



Monitoring and Controlling Failover

March 2015

This chapter contains these topics:

- [Failover Alerts, page 3-1](#)
- [Monitoring Failover from Cisco DMM, page 3-5](#)
- [Monitoring Failover from AAI, page 3-7](#)
- [Forcing an Appliance to Fail Over, page 3-10](#)

Failover Alerts


Two additional alerts have been added to the Cisco DMM Administration > Alerts > Notification Rules page to support failover:

- **Cluster node is deactivated**—When configured, this alert is triggered whenever an appliance in a failover configuration goes offline.
- **Cluster Node is activated**—When configured, this alert is triggered whenever an appliance in a failover configuration comes online.

When an appliance in a failover configuration fails, you will receive a cluster node down notification.

If you reboot an appliance, you will receive a cluster down notification followed by a cluster node activated notification for that appliance as the appliances reboots into the standby state.

Figure 3-1 Failover Alerts



See the *Events and Notifications Chapter* in the *User Guide for Cisco Digital Media Manager 5.3.x* for information about enabling events, configuring your SNMP server, and populating your MIB browser.

http://www.cisco.com/c/en/us/td/docs/video/digital_media_systems/5_x/5_3/dmm/user/guide/dmsug53x/eventnotify.html

For more information about each type of alert, see the following topics:

- [SNMP Alerts, page 3-2](#)
- [Syslog Alerts, page 3-3](#)
- [E-Mail Alerts, page 3-3](#)

SNMP Alerts

The following traps pertain to appliance Up/Down events:

- .1.3.6.1.4.1.9.9.655.0.6—cluster node down
- .1.3.6.1.4.1.9.9.655.0.5—cluster node up

Syslog Alerts

The following are sample UP/DOWN syslog alerts:

```
05-17-201110:56:42Local7.Debug10.0.0.1May 16 22:54:51 dmm.example.com
%DMS-1-ClusterNodeDownEvent: Cluster node dmm1.example.com is DOWN[DmmCluster] [
Original severity = severityCATASTROPHIC ]
```

```
05-17-201110:58:11Local7.Debug10.194.51.45May 16 22:56:21 dmm1.example.com
%DMS-1-ClusterNodeUpEvent: Cluster node dmm1.example.com is UP[DmmCluster] [ Original
severity = severityINFO ]
```

E-Mail Alerts

Figure 3-2 shows a typical event e-mail notification.

Figure 3-2 A Failover Node Outage Notification

From: root@dmm.example.com [mailto:root@dmm.example.com]
Sent: Monday, May 16, 2011 10:20 AM
To: System Admin (sysadmin@example.com)
Subject: DMS Alert 'ClusterNodeDownEvent'

This is an alarm from *Digital Media Systems* with the following details:

- Alarm Type: ClusterNodeDownEvent
- Alarm Source: DmmCluster
- Cluster Virtual FQDN: dmm.example.com
- Cluster Node FQDN: dmm1.example.com
- Severity: severityCATASTROPHIC
- When originated: Mon May 16 10:20:07 PDT 2011
- Comments: Cluster node dmm1.example.com is UNKNOWN

Digital Media Manager Administration Module

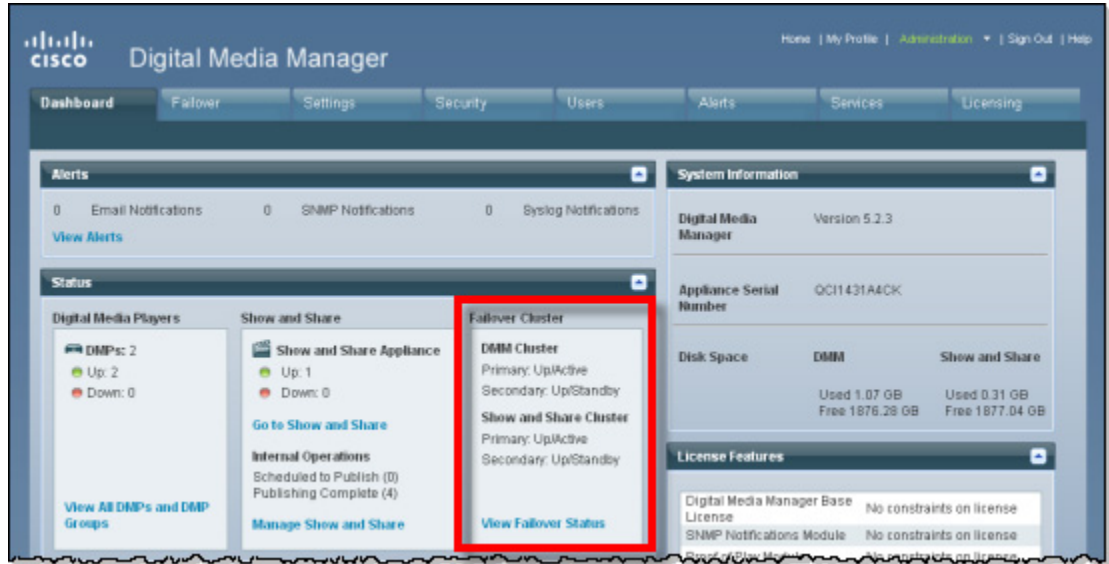
Table 3-1 shows the information that is sent by e-mail:

Table 3-1 Event E-Mail Notification Fields

Field	Description
Alarm Type	<ul style="list-style-type: none"> ClusterNodeDownEvent—The appliance failed or has been taken offline. ClusterNodeUpEvent—The appliance has come online and has entered the active or standby state.
Alarm Source	<ul style="list-style-type: none"> DmmCluster—The alarm came from a Cisco DMM appliance. VpCluster—The alarm came from a Cisco Show and Share appliance.
Cluster Virtual FQDN	The virtual FQDN of the appliance cluster.
Cluster Node FQDN:	The dedicated FQDN of the appliance.
Severity	<ul style="list-style-type: none"> severityCATASTROPHIC—The appliance has experienced a failover event. severityINFO—The message is an informational event (such as an UP message).
Comments:	<p>The comment takes the form of:</p> <p>Cluster node <i>dedicated_fqdn</i> is <i>status</i></p> <p>The <i>status</i> is one of the following values:</p> <ul style="list-style-type: none"> UNKNOWN—The appliance is transitioning between states. UP—The appliance is up and in the active state. DOWN—The appliance has failed. <p>STANDBY—The appliance is up and in the standby state.</p>

Monitoring Failover from Cisco DMM

The Cisco DMM home page displays a summary status of the failover cluster.



Click **View Failover Status** to go to the **Administration > Failover > Status** page.

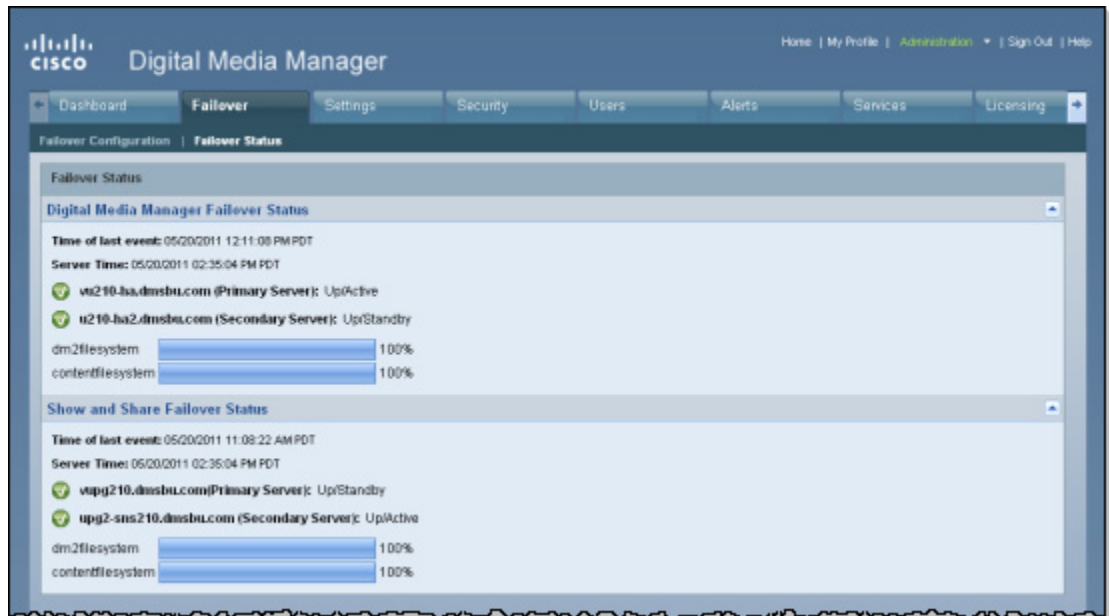


Table 3-2 provides the Failover Status information.

Table 3-2 Failover Status

Field	Description
Time of last event	The time (determined by the appliance time) of the last failover event.
Server Time	The time on the appliance.
Server status	For each server (Primary and Secondary), one of the following states: <ul style="list-style-type: none"> • Up/Active—The appliance is operating normally and is in the active state. • Up/Standby—The appliance is operating normally and is in the standby state. • Down—The appliance experienced a failover event and is currently in a failed state. Depending upon the failure, you may be able to access the appliance AAI interface. • Unknown—The appliance is transitioning between the UP and DOWN states.
Replication Status	The percentage complete the replication of information between the primary and secondary appliance. During initial activation, this value will be below 100 percent and the failover cluster is configured. During normal operation, this value should remain at 100 percent.

What to Look For on This Page

The following conditions indicate abnormal operation and should be investigated:

- An appliance in the Down state. Use the [Cluster Resource Status, page 3-8](#) page to determine which resources have failed.
- An appliance in the Unknown state. This state indicates that the appliance is transitioning between UP and DOWN.
- One node down and the message “No sync in progress”. There can be several causes for this. The failover cluster may be in Split Brain mode (see [Split Brain Recovery, page 4-3](#), for information on how to confirm and recover from split brain)

The active mode may have had a disk fail but not failed over. In this case, you can force a failover (see [Forcing an Appliance to Fail Over, page 3-10](#)) and then proceed with the recovery procedure (see [Recovering from a Failover, page 4-1](#)).

Monitoring Failover from AAI

You can monitor the following by using AAI:

- [Replication Status, page 3-7](#)
- [Cluster Resource Status, page 3-8](#)

Replication Status

The AAI replication status screen provides you with the same information as the Cisco DMM **Administration > Failover > Failover Status** page. You can use this screen to track the progress of data replication.

```

                                REPLICATION STATUS
REPLICATION STATUS
  0:dm2filesystem      Connected Primary/Secondary UpToDate/UpToDate C r----
 /dm2      ext3 17G  1.2G 15G  8%
  1:contentfilesystem Connected Primary/Secondary UpToDate/UpToDate C r----
 /content ext3 1.9T  1.1G 1.8T 1%
  
```

< OK >

Procedure

To access the Replication Status screen, do the following:

-
- Step 1** Log in to AAI.
- Step 2** Choose FAIL_OVER > STATUS > REPLICATION.
-

Cluster Resource Status

The cluster resource status screen displays the status of the monitored components and services. When determining the cause of a failover, use this screen to check the status of the monitored services.

- Services with a status of “Started” are operating normally.
- Services with a status of “Stopped” have failed.

```

CLUSTER RESOURCE STATUS
=====
Last updated: Fri May 20 14:43:02 2011
Stack: Heartbeat
Current DC: vu210-ha.dmsbu.com (717b1ad4-f632-49dc-8455-a6384ae0b9ee) -
partition with quorum
Version: 1.0.9-89bd754939df5150de7cd76835f98fe90851b677
2 Nodes configured, unknown expected votes
4 Resources configured.
=====

Online: [ vu210-ha.dmsbu.com u210-ha2.dmsbu.com ]

Resource Group: DMS-DMM-group
  external-addr-ip (ocf::heartbeat:IPaddr2): Started vu210-ha.dmsbu.com
  unmountWAASAtStartup (lsb:waasUnmounterAtStartup): Started
vu210-ha.dmsbu.com
  dm2 (ocf::heartbeat:Filesystem): Started vu210-ha.dmsbu.com
  content (ocf::heartbeat:Filesystem): Started vu210-ha.dmsbu.com
  unmountWAASAtShutdown (lsb:waasUnmounterAtShutdown): Started
42%
< OK >

```

When a service is shown as “unmanaged” or “failed”, the nodes should be restarted according to the following:

- UNMANAGED FAILED—Both nodes should be restarted, starting first with the node showing unmanaged, then the other node.
- FAILED—The node on which the resource is shown as Failed should be restarted.


```

CLUSTER RESOURCE STATUS
=====
Last updated: Fri May 13 21:27:15 2011
Stack: Heartbeat
Current DC: crepe.cisco.com (e7c13165-ccd8-45ec-bf97-52d626d91d1b) - partition with quorum
Version: 1.0.9-89bd754939df5150de7cd76835f98fe90851b677
2 Nodes configured, unknown expected votes
4 Resources configured.
=====

Online: [ roti.cisco.com crepe.cisco.com ]

Resource Group: DMS-DMM-group
external-addr-ip (ocf::heartbeat:IPaddr2): Stopped
unmountNAASAtStartup (lsb:waasUnmounterAtStartup): Stopped
dm2 (ocf::heartbeat:Filesystem): Stopped
content (ocf::heartbeat:Filesystem): Stopped
unmountNAASAtShutdown (lsb:waasUnmounterAtShutdown): Started crepe.cisco.com (unmanaged) FAILED
activemq (ocf::dms:activemq): Stopped
DmsNodeDeactivationNotifier (ocf::dms:DmsNodeDeactivationNotifier): Stopped
DmsFlashPolicyDaemon (ocf::dms:DmsFlashPolicyDaemon): Stopped
pgsql (ocf::heartbeat:pgsql): Stopped
ems (ocf::dms:ems): Stopped
apache (ocf::heartbeat:apache): Stopped
tomcat (ocf::dms:tomcat): Stopped

```

The fail count for each service appears in the Migration summary section at the bottom of the screen:

```

apache (ocf::heartbeat:apache): Started vu210-ha.dmsbu.com
tomcat (ocf::dms:tomcat): Started vu210-ha.dmsbu.com
scheduleBackup (lsb:scheduleBackup): Started vu210-ha.dmsbu.com
dmpdiscoverer (lsb:dmpdiscoverer): Started vu210-ha.dmsbu.com
rsyslog (lsb:rsyslog): Started vu210-ha.dmsbu.com
DmsNodeActivationNotifier (ocf::dms:DmsNodeActivationNotifier): Started
vu210-ha.dmsbu.com
Master/Slave Set: ms_drbd_contentfilesystem
Masters: [ vu210-ha.dmsbu.com ]
Slaves: [ u210-ha2.dmsbu.com ]
Master/Slave Set: ms_drbd_dm2filesystem
Masters: [ vu210-ha.dmsbu.com ]
Slaves: [ u210-ha2.dmsbu.com ]
Clone Set: connected
Started: [ u210-ha2.dmsbu.com vu210-ha.dmsbu.com ]

Migration summary:
* Node u210-ha2.dmsbu.com: pingd=1
* Node vu210-ha.dmsbu.com: pingd=1

```

Procedure

To access the Cluster Resource Status screen, do the following:

-
- Step 1** Log in to AAI.
 - Step 2** Choose FAIL_OVER > STATUS > CLUSTER_RESOURCE.
 - Step 3** Use the up and down arrow keys to scroll through the displayed information.
-

Forcing an Appliance to Fail Over

To force an appliance to fail over, do the following:

-
- Step 1** Log in to the active appliance AAI interface. Use the virtual FQDN or IP address to ensure that you are accessing the active appliance.
- Step 2** Choose APPLIANCE_CONTROL > RESTART_OPTIONS > RESTART_WEB_SERVICES.
- Restarting the web services on the active appliance triggers a failover to the secondary appliance. The appliance reboots to the standby state and uses the dedicated FQDN and IP address.
-