information to help you diagnose and fix the problem. For troubleshooting information, see Appendix A, "Troubleshooting Package Installation and Update," in *Updating Systems and Adding Software in Oracle Solaris 11.4*.

3. **Install the HA for WebSphere MQ software package.**

```
# pkg install ha-cluster/data-service/websphere-message-queue
```

4. **Verify that the package installed successfully.**

```
$ pkg info ha-cluster/data-service/websphere-message-queue
```

Installation is successful if output shows that `State` is `Installed`.

5. **Perform any necessary updates to the Oracle Solaris Cluster software.**

For instructions on updating your software, see *Updating Your Oracle Solaris Cluster 4.4 Environment*.

Registering and Configuring HA for WebSphere MQ

This section contains the procedures you need to configure HA for WebSphere MQ.

Some procedures within this section require you to use certain Oracle Solaris Cluster commands. Refer to the relevant Oracle Solaris Cluster command man page for more information about these command and their parameters.

How to Register and Configure HA for WebSphere MQ

Determine if a single or multiple IBM MQ instances will be deployed.

Refer to "Restriction for Multiple IBM MQ Instances" on page 11 to determine how to deploy a single or multiple IBM MQ instances.

How to Register and Configure HA for WebSphere MQ in a Failover Resource Group

This procedure assumes that you installed the data service packages during your initial Oracle Solaris Cluster installation.

If you did not install the HA for WebSphere MQ packages as part of your initial Oracle Solaris Cluster installation, go to "How to Install the HA for WebSphere MQ Package" on page 23.

Note - Perform this procedure on one node of the cluster only.

---

1. **On a cluster member, become an administrator that provides `solaris.cluster.modify` authorization.**

2. **Register the following resource types.**

```
# clresourcetype register SUNW.HAStoragePlus
# clresourcetype register SUNW.gds
```

3. **Create a failover resource group for IBM MQ.**

Note - Refer to "Restriction for Multiple IBM MQ Instances" on page 11 for more information on the *nodelist* entry.

```
# clresourcegroup create -n nodelist websphere-mq-resource-group
```

4. **Create a resource for the IBM MQ Logical Hostname.**

```
# clreslogicalhostname create -g websphere-mq-resource-group \
-h websphere-mq-logical-hostname \
websphere-mq-logical-hostname-resource
```

5. **Create a resource for the IBM MQ Disk Storage.**

   a. **If a ZFS highly available local file system is being used.**

```
# clresource create -g websphere-mq-resource-group  \
-t SUNW.HAStoragePlus \
```

```
-p Zpools=websphere-mq-zspool \
websphere-mq-hastorage-resource
```

b. If a cluster file system or a non-ZFS highly available local file system is being used.

```
# clresource create -g websphere-mq-resource-group  \
-t SUNW.HAStoragePlus \
-p FilesystemMountPoints=websphere-mq-filesystem-mountpoint \
websphere-mq-hastorage-resource
```

6. Bring online the failover resource group for IBM MQ that now includes the Logical Hostname and Disk Storage resources.

```
# clresourcegroup online -M websphere-mq-resource-group
```

7. Create a resource for the IBM MQ queue manager.

Edit `/opt/SUNWscmqs/mgr/util/mgr_config` and follow the comments within that file. After you have edited `mgr_config`, you must register the resource.

```
# cd /opt/SUNWscmqs/mgr/util
# vi mgr_config
# ./mgr_register
```

The following example shows the `/opt/SUNWscmqs/mgr/util/mgr_config` file that has been edited to configure a queue manager resource.

```
# cat > /var/tmp/mgr1_config <<-EOF
# +++ Required parameters +++
RS=wmq1-qmgr
RG=wmq1-rg
QMGR=qmgr1
LH=wmq1-lh
HAS_RS=wmq1-ZFShas
LSR_RS=
CLEANUP=YES
SERVICES=NO
USERID=mqm

# +++ Optional parameters +++
DB2INSTANCE=
ORACLE_HOME=
ORACLE_SID=
START_CMD=
STOP_CMD=
RS_ZONE=
PROJECT=default
```

```
TIMEOUT=300
EOF

# /opt/SUNWscmqs/mgr/util/mgr_register -f /var/tmp/mgr1_config
```

8. **Enable the resource.**

```
# clresource enable websphere-mq-resource
```

9. **Create and register a resource for any other IBM MQ components.**

Repeat this step for each IBM MQ component that is required.

Edit `/opt/SUNWscmqs/xxx/util/xxx_config` and follow the comments within that file. Where `xxx` represents one of the following IBM MQ components:

```
chi  Channel Initiator
csv  Command Server
lsr  Listener
trm  Trigger Monitor
```

After you have edited *xxx*`_config`, you must register the resource.

```
# cd /opt/SUNWscmqs/xxx/util/
# vi xxx_config
# ./xxx_register
```

The following example shows the `/opt/SUNWscmqs/lsr/util/lsr_config` file that has been edited to configure a listener resource.

```
# cat > /var/tmp/lsr1_config <<-EOF
# +++ Required parameters +++
RS=wmq1-lsr
RG=wmq1-rg
QMGR=qmgr1
PORT=1414
IPADDR=
BACKLOG=100
LH=wmq1-lh
QMGR_RS=wmq1-qmgr
USERID=mqm
RS_ZONE=
PROJECT=default
EOF

# /opt/SUNWscmqs/lsr/util/lsr_register -f /var/tmp/lsr1_config
```

10. **Enable the IBM MQ component resources.**

```
# clresource enable websphere-mq-resource
```

Next Steps See “Verifying the HA for WebSphere MQ Installation and Configuration” on page 28.

Verifying the HA for WebSphere MQ Installation and Configuration

This section contains the procedure you need to verify that you installed and configured your data service correctly.

▼ How to Verify the HA for WebSphere MQ Installation and Configuration

1. **On a cluster member, become an administrator that provides `solaris.cluster.modify` authorization.**

2. **Ensure all the IBM MQ resources are online.**

   ```
   # cluster status
   ```

   Enable any IBM MQ resources that are not online.

   ```
   # clresource enable websphere-mq-resource
   ```

3. **Switch the IBM MQ resource group to another cluster node.**

   ```
   # clresourcegroup switch -n websphere-mq-resource-group
   ```

Upgrading HA for WebSphere MQ

Upgrade the HA for WebSphere MQ data service if the following conditions apply:

- You are upgrading from an earlier version of the HA for WebSphere MQ data service.
- You need to use the new features of this data service.

▼ How to Migrate Existing Resources to a New Version of HA for WebSphere MQ

Perform steps 1, 2, 3 and 6 if you have an existing HA for WebSphere MQ deployment and wish to upgrade to the new version. Complete all steps if you need to use the new features of this data service.

1. **On a cluster member, become an administrator that provides `solaris.cluster.modify` authorization.**

2. **Disable the IBM MQ resources.**

   `# clresource disable` *websphere-mq-resource*

3. **Install the new version of HA for WebSphere MQ to each cluster**

   Refer to "How to Install the HA for WebSphere MQ Package" on page 23 for more information.

4. **Delete the IBM MQ resources, if you want to use new features that have been introduced in the new version of HA for WebSphere MQ.**

   `# clresource delete` *websphere-mq-resource*

5. **Reregister the IBM MQ resources, if you want to use new features that have been introduced in the new version of HA for WebSphere MQ.**

   Refer to "How to Register and Configure HA for WebSphere MQ" on page 24 for more information.

6. **Enable the IBM MQ resources**

   If you have only performed steps 1, 2 and 3 you will need to re-enable the IBM MQ resources.

   `# clresource enable` *websphere-mq-resource*

Understanding the HA for WebSphere MQ Fault Monitor

This section describes the HA for WebSphere MQ fault monitor probing algorithm or functionality, states the conditions, and recovery actions associated with unsuccessful probing.

For conceptual information on fault monitors, see the *Concepts for Oracle Solaris Cluster 4.4*.

Resource Properties

The HA for WebSphere MQ fault monitor uses the same resource properties as resource type `SUNW.gds`. Refer to the `SUNW.gds(5)` man page for a complete list of resource properties used.

Probing Algorithm and Functionality

The HA for WebSphere MQ fault monitor is controlled by the extension properties that control the probing frequency. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Oracle Solaris Cluster installations. Therefore, you should tune the HA for WebSphere MQ fault monitor *only* if you need to modify this preset behavior.

- Setting the interval between fault monitor probes (`Thorough_probe_interval`)
- Setting the time-out for fault monitor probes (`Probe_timeout`)
- Setting the number of times the fault monitor attempts to restart the resource (`Retry_count`)

The HA for WebSphere MQ fault monitor checks the `queue manager` and other components within an infinite loop. During each cycle the fault monitor will check the relevant component and report either a failure or success.

If the fault monitor is successful it returns to its infinite loop and continues the next cycle of probing and sleeping.

If the fault monitor reports a failure a request is made to the cluster to restart the resource. If the fault monitor reports another failure another request is made to the cluster to restart the resource. This behavior will continue whenever the fault monitor reports a failure.

If successive restarts exceed the `Retry_count` within the `Thorough_probe_interval` a request to failover the resource group onto a different node is made.

Operations of the Queue Manager Probe

The IBM MQ queue manager probe checks the queue manager by using a program named `create_tdq` which is included in the HA for WebSphere MQ data service.

The `create_tdq` program connects to the queue manager, creates a temporary dynamic queue, puts a message to the queue and then disconnects from the queue manager.

Operations of the Channel Initiator, Command Server, Listener and Trigger Monitor Probes

The IBM MQ probe for the channel initiator, command server, listener and trigger monitor all operate in a similar manner and will simply restart any component that has failed.

The process monitor facility will request a restart of the resource as soon as any component has failed.

The channel initiator, command server and trigger monitor are all dependent on the queue manger being available. The listener has an optional dependency on the queue manager that is set when the listener resource is configured and registered. Therefore if the queue manager fails the channel initiator, command server, trigger monitor and optional dependent listener will be restarted when the queue manager is available again.

Debugging HA for WebSphere MQ

▼ How to Turn on Debugging for HA for WebSphere MQ

HA for WebSphere MQ can be used by multiple IBM MQ instances. It is possible to turn debug on for all IBM MQ instances or a particular IBM MQ instance.

A `config` file exists under `/opt/SUNWscmqs/xxx/etc`, where `xxx` can be `mgr` (Queue Manager), `chi` (Channel Initiator), `csv` (Command Server), `lsr` (Listener) and `trm` (Trigger Monitor).

These files allow you to turn on debug for all IBM MQ instances or for a specific IBM MQ instance on a particular node within the cluster. If you require debug to be turned on for HA for WebSphere MQ across the whole cluster, repeat this step on all nodes within the cluster.

1. **Edit `/etc/syslog.conf` and change `daemon.notice` to `daemon.debug`.**

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.notice;mail.crit        /var/adm/messages
*.alert;kern.err;daemon.err                     operator
#
```

Change the `daemon.notice` to `daemon.debug` and restart `syslogd`. Note that the output below, from `grep daemon /etc/syslog.conf`, shows that `daemon.debug` has been set.

```
# grep daemon /etc/syslog.conf
*.err;kern.debug;daemon.debug;mail.crit        /var/adm/messages
*.alert;kern.err;daemon.err                    operator
```

Restart the syslog daemon.

```
# svcadm disable system-log
# svcadm enable system-log
```

2. **Edit `/opt/SUNWscmqs/xxx/etc/config`.**

   Perform this step for each component that requires debug output, on each node of Oracle Solaris Cluster as required.

   Edit `/opt/SUNWscmqs/xxx/etc/config` and change `DEBUG=` to `DEBUG=ALL` or `DEBUG=`*resource*.

```
# cat /opt/SUNWscmqs/mgr/etc/config
#
# Copyright 2006 Oracle Corporation.  All rights reserved.
# Use is subject to license terms.
#
##ident   "@(#)config 1.2     06/03/08 SMI"
#
# Usage:
#       DEBUG=<RESOURCE_NAME> or ALL
#
DEBUG=ALL
```

Note - To turn off debug, reverse the steps above.

♦ ♦ ♦ APPENDIX A

Deployment Example: Installing an IBM MQ Queue Manager in the Global Cluster

How to Install an IBM MQ Queue Manager in the Global Cluster

Target Cluster Configuration

This example uses a two-node cluster, pnode1 and pnode2.

Software Configuration

This deployment example uses the following software products:

- Oracle Solaris for SPARC or x86 platforms
- Oracle Solaris Cluster 4.4 software
- HA for WebSphere MQ data service
- IBM MQ v8.0 for Oracle Solaris

This example assumes that you have already installed and configured Oracle Solaris Cluster. It illustrates the installation and configuration of the data service only.

Assumptions

The instructions in this example have been developed with the following assumptions:

- Shell environment: All commands and the environment setup in this example are for the Korn shell environment. If you use a different shell, replace any Korn shell-specific information or instructions with the appropriate information for you preferred shell environment.
- User login: Unless otherwise specified, perform all procedures by assuming a root role that provides solaris.cluster.admin, solaris.cluster.modify, and solaris.cluster.read authorization.

Installing and Configuring IBM MQ v8.0

```
The instructions within the deployment example will install IBM MQ v8.0 for Oracle
 Solaris on Oracle Solaris 11.2 and will configure IBM MQ on an Oracle ZFS highly
 available local file system.

1. Install and configure the cluster as instructed in the Oracle Solaris Cluster 4.4
 Software Installation Guide.
It is recommended that you install the ha-cluster-full IPS package.

2. Create the IBM MQ group and userid.

Perform on each node of the global cluster.

root@pnode1:~# groupadd -g 1002 mqm
root@pnode1:~# useradd -u 1002 -g 1002 -d /var/mqm mqm
root@pnode1:~#
root@pnode1:~# projadd -c "WebSphere MQ default settings" -K "process.max-file-
descriptor=(basic,10000,deny)" -K "project.max-shm-memory=(priv,4GB,deny)" -K
 "project.max-shm-ids=(priv,1024,deny)" -K "project.max-sem-ids=(priv,128,deny)"
 group.mqm
root@pnode1:~#

3. Add a logical host name to /etc/hosts.

Perform on each node of the global cluster.

The following output shows a logical host entry for Queue Manager qmgr1.

root@pnode1:~# grep qmgr1 /etc/hosts
10.134.84.62    qmgr1
root@pnode1:~#

4. Create a ZFS pool and ZFS file systems on the shared storage.

root@pnode1:~# scdidadm -L d1
1        pnode2:/dev/rdsk/c0t600009700001968007955330303031 46d0 /dev/did/rdsk/d1
1        pnode1:/dev/rdsk/c0t600009700001968007955330303031 46d0 /dev/did/rdsk/d1
root@pnode1:~#
root@pnode1:~# zpool create -m /ZFSwmq1 wmq1 /dev/did/dsk/d1s0
root@pnode1:~# zfs create wmq1/log
root@pnode1:~# zfs create wmq1/qmgrs

5. Create appropriate mount points and symlinks for the queue manager.

Note: Some commands are issued on both the global cluster nodes.

root@pnode1:~# mkdir -p /var/mqm/log /var/mqm/qmgrs
root@pnode1:~# ln -s /ZFSwmq1/log /var/mqm/log/qmgr1
root@pnode1:~# ln -s /ZFSwmq1/qmgrs /var/mqm/qmgrs/qmgr1
```

```
root@pnode1:~# chown -R mqm:mqm /var/mqm
root@pnode1:~# chown -R mqm:mqm /ZFSwmq1

root@pnode2:~# mkdir -p /var/mqm/log /var/mqm/qmgrs
root@pnode2:~# ln -s /ZFSwmq1/log /var/mqm/log/qmgr1
root@pnode2:~# ln -s /ZFSwmq1/qmgrs /var/mqm/qmgrs/qmgr1
root@pnode2:~# chown -R mqm:mqm /var/mqm

6. Install the IBM MQ software.

Note: You must install IBM MQ in /opt/mqm.

Perform on each node of the global cluster.

root@pnode1:~# cd software install directory
root@pnode1:~# ./mqlicense.sh
root@pnode1:~# pkgadd -d

7. Set the primary instance to use /opt/mqm.

Perform on each node of the zone cluster.

root@pnode1:~# /opt/mqm/bin/setmqinst -i -p /opt/mqm
90 of 90 tasks have been completed successfully.
'Installation1' (/opt/mqm) set as the primary installation.
root@pnode1:~#

8. Verify the IBM MQ installation

a) Create the Queue Manager qmgr1 on pnode1.

root@pnode1:~# su - mqm
-bash-4.1$
-bash-4.1$ crtmqm qmgr1
WebSphere MQ queue manager created.
Directory '/var/mqm/qmgrs/qmgr1' created.
The queue manager is associated with installation 'Installation1'.
Creating or replacing default objects for queue manager 'qmgr1'.
Default objects statistics : 79 created. 0 replaced. 0 failed.
Completing setup.
Setup completed.
-bash-4.1$

b) Start the Queue Manager qmgr1 on pnode1.

-bash-4.1$ strmqm qmgr1
WebSphere MQ queue manager 'qmgr1' starting.
The queue manager is associated with installation 'Installation1'.
```

```
5 log records accessed on queue manager 'qmgr1' during the log replay phase.
Log replay for queue manager 'qmgr1' complete.
Transaction manager state recovered for queue manager 'qmgr1'.
WebSphere MQ queue manager 'qmgr1' started using V8.0.0.0.
-bash-4.1$

c) Create a test local queue on pnode1.

-bash-4.1$ runmqsc qmgr1
5724-H72 (C) Copyright IBM Corp. 1994, 2014.
Starting MQSC for queue manager qmgr1.


def ql(sc3test) defpsist(yes)
     1 : def ql(sc3test) defpsist(yes)
AMQ8006: WebSphere MQ queue created.
end
     2 : end
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
-bash-4.1$

d) Put a message to the test local queue on pnode1.

-bash-4.1$ /opt/mqm/samp/bin/amqsput SC3TEST qmgr1
Sample AMQSPUT0 start
target queue is SC3TEST
test test test
?^C
-bash-4.1$

e) Stop the Queue Manager qmgr1 on pnode1.

-bash-4.1$ endmqm -i qmgr1
WebSphere MQ queue manager 'qmgr1' ending.
WebSphere MQ queue manager 'qmgr1' ended.
-bash-4.1$
-bash-4.1$ exit
logout
root@pnode1:~#

f) Export the zpool from pnode1.

root@pnode1:~# zpool export wmq1

g) Copy the queue manager definition from pnode1 to pnode2.
```

```
Since the Queue Manager qmgr1 was created on node pnode1, we need to ensure that qmgr1
 is known on node pbikns2. Either copy /var/mqm/mqs.ini between nodes or copy just the
 Queue Manager definition.

root@pnode1:~# su – mqm
-bash-4.1$
-bash-4.1$ cat /var/mqm/mqs.ini


<... snipped ...>

QueueManager:
   Name=qmgr1
   Prefix=/var/mqm
   Directory=qmgr1
   InstallationName=Installation1
-bash-4.1$

h) Add the Queue Manager entry to /var/mqm/mqs.ini on node pnode2.

root@pnode2:~# su – mqm
-bash-4.1$
-bash-4.1$ cat /var/mqm/mqs.ini

<... snipped ...>

QueueManager:
   Name=qmgr1
   Prefix=/var/mqm
   Directory=qmgr1
   InstallationName=Installation1
-bash-4.1$

i) Import the zpool on pnode2.

root@pnode2:~# zpool import wmq1

j) Start the Queue Manager qmgr1 on pnode2.

root@pnode2:~# su – mqm
-bash-4.1$
-bash-4.1$ strmqm qmgr1
WebSphere MQ queue manager 'qmgr1' starting.
The queue manager is associated with installation 'Installation1'.
5 log records accessed on queue manager 'qmgr1' during the log replay phase.
Log replay for queue manager 'qmgr1' complete.
Transaction manager state recovered for queue manager 'qmgr1'.
WebSphere MQ queue manager 'qmgr1' started using V8.0.0.0.
```

```
-bash-4.1$

k) Get the message from the test local queue on pnode2.

-bash-4.1$ /opt/mqm/samp/bin/amqsget SC3TEST qmgr1
Sample AMQSGET0 start
message test
?^C
-bash-4.1$

l) Delete the test local queue on pnode2.

-bash-4.1$ runmqsc qmgr1
5724-H72 (C) Copyright IBM Corp. 1994, 2014.
Starting MQSC for queue manager qmgr1.


delete ql(sc3test)
     1 : delete ql(sc3test)
AMQ8007: WebSphere MQ queue deleted.
end
     2 : end
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
-bash-4.1$

m) Stop the Queue Manager qmgr1 on pnode2.

-bash-4.1$ endmqm -i qmgr1
WebSphere MQ queue manager 'qmgr1' ending.
WebSphere MQ queue manager 'qmgr1' ended.
-bash-4.1$
-bash-4.1$ exit
logout
root@pnode1:~#

n) Export the zpool from pnode2.

root@pnode2:~# zpool export wmq1

9. Create the Oracle Solaris Cluster resources for IBM MQ.

Perform on one node of the global cluster.

a) Register the required resource types.

root@pnode1:~# clrt register SUNW.HAStoragePlus
```

```
root@pnode1:~# clrt register SUNW.gds

b) Create a failover resource group.

root@pnode1:~# clrg create wmq1-rg

c) Create a logical host resource.

root@pnode1:~# clrslh create -g wmq1-rg -h qmgr1 wmq1-lh

d) Create a storage resource.

root@pnode1:~# clrs create -g wmq1-rg -t SUNW.HAStoragePlus -p Zpools=wmq1 wmq1-has

e) Enable the resource groups and its resources.

root@pnode1:~# clrg online -eM wmq1-rg
root@pnode1:~# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-has            pnode2          Online       Online
                    pnode1          Offline      Offline

wmq1-lh             pnode2          Online       Online - LogicalHostname online.
                    pnode1          Offline      Offline

root@pnode1:~#

f) Create a Queue Manager resource.

root@pnode1:~# cd /opt/SUNWscmqs/mgr/util
root@pnode1:/opt/SUNWscmqs/mgr/util# vi mgr_config

g) Set the following entries:

RS=wmq1-rs
RG=wmq1-rg
QMGR=qmgr1
LH=wmq1-lh
HAS_RS=wmq1-has

h) Register the Queue Manager resource.

root@pnode1:/opt/SUNWscmqs/mgr/util# ./mgr_register
```

```
i) Enable the Queue Manager resource.

root@pnode1:/opt/SUNWscmqs/mgr/util# clrs enable wmq1-rs
root@pnode1:/opt/SUNWscmqs/mgr/util# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-rs             pnode2         Online       Online
                    pnode1         Offline      Offline

wmq1-has            pnode2         Online       Online
                    pnode1         Offline      Offline

wmq1-lh             pnode2         Online       Online - LogicalHostname online.
                    pnode1         Offline      Offline

root@pnode1:/opt/SUNWscmqs/mgr/util#

j) Create a Listener resource.

root@pnode1:~# cd /opt/SUNWscmqs/lsr/util
root@pnode1:/opt/SUNWscmqs/lsr/util# vi lsr_config

k) Set the following entries:

RS=wmq1-lsr
RG=wmq1-rg
QMGR=qmgr1
PORT=1414
IPADDR=10.134.84.62
BACKLOG=100
LH=wmq1-lh
QMGR_RS=wmq1-rs
USERID=mqm

l) Register the Listener resource.

root@pnode1:/opt/SUNWscmqs/lsr/util# ./lsr_register

m) Enable the Listener resource.

root@pnode1:/opt/SUNWscmqs/lsr/util# clrs enable wmq1-lsr
root@pnode1:/opt/SUNWscmqs/lsr/util# clrs status -g wmq1-rg

=== Cluster Resources ===
```

```
Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-rs             pnode2         Online       Online
                    pnode1         Offline      Offline

wmq1-has            pnode2         Online       Online
                    pnode1         Offline      Offline

wmq1-lh             pnode2         Online       Online - LogicalHostname online.
                    pnode1         Offline      Offline

root@pnode1:/opt/SUNWscmqs/lsr/util#
```

10. Switch the resource group between nodes.

```
root@pnode1:~# clrg switch -n pnode1 wmq1-rg
root@pnode1:~# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-lsr            pnode2         Offline      Offline
                    pnode1         Online       Online

wmq1-rs             pnode2         Offline      Offline
                    pnode1         Online       Online

wmq1-has            pnode2         Offline      Offline
                    pnode1         Online       Online

wmq1-lh             pnode2         Offline      Offline - LogicalHostname offline.
                    pnode1         Online       Online - LogicalHostname online.

root@pnode1:~#
root@pnode1:~# clrg switch -n pnode2 wmq1-rg
root@pnode1:~# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-lsr            pnode2         Online       Online
                    pnode1         Offline      Offline

wmq1-rs             pnode2         Online       Online
                    pnode1         Offline      Offline
```

```
wmq1-has            pnode2         Online       Online
                    pnode1         Offline      Offline

wmq1-lh             pnode2         Online       Online - LogicalHostname online.
                    pnode1         Offline      Offline - LogicalHostname offline.

root@pnode1:~#

(Optional) Repeat for each IBM MQ component that is required. Edit /opt/SUNWscmqs/xxx/
util/xxx_config and follow the comments within that file. Where
xxx represents one of the following IBM MQ components:
                 chi  Channel Initiator
                 csv  Command Server
                 trm  Trigger Monitor

After you edit the xxx_config file, you must register the resource.
# cd /opt/SUNWscmqs/xxx/util/
# vi xxx_config
# ./xxx_register
```

♦ ♦ ♦ APPENDIX B

Deployment Example: Installing an IBM MQ Queue Manager in the Zone Cluster

How to Install an IBM MQ Queue Manager in the Zone Cluster

Target Cluster Configuration

This example uses a two-node cluster, pnode1 and pnode2, that has a zone cluster named zc1. Zone cluster nodes are named vznode1 and vznode2.

```
root@pnode1:~# clzc status

=== Zone Clusters ===

--- Zone Cluster Status ---

Name   Brand     Node Name   Zone Host Name   Status   Zone Status
----   -----     ---------   --------------   ------   -----------
zc1    solaris   pnode1      vznode1          Online   Running
                 pnode2      vznode2          Online   Running

root@pnode1:~#
```

Software Configuration

This deployment example can use the following software products:

- Oracle Solaris for SPARC or x86 platforms
- Oracle Solaris Cluster 4.4 software
- HA for WebSphere MQ data service
- IBM MQ v8.0 for Oracle Solaris

This example assumes that you have already installed and configured Oracle Solaris Cluster. It shows the installation and configuration of the data service only.

Assumptions

The instructions in this example have been developed with the following assumptions:

- Shell environment: All commands and the environment setup in this example are for the Korn shell environment. If you use a different shell, replace any Korn shell-specific information or instructions with the appropriate information for you preferred shell environment.
- User login: Unless otherwise specified, perform all procedures by assuming a root role that provides solaris.cluster.admin, solaris.cluster.modify, and solaris.cluster.read authorization.

Installing and Configuring IBM WebSphere MQ v8.0

The instructions within the deployment example will install IBM MQ v8.0 for Solaris on Oracle Solaris 11.2 and will configure IBM MQ on an Oracle ZFS highly available local file system.

1. Install and configure the cluster as instructed in the Oracle Solaris Cluster 4.4 Software Installation Guide.
It is recommended that you install the ha-cluster-full IPS package.

2. Create the zone cluster.

Later on Queue Manager qmgr1 will be deployed within a zone cluster. Once created, IBM MQ will be installed within the zone cluster nodes.
Note that the logical host qmgr1 and Oracle ZFS Storage Pool wmq1 have been included within the zone cluster configuration.

Perform on one node of the global cluster.

```
root@pnode1:~# cat /var/tmp/zc1.txt
create
set zonepath=/zones/zc1
set autoboot=true
add node
set physical-host=pndoe1
set hostname=vznode1
add net
set address=10.134.84.88
set physical=sc_ipmp0
end
end
add node
set physical-host=pnode2
set hostname=vznode2
add net
set address=10.134.84.90
```

```
set physical=sc_ipmp0
end
end
add net
set address=qmgr1
end
add dataset
set name=wmq1
end
commit
exit
root@pnode1:~#
root@pnode1:~# clzc configure -f /var/tmp/zc1.txt zc1
root@pnode1:~#clzc install zc1
Waiting for zone install commands to complete on all the nodes of the zone cluster
 "zc1"...
root@pnode1:~#
root@pnode1:~# clzc boot zc1
Waiting for zone boot commands to complete on all the nodes of the zone cluster "zc1"...
root@pnode1:~#
root@pnode1:~# clzc status

=== Zone Clusters ===

--- Zone Cluster Status ---

Name   Brand     Node Name   Zone Host Name   Status   Zone Status
----   -----     ---------   --------------   ------   -----------
zc1    solaris   pnode1      vznode1        Online   Running
                 pnode2      vznode2        Online   Running

root@pnode1:~#
```

3. Create the IBM MQ group and userid.

Perform on each node of the zone cluster.

```
root@pnode1:~# zlogin zc1
[Connected to zone 'zc1' pts/2]
?Oracle Corporation SunOS 5.11 11.2 July 2015
root@vznode1:~#
root@vznode1:~# groupadd -g 1002 mqm
root@vznode1:~# useradd -u 1002 -g 1002 -d /var/mqm mqm
root@vznode1:~#
root@vznode1:~# projadd -c "WebSphere MQ default settings" -K "process.max-file-
descriptor=(basic,10000,deny)" -K "project.max-shm-memory=(priv,4GB,deny)" -K
 "project.max-shm-ids=(priv,1024,deny)" -K "project.max-sem-ids=(priv,128,deny)"
 group.mqm
```

```
root@vznode1:~#

4. Add a logical host name to /etc/hosts.

Perform on each node of the zone cluster.

The following output shows a logical host entry for Queue Manager qmgr1.

root@vznode1:~# grep qmgr1 /etc/hosts
10.134.84.62    qmgr1
root@vznode1:~#

5. Create a ZFS pool and ZFS file systems on a shared storage on a global cluster node.

Perform this step on one node of the global cluster.

root@pnode1:~# scdidadm -L d1
1        pnode2:/dev/rdsk/c0t600009700001968007955330303031 46d0 /dev/did/rdsk/d1
1        pnode1:/dev/rdsk/c0t600009700001968007955330303031 46d0 /dev/did/rdsk/d1
root@pnode1:~#
root@pnode1:~# zpool create -m /ZFSwmq1 wmq1 /dev/did/dsk/d1s0
root@pnode1:~# zfs create wmq1/log
root@pnode1:~# zfs create wmq1/qmgrs

6. Create Oracle Solaris Cluster logical host and storage resources.

Perform on one node of the zone cluster.

a) Register the required resource types.

root@vznode1:~# export PATH=$PATH:/usr/cluster/bin
root@vznode1:~# clrt register SUNW.HAStoragePlus
root@vznode1:~# clrt register SUNW.gds

b) Create a failover resource group.

root@vznode1:~# clrg create wmq1-rg

c) Create a logical host resource.

root@vznode1:~# clrslh create -g wmq1-rg -h qmgr1 wmq1-lh

d) Create a storage resource.

root@vznode1:~# clrs create -g wmq1-rg -t SUNW.HAStoragePlus -p Zpools=wmq1 wmq1-has

e) Enable the resource groups and its resources.
```

```
root@vznode1:~# clrg online -eM wmq1-rg
root@vznode1:~# clrg online -eM -n vznode1 wmq1-rg
root@vznode1:~# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-has            vznode2        Offline      Offline
                    vznode1        Online       Online

wmq1-lh             vznode2        Offline      Offline - LogicalHostname offline.
                    vznode1        Online       Online - LogicalHostname online.

root@vznode1:~#
```

7. Create appropriate mount points and symlinks for the queue manager.

Note: Some commands are issued on both the nodes of the zone cluster.

```
root@vznode1:~# mkdir -p /var/mqm/log /var/mqm/qmgrs
root@vznode1:~# ln -s /ZFSwmq1/log /var/mqm/log/qmgr1
root@vznode1:~# ln -s /ZFSwmq1/qmgrs /var/mqm/qmgrs/qmgr1
root@vznode1:~# chown -R mqm:mqm /var/mqm
root@vznode1:~# chown -R mqm:mqm /ZFSwmq1

root@vznode2:~# mkdir -p /var/mqm/log /var/mqm/qmgrs
root@vznode2:~# ln -s /ZFSwmq1/log /var/mqm/log/qmgr1
root@vznode2:~# ln -s /ZFSwmq1/qmgrs /var/mqm/qmgrs/qmgr1
root@vznode2:~# chown -R mqm:mqm /var/mqm
```

8. Install IBM WebSphere MQ software.

Note: You must install IBM MQ in /opt/mqm.

Perform on each node of the zone cluster.

```
root@vznode1:~# cd software intall directory
root@vznode1:~# ./mqlicense.sh
root@vznode1:~# pkgadd -d .
```

9. Set the primary instance to use /opt/mqm.

Perform on each node of the zone cluster.

```
root@vznode1:~# /opt/mqm/bin/setmqinst -i -p /opt/mqm
```

```
90 of 90 tasks have been completed successfully.
'Installation1' (/opt/mqm) set as the primary installation
root@vznode1:~#

10. Verify the IBM MQ installation

a) Create the Queue Manager qmgr1 on vznode1.

root@vznode1:~# su - mqm
-bash-4.1$
-bash-4.1$ crtmqm qmgr1
WebSphere MQ queue manager created.
Directory '/var/mqm/qmgrs/qmgr1' created.
The queue manager is associated with installation 'Installation1'.
Creating or replacing default objects for queue manager 'qmgr1'.
Default objects statistics : 79 created. 0 replaced. 0 failed.
Completing setup.
Setup completed.
-bash-4.1$

b) Start the Queue Manager qmgr1 on vznode1.

-bash-4.1$ strmqm qmgr1
WebSphere MQ queue manager 'qmgr1' starting.
The queue manager is associated with installation 'Installation1'.
5 log records accessed on queue manager 'qmgr1' during the log replay phase.
Log replay for queue manager 'qmgr1' complete.
Transaction manager state recovered for queue manager 'qmgr1'.
WebSphere MQ queue manager 'qmgr1' started using V8.0.0.0.
-bash-4.1$

c) Create a test local queue on vznode1.

-bash-4.1$ runmqsc qmgr1
5724-H72 (C) Copyright IBM Corp. 1994, 2014.
Starting MQSC for queue manager qmgr1.


def ql(sc3test) defpsist(yes)
     1 : def ql(sc3test) defpsist(yes)
AMQ8006: WebSphere MQ queue created.
end
     2 : end
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
-bash-4.1$
```

```
d) Put a message to the test local queue on vznode1.

-bash-4.1$ /opt/mqm/samp/bin/amqsput SC3TEST qmgr1
Sample AMQSPUT0 start
target queue is SC3TEST
test test test
?^C
-bash-4.1$

e) Stop the Queue Manager qmgr1 on vznode1.

-bash-4.1$ endmqm -i qmgr1
WebSphere MQ queue manager 'qmgr1' ending.
WebSphere MQ queue manager 'qmgr1' ended.
-bash-4.1$
-bash-4.1$ exit
logout
root@vznode1:~#

f) Switch the Oracle Solaris Cluster resource group to vznode2.

root@vznode1:~# clrg switch -n vznode2 wmq1-rg
root@vznode1:~# clrs status -g wmq1-rg +

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-has            vznode2        Online       Online
                    vznode1        Offline      Offline

wmq1-lh             vznode1        Online       Online - LogicalHostname online.
                    vznode1        Offline      Offline - LogicalHostname offline.

root@vznode1:~#

g) Copy the queue manager definition from vznode1 to vznode2.

Since the Queue Manager qmgr1 was created on node vznode1, ensure that qmgr1 is known on
 node vznode2. Either copy /var/mqm/mqs.ini between nodes or just copy the Queue Manager
 definition.

root@vznode1:~# su – mqm
-bash-4.1$
-bash-4.1$ cat /var/mqm/mqs.ini


<... snipped ...>
```

```
QueueManager:
   Name=qmgr1
   Prefix=/var/mqm
   Directory=qmgr1
   InstallationName=Installation1
-bash-4.1$

h) Add the Queue Manager entry to /var/mqm/mqs.ini on node vznode2.

root@vznode2:~# su – mqm
-bash-4.1$
-bash-4.1$ cat /var/mqm/mqs.ini


<... snipped ...>

QueueManager:
   Name=qmgr1
   Prefix=/var/mqm
   Directory=qmgr1
   InstallationName=Installation1
-bash-4.1$


i) Start the Queue Manager qmgr1 on vznode2.

-bash-4.1$ strmqm qmgr1
WebSphere MQ queue manager 'qmgr1' starting.
The queue manager is associated with installation 'Installation1'.
5 log records accessed on queue manager 'qmgr1' during the log replay phase.
Log replay for queue manager 'qmgr1' complete.
Transaction manager state recovered for queue manager 'qmgr1'.
WebSphere MQ queue manager 'qmgr1' started using V8.0.0.0.
-bash-4.1$

j) Get the message from the test local queue on vznode2.

-bash-4.1$ /opt/mqm/samp/bin/amqsget SC3TEST qmgr1
Sample AMQSGET0 start
message test
?^C
-bash-4.1$

k) Delete the test local queue on vznode2.

-bash-4.1$ runmqsc qmgr1
5724-H72 (C) Copyright IBM Corp. 1994, 2014.
```

```
Starting MQSC for queue manager qmgr1.


delete ql(sc3test)
     1 : delete ql(sc3test)
AMQ8007: WebSphere MQ queue deleted.
end
     2 : end
One MQSC command read.
No commands have a syntax error.
All valid MQSC commands were processed.
-bash-4.1$

l) Stop the Queue Manager qmgr1 on vznode2.

-bash-4.1$ endmqm -i qmgr1
WebSphere MQ queue manager 'qmgr1' ending.
WebSphere MQ queue manager 'qmgr1' ended.
-bash-4.1$
-bash-4.1$ exit
logout
root@vznode2:~#


11. Create Oracle Solaris Cluster resources for IBM MQ.

Perform on one node of the zone cluster.

a) Create a Queue Manager resource.

root@vznode2:~# cd /opt/SUNWscmqs/mgr/util
root@vznode2:/opt/SUNWscmqs/mgr/util# vi mgr_config

b) Set the following entries:

RS=wmq1-rs
RG=wmq1-rg
QMGR=qmgr1
LH=wmq1-lh
HAS_RS=wmq1-has

c) Register the Queue Manager resource.

root@vznode2:/opt/SUNWscmqs/mgr/util# ./mgr_register

d) Enable the Queue Manager resource.

root@vznode2:/opt/SUNWscmqs/mgr/util# clrs enable wmq1-rs
```

```
root@vznode2:/opt/SUNWscmqs/mgr/util# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-rs             vznode2        Online       Online
                    vznode1        Offline      Offline

wmq1-has            vznode2        Online       Online
                    vznode1        Offline      Offline

wmq1-lh             vznode2        Online       Online - LogicalHostname online.
                    vznode1        Offline      Offline

root@vznode2:/opt/SUNWscmqs/mgr/util#

e) Create a Listener resource.

root@vznode2:~# cd /opt/SUNWscmqs/lsr/util
root@vznode2:/opt/SUNWscmqs/lsr/util# vi lsr_config

f) Set the following entries:

RS=wmq1-lsr
RG=wmq1-rg
QMGR=qmgr1
PORT=1414
IPADDR=10.134.84.62
BACKLOG=100
LH=wmq1-lh
QMGR_RS=wmq1-rs
USERID=mqm

g) Register the Listener resource.

root@vznode2:/opt/SUNWscmqs/lsr/util# ./lsr_register

h) Enable the Listener resource.

root@vznode2:/opt/SUNWscmqs/lsr/util# clrs enable wmq1-lsr
root@vznode2:/opt/SUNWscmqs/lsr/util# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-rs             vznode2        Online       Online
```

```
                    vznode1        Offline      Offline

wmq1-has            vznode2        Online       Online
                    vznode1        Offline      Offline

wmq1-lh             vznode2        Online       Online - LogicalHostname online.
                    vznode1        Offline      Offline

root@vznode2:/opt/SUNWscmqs/lsr/util#


12. Switch the resource group between nodes.

root@vznode2:~# clrg switch -n vznode1 wmq1-rg
root@vznode2:~# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-lsr            vznode2          Offline      Offline
                    vznode1          Online       Online

wmq1-rs             vznode2          Offline      Offline
                    vznode1          Online       Online

wmq1-has            vznode2          Offline      Offline
                    vznode1          Online       Online

wmq1-lh             vznode2          Offline      Offline - LogicalHostname offline.
                    vznode1          Online       Online - LogicalHostname online.

root@vznode2:~#
root@vznode2:~# clrg switch -n vznode2 wmq1-rg
root@vznode2:~# clrs status -g wmq1-rg

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
wmq1-lsr            vznode2        Online       Online
                    vznode1        Offline      Offline

wmq1-rs             vznode2        Online       Online
                    vznode1        Offline      Offline

wmq1-has            vznode2        Online       Online
                    vznode1        Offline      Offline
```

```
wmq1-lh             vznode2        Online       Online - LogicalHostname online.
                    vznode1        Offline      Offline - LogicalHostname offline.

root@vznode2:~#

(Optional) Repeat for each IBM MQ component that is required. Edit /opt/SUNWscmqs/xxx/
util/xxx_config and follow the comments within that file. Where
xxx represents one of the following IBM MQ components:
                  chi  Channel Initiator
                  csv  Command Server
                  trm  Trigger Monitor

After you edit the xxx_config file, you must register the resource.
# cd /opt/SUNWscmqs/xxx/util/
# vi xxx_config
# ./xxx_register
```

Index