Hitless Failover and Hitless Upgrade User Guide

This guide describes hitless failover and hitless upgrade, including:

T-sync is a term used to describe the hitless failover and hitless upgrade features available on the BlackDiamond® Management Switch Module 3 (MSM-3). In simple terms, *hitless failover* transfers switch management control from the master MSM-3 to the slave MSM-3 without causing traffic to be dropped. *Hitless upgrade* allows an ExtremeWare® software upgrade on a BlackDiamond 6800 series chassis without taking it out of service or losing traffic.

To configure hitless failover or hitless upgrade, you must install MSM-3 modules in your BlackDiamond chassis; MSM64i modules do not support hitless failover or hitless upgrade.

If you enable T-sync and normally use scripts to configure your switch, Extreme Networks recommends using the `download configuration incremental` command instead.

NOTE

To use the T-sync features available on the MSM-3 modules, you must install and run ExtremeWare 7.1.1 or later and BootROM 8.1 or later.

Causes and Behaviors of MSM Failover

This section describes the events that cause an MSM failover and the behavior of the system after failover occurs.

The following events cause an MSM failover:

- Operator command

- Software exception
- Watchdog timeout
- Keepalive failure
- Diagnostic failure
- Hot-removal of the master MSM
- Hard-reset of the master MSM

NOTE

Operator command and software exception support hitless failover.

Operator Command and Software Exception. Of the listed events, only operator command and software exception result in a hitless failover. The remaining sections of this guide describe T-sync, including:

- Supported features
- How to configure the T-sync features
- The behavior surrounding hitless failover and hitless upgrade

Watchdog Timeout and Keepalive Failure. Both the watchdog timeout and the keepalive failure are long duration events, thus they are not hitless. If one of these events occur:

- All saved operational state information is discarded
- The failed master is hard reset
- The slave uses its own flash configuration file

Diagnostic Failure, Hot-removal, or Hard-reset of the Master MSM. If the master MSM-3 experiences a diagnostic failure or you hot-remove it, a “partial” hitless failover function is performed and some traffic flows will not experience traffic hits. The switch cannot perform a completely hitless failover because it lost hardware that it uses during normal operation.

To understand how traffic is affected when MSM-3 hardware is lost, a brief explanation of the switch fabric is given. Each MSM-3 has switching logic that provides bandwidth to each I/O module. When two MSM-3s are present, both provide bandwidth so that twice the amount of bandwidth is available. For each traffic flow that requires inter-module data movement, the I/O module chooses an MSM-3 to switch the data for that flow. When an MSM-3 is lost, the remaining MSM-3 eventually instructs the I/O module that all inter-module traffic is to use the switching logic of the remaining MSM-3. In the time between the loss of an MSM-3 and the reprogramming of the I/O module, traffic destined for the lost MSM-3 switching logic is dropped.

The I/O module also switches some traffic flows directly between its own ports without MSM-3 involvement.

If you hot-remove the master MSM-3, only half of the switch fabric remains operational. The slave becomes the master and reprograms each I/O module to send all traffic through it’s own switch fabric logic. In the time between the failure and the reprogramming of the I/O module, traffic destined for the removed MSM-3’s switching logic is lost. After the new master recovers, it reprograms the I/O module so that all traffic uses the available MSM-3 switching logic.

If you hard-reset the master MSM-3 (using the recessed reset button on the MSM-3), all of the master’s switch programming is lost. As a result, traffic that the I/O module forwards to the master is also lost.

After a failover occurs, the new master reprograms the “reset” MSM-3’s switch fabric and the switching logic of both MSM-3s is available again. In this case, the “Cause of last MSM failover” displayed by the `show msm-failover` command indicates “removal,” and a “partial” hitless failover has occurred.

A “partial” hitless failover preserves:

- Data flows in the hardware and software, layer 2 protocol states, configurations, etc.
- All of the software states and the hardware states that are not interrupted by the diagnostic failure, hot-removal, or hard-reset.

After a failover caused by hot-removal or diagnostic failure, the I/O modules are reprogrammed to use only the switching logic of the remaining MSM-3. After a failover caused by a hard-reset of the master MSM-3, the reset MSM-3’s switch fabric is reprogrammed and placed into full operation. Thus, a data hit of several seconds occurs for flows that were directed to the failed MSM-3. For flows that were directed to the currently active MSM-3, or for inter-module flows, there is no hit.

Summary of Supported Features

This section describes the features supported by T-sync. If the information in the release notes differ from the information in this guide, follow the release notes.

- Preserves unsaved configurations across a failover
- Load sharing
- Learned MAC address
- ARP
- STP
- EAPSv1
- IP FDB entries
- Access lists
- ESRP
- SNMP trap failover
- Configuration via the web, CLI, and SNMP

NOTE

T-sync does not support EAPSv2.

Overview of Hitless Failover

When you install two MSM-3 modules in a BlackDiamond chassis, one MSM-3 assumes the role of master and the other assumes the role of slave. The master executes the switch’s management function, and the slave acts in a standby role. Hitless failover is a mechanism to transfer switch management control from the master to the slave.

When there is a software exception in the master, the slave may be configured to take over as the master. Without T-sync, a software exception results in a traffic “hit” because the hardware is

reinitialized and all FDB information is lost. The modules require seconds to complete the initialization, but it may take minutes to relearn the forwarding information from the network. With T-sync, it is possible for this transition to occur without interrupting existing unicast traffic flows.

During failover, the master passes control of all system management functions to the slave. In addition, hitless failover preserves layer 2 data and layer 3 unicast flows for recently routed packets. When a hitless failover event occurs, the failover timer begins and all previously established traffic flows continue to function without packet loss. Hitless failover also preserves the:

- Master's active configuration (both saved and unsaved)
- Forwarding and resolution database entries (layer 2, layer 3, and ARP)
- Loop redundancy and protocol states (STP, EAPS, ESRP, and others)
- Load shared ports
- Access control lists

NOTE

Hitless failover does not preserve the full route table, routing protocol databases for OSPF, BGP, RIP, etc., or ICMP traffic.

Hitless Failover Concepts

T-sync preserves the current active configuration across a hitless failover. When you first boot up your BlackDiamond switch, it uses the master MSM-3 configuration. During the initialization of the slave, the master's active configuration is relayed to the slave. As you make configuration changes to the master, the master relays those individual changes to the slave. When a failover occurs, the slave continues to use the master's configuration. Regardless of the number of failovers, the active configuration remains in effect provided the slave can process it.

NOTE

It is important to save any switch configuration changes that you make. Configuration changes made in real-time must be saved on the master MSM-3 to guarantee hitless failover and hitless upgrade operation. Failure to save the configuration may result in an unstable environment after the hitless failover or upgrade operation is complete.

If a hitless failover occurs before you can save the changes, the changes are still in effect on the new master MSM-3. The asterisk appears in front of the command line if unsaved configuration changes are present after a hitless failover. To save your changes after a hitless failover, use the save command.

NOTE

If you have a BlackDiamond 6816 switch populated with four MSM-3 modules, the MSMs in slots C and D provide extra switch bandwidth; they do not participate in switch management functions.

Configuring Hitless Failover

You can configure failover so that one of the following occurs:

- All links are forced to be in a down state (nothing is preserved)
- Only the configuration is preserved
- Only the link up / down state is preserved
- The configuration and link up / down states are preserved
- The configuration, link up / down states, and layer 2 FDB and states (STP, EAPS, and ESRP) are preserved
- The configuration, link up / down states, layer 2 FDB and states, and the layer 3 FDB and ARP table are preserved

Hitless failover operation utilizes the last two options. To enable hitless failover, see the following section, "Enabling Hitless Failover."

You can also configure ESRP hitless failover behavior. See "Configuring ESRP for Hitless Failover" on page 8 for more information.

To use the hitless failover feature, you must have a BlackDiamond 6800 series chassis installed with MSM-3 modules running ExtremeWare 7.1.1 or later and BootROM 8.1 or later.

Enabling Hitless Failover

To enable hitless failover, you need to:

- Configure the system recovery level to automatically reboot after a software exception
- Enable the slave MSM-3 to "inherit" its configuration from the master MSM-3
- Configure the external ports to remain active when a failover occurs
- Enable the preservation of layer 2 and / or layer 3 state in the slave MSM-3

NOTE

If you have an active Telnet session and initiate a hitless failover on that switch, the session disconnects when failover occurs.

Configuring the System Recovery Level

You must configure the slave MSM-3 to take over control of the switch if there is a software exception on the master. To configure the slave to assume the role of master, use the following command:

```
configure sys-recovery-level [all | critical] msm-failover
```

where the following is true:

- `all`—Configures ExtremeWare to log an error into the syslog and automatically reboot the system after any task exception
- `critical`—Configures ExtremeWare to log an error into the syslog and automatically reboot the system after a critical task exception

For more information about this command, see the following chapters:

- "Status Monitoring and Statistics" in the *ExtremeWare Software User Guide*
- "Commands for Status Monitoring and Statistics" in the *ExtremeWare Command Reference Guide*

Inheriting the Master's Software Configuration

To enable the slave MSM-3 to inherit the software configuration from the master MSM-3, use the following command:

```
configure msm-failover slave-config [inherited | flash]
```

where the following is true:

- `inherited`—Specifies that the slave MSM-3 inherits the software configuration maintained by the current master MSM-3 (this supports hitless failover)
- `flash`—Specifies that the slave MSM-3 use the configuration stored in its flash memory (this is the default and does not support hitless failover)

  If you enter the `flash` keyword, you cannot use the `preserve-state` option of the `configure msm-failover link-action` command.

Configuring Port Behavior and Preserving Layer 2 and Layer 3 States

In addition to enabling the use of the inherited software configuration, you need to configure the behavior of the external ports when a failover occurs. To configure the external port response, use the following command:

```
configure msm-failover link-action [keep-links-up {preserve-state [l2 | l2_l3]} |
take-links-down]
```

where the following is true:

- `keep-links-up`—Configures the external ports to not be reset when MSM failover occurs
- `preserve-state`—Configures the preservation of the link up and down states

NOTE

If you do not enter the `preserve-state` *keyword, layer 2 and layer 3 states are not preserved, and the failover or upgrade is not hitless.*

Enter one of the following options to preserver layer 2 and/or layer 3 states:

— `l2`—Preserves layer 2 FDB and states (MAC FDB, load sharing, STP, EAPS, and ESRP)

If you enter `l2`, additional ESRP configuration is required to preserve the ESRP state. By default, the ESRP failover action is *none*. Configuring the slave to "inherit" the master's configuration and setting the preserve state to l2 or l2_l3 is not sufficient to preserve the ESRP state.

See "Configuring ESRP for Hitless Failover" on page 8 for more information.

— `l2_l3`—Preserves layer 2 FDB and states plus layer 3 states (IPFDB, ARP, access lists, etc.)

- If you enter `l2_l3`, the unicast hardware IP FDB is also preserved, but the full route table and routing protocol databases for OSPF, BGP, RIP, etc. are not preserved.

  After a hitless failover has completed, the routing protocols initialize like they do on a full reboot. The neighboring routers see that the router has restarted and the neighbors re-route IP

traffic to alternate routes until the switch has reestablished it's routing databases.

Since existing IP traffic flows are preserved in the FDB, data continues to be forwarded for these flows during the start of the hitless failover and the traffic re-route. This has the effect of shortening or eliminating traffic hits for these flows.

The design of the neighboring router and/or the network traffic load determines whether a network re-routing operation is or is not hitless.

- If you enter `l2_l3`, you also need to configure ESRP for hitless failover to preserve the ESRP state. See "Configuring ESRP for Hitless Failover" on page 8 for more information.

- `take-links-down`—Configures the external ports to be reset when MSM failover occurs (this is the default and does not support hitless failover)

Configuring Timers

For switch management functions to hitlessly transition between the master and the slave, timer expiration is required. When you initiate hitless failover, the failover timer begins.

The failover timer configures the time it takes from when hitless failover begins until the relearned layer 3 databases are linked to the FDB. All FDB entries that are not linked to one of the databases at the timeout are deleted.

To configure the failover timer, use the following command:

```
configure msm-failover timeout <time>
```

The `time` parameter specifies the failover time. By default, the failover time is 60 seconds, and the range is 30 to 300 seconds.

Disabling Hitless Failover

To disable hitless failover, and return to the factory defaults, use the following command:

```
unconfigure msm-failover
```

The following occurs after you execute this command:

- The external ports are reset when an MSM failover occurs
- No state is preserved when a failover occurs
- The MSM failover timeout returns to 60 seconds
- The new master uses the configuration file kept in its flash memory upon failover

Displaying Hitless Failover Statistics

To display hitless failover statistics, use the following command:

```
show msm-failover
```

The output displays the following:

- Current state of the MSM
- Software image information (primary/secondary image, version)

- Cause of the last failover
- Failover configuration (link action, preserve state, slave configuration, timeout)
- ESRP failover mode
- Failover status for the supported subsystems (configuration, layer 2 hardware, layer 3 hardware, STP, EAPS, ARP, ESRP)

  Each of the supported subsystems display one of the following states:

  — disable—Hitless failover is disabled. This is also the initial state.

  — initial—Hitless failover is enabled, but the downloading of the subsystem state has not yet started for a particular subsystem.

  — xfr—The subsystem's state is in the process of being transferred to the slave. The state transfer includes all of the state for that subsystem.

  — ready—The subsystem has received its state download. In the ready state, it may receive updates to its internal states.

  — failed—The subsystem encountered a failure. To clear the failure, reboot the slave MSM.

  — unknown—If this state is displayed, contact Extreme Networks® Technical Support.

  — <not available>—The state and reason for the current slave shows this if the slave is in the process of being rebooted or is not present in the chassis.

After a reboot or insertion of a slave MSM-3, use this command to ensure that the slave is ready before initiating a hitless failover.

Configuring ESRP for Hitless Failover

Extreme Standby Router Protocol (ESRP) operates at both the layer 2 and layer 3 levels. An ESRP instance has the following states:

- Neutral—The initial state when ESRP is enabled.
- Slave—The slave switch is available to assume the responsibilities of the master switch if the master becomes unavailable or criteria for ESRP changes. Forwarding is disabled.
- Pre-master—The ESRP pre-master switch is ready to be master but is going through possible loop detection. Forwarding is disabled.
- Master—The ESRP master switch is responsible for responding to clients for layer 3 routing and layer 2 switching for the VLAN. Forwarding is enabled.

During the initialization of hitless failover, an ESRP instance is placed in the neutral state. Therefore, a switch that is in the master state and experiences a non-hitless failover is placed in the neutral state. This may result in a loss of traffic and the election of a new master.

To prevent standby nodes from renegotiating when the master node attempts a hitless failover, the master switch sends a notification to the standby nodes indicating a hitless failover attempt. The standby nodes increase their timeout values so they do not elect a new master. After the master recovers, it resumes normal communication with the standby nodes, and the standby nodes recognize the continued presence of the master. All unicast flows are preserved, and the ports retain the same state throughout the failover.

To configure the desired operation of hitless failover when ESRP is in use, use the following command:

```
configure msm-failover esrp-failover-mode [none | rapid-reelection |
remain-esrp-master-in-l2-domains {<reelect-timeout>}]
```

where the following is true:

- none—Specifies that ESRP does not participate in hitless failover. The master switch blocks its ports during a failover and performs a full initialization of ESRP. This is the default.
- rapid-reelection—Specifies that ESRP behaves as if none was selected. In addition, if a failover occurs when the switch is the master in an ESRP domain, the switch sends a notification that the standby nodes should elect a new master as soon as possible. This facilitates a faster ESRP master reelection than none.
- remain-esrp-master-in-l2-domains—Specifies that an ESRP master notifies the standby nodes of the failover and wishes to remain the master. Along with the notification, it sends the amount of time in seconds that the standby nodes should wait before beginning reelection.

  — reelect-timeout—Specifies the amount of time the standby nodes should wait before beginning reelection. The default is 30 seconds, and the range is 15 - 180 seconds.

  When the master finishes initializing, it resumes communication with the standby nodes, and the standby nodes revert to their standard timeout value.

ESRP Domain Behavior

The configure msm-failover esrp-failover-mode command affects all ESRP domains. However, individual domains may respond differently to hitless failover depending on circumstances and configurations.

When using the remain-esrp-master-in-l2-domains option, the behavior is hitless within an ESRP domain whenever that domain is configured with layer 2 tracking options only. If you have an ESRP domain with layer 3 tracking options, or you configure an ESRP VLAN to have both layer 2 and layer 3 tracking options, the remain-esrp-master-in-l2-domains option is overridden. Rather, the ESRP domain or VLAN assumes the behavior of the rapid-reelection option.

Table 1: Hitless failover support for ESRP tracking options

Tracking Option	Hitless Support
Diagnostic	Yes
Environment	Yes
VLAN	Yes
OSPF, BGP, RIP	No
IP Route	No
Ping	No

For more information about ESRP and ESRP tracking, see the chapter “Extreme Standby Routing Protocol” in the *ExtremeWare Software User Guide*.

Displaying ESRP Hitless Failover Statistics

To display ESRP hitless failover statistics, use the following command:

```
show esrp detail
```

The output varies depending upon the configuration and the state of the switch:

- Standby switch—Information about the impending failover and the timeout is displayed
- Layer 3 tracking in use and the failover mode is `remain-esrp-master-in-l2-domain`—Information about rapid reelection and layer 3 tracking is displayed
- Layer 3 tracking is not in use and the failover mode is `remain-esrp-master-in-l2-domain`—Information about remaining the master is displayed
- `rapid-reelection`—Information about rapid reelection is displayed
- none—Information about not participating in hitless failover is displayed

Overview of Hitless Upgrade

As described previously, when you install two MSM-3 modules in a BlackDiamond chassis, one assumes the role of master and the other assumes the role of slave. The master executes the switch's management function, and the slave acts in a standby role. Hitless upgrade (a component of T-sync) is a mechanism that allows an upgrade of the ExtremeWare version running on a BlackDiamond chassis without:

- Taking the switch out of service
- Losing traffic
- Interrupting network operation

NOTE

It is important to save any switch configuration changes that you make. Configuration changes made in real-time must be saved on the master MSM-3 to guarantee hitless failover and hitless upgrade operation. Failure to save the configuration may result in an unstable environment after the hitless failover or upgrade operation is complete.

You perform a hitless upgrade by downloading the new software image, selecting it, and then forcing a hitless failover to occur. This guide describes two methods that you can use to perform a hitless upgrade:

- Standard
- Conservative

Each method results in an upgrade to the new version of ExtremeWare; the difference is which version is executed if there is another hitless failover. When you perform any upgrade, pre-established flows remain active, and new flows take additional time.

NOTE

If you have a BlackDiamond 6816 switch populated with four MSM-3 modules, the MSMs in slots C and D provide extra switch bandwidth; they do not participate in switch management functions.

To use the hitless upgrade feature, you must have a BlackDiamond 6800 series chassis installed with MSM-3 modules running ExtremeWare 7.1.1 or later and BootROM 8.1 or later.

Standard Upgrade

One method of hitless upgrade is a standard upgrade. A *standard* upgrade causes the new version of ExtremeWare to execute on any failover that occurs *after* the software upgrade. Perform a standard upgrade after you qualify the new ExtremeWare release for your network.

Use the standard approach after you test the new software release and you are ready to implement the new software across your entire network.

Conservative Upgrade

Another method of hitless upgrade is a conservative upgrade. A *conservative* upgrade causes the previous version of ExtremeWare to execute on any failover that occurs *after* the upgrade. Perform a conservative upgrade when you want to "try" a new version of ExtremeWare.

Extreme Networks recommends using the conservative approach to:

- Test a new software release before deploying it across your entire network
- Test a new software patch
- Qualify the new ExtremeWare release for your network

If you are not ready to deploy the new software, you can hitlessly restore the previous software version and network configuration. Use the conservative method to test a new ExtremeWare release. Do not use this method to run two different versions of ExtremeWare on the master and slave for an extended period of time.

Performing a Hitless Upgrade

You can perform a hitless software upgrade on an MSM-3 without interrupting network operation. This section describes two software upgrade methods and includes a troubleshooting section:

- Standard Software Upgrade on page 11
- Conservative Software Upgrade on page 13
- Troubleshooting on page 14

Standard Software Upgrade

The steps described in this section assume the following:

- MSM-3 installed in slot A is the master
- Primary image is in use

- MSM-A is the MSM-3 installed in slot A
- MSM-B is the MSM-3 installed in slot B
- You are running ExtremeWare 7.1.1 or later and BootROM 8.1 or later (see the *ExtremeWare 7.1.1 Release Notes* for more information)
- You configured the system for hitless failover operation (see "Configuring Hitless Failover" on page 5 for more information)

NOTE

If you have an active Telnet session and initiate a hitless failover on that switch, the session disconnects when failover occurs.

To failover to the same, new software image, do the following:

1. Download the new software image to the primary image space using the following command:

   `download image [<hostname> | <ip address>] <filename>`

   where the following is true:

   - `hostname`—Specifies the hostname of the TFTP server from which the image should be obtained (DNS must be enabled to use this option)
   - `ip address`—Specifies the IP address of the TFTP server from which the image should be obtained
   - `filename`—Specifies the filename of the new software image

   The primary image loads on both the master and slave MSM.

2. To use the new software image on the slave, you must reboot the slave before the failover can take place.

   Reboot MSM-B using the following command:

   `reboot slot msm-b`

   where `msm-b` specifies the slave MSM-3 installed in slot B.

   After you reboot the slave, the new software image begins running on the slave.

   The master downloads it's configurations, FDB entries, etc. to the slave. After the master finishes it's download, the following message is logged:

   `Slave MSM initialized for hitless failover operation.`

3. Verify that the correct BootROM and ExtremeWare version are loaded using the `show switch` and `show version` commands.

4. Initiate failover using the following command:

   `run msm-failover`

   MSM-B becomes the master and runs the new software image. MSM-A becomes the slave and also runs the new software image.

5. Verify the slave state using the following command:

   `show msm-failover`

   You can also verify the slave state by viewing the syslog. To view the syslog, use the following command:

   `show log`

Conservative Software Upgrade

The steps described in this section assume the following:

- MSM-3 installed in slot A is the master
- Primary image is in use
- MSM-A is the MSM-3 installed in slot A
- MSM-B is the MSM-3 installed in slot B
- You are running ExtremeWare 7.1.1 or later and BootROM 8.1 or later (see the *ExtremeWare 7.1.1 Release Notes* for more information)
- You configured the system for hitless failover operation (see "Configuring Hitless Failover" on page 5 for more information)

NOTE

If you have an active Telnet session and initiate a hitless failover on that switch, the session disconnects when failover occurs.

To failback to the previous software image, do the following:

1. Choose the secondary image of MSM-A using the following command:

   `use image secondary`

2. Download the new software image to the secondary image space using the following command:

   `download image [<hostname> | <ip address>] <filename>`

   where the following is true:

   - `hostname`—Specifies the hostname of the TFTP server from which the image should be obtained (DNS must be enabled to use this option)
   - `ip address`—Specifies the IP address of the TFTP server from which the image should be obtained
   - `filename`—Specifies the filename of the new software image

3. Verify that the correct BootROM and ExtremeWare version are loaded using the `show switch` and `show version` commands.

4. Force the slave to boot up using the new software version. This must happen before failover can occur.

   Reboot MSM-B using the following command:

   `reboot slot msm-b`

   where `msm-b` specifies the slave MSM-3 installed in slot B.

   After the slave is reset, a series of log messages are displayed:

   - The first message indicates that MSM-B has initialized as a slave
   - Another message appears for each database download (configurations, FDB entries, etc.)
   - When all of the databases are downloaded, a message indicates that the slave is fully prepared for execution of a failover

5. Select the primary image (that now contains the older software release) using the following command:

   `use image primary`

6 Initiate failover using the following command:

```
run msm-failover
```

The failover allows MSM-B to become the master running the new software release. Because the current image selected is primary, MSM-A reboots to the old release. MSM-B then downloads the original configuration from MSM-A back to MSM-A.

NOTE

If you issue any configuration command (configuration, enable, disable, create, etc.) after a conservative upgrade, you are unable to downgrade to the previous software release.

By using the conservative method, you can upgrade or downgrade software without disrupting the network. This method allows you to test new software releases because you:

- Will not disrupt the network
- Can failback to your previous software version and configurations saved to MSM-A

If you want to failback to MSM-A after you test a software release, enter the `run msm-failover` command. MSM-A becomes the master and runs the software and configurations you had in place before the upgrade.

After you qualify the new software release and determine that you do not want to failback to the previous software release, enter the `use image secondary` command. The next time you reboot the switch, both MSM-3s run the new software image.

Troubleshooting

If you encounter problems during a hitless upgrade, this section may be helpful. If you have a problem not listed here or in the release notes, contact your Extreme Networks Technical Support representative.

MSM-3 Initialization

If the slave MSM-3 does not initialize properly, you can restart the MSM hardware, including the switch fabric, using the following command:

```
reboot slot [msm-a | msm-b] hard-reset
```

where the following is true:

- `msm-a`—Specifies the slave MSM-3 module installed in slot A.
- `msm-b`—Specifies the slave MSM-3 module installed in slot B.
- `hard-reset`—Restarts the MSM hardware, processor, and switch fabric. If you select this option, you will experience some traffic loss.

NOTE

If you enter the `hard-reset` *option, a hitless upgrade does not occur.*

After you enter the `hard-reset` option, a warning message is displayed that indicates some of the switch forwarding operations may be briefly interrupted, and you are asked to confirm the operation. If you do not want to interrupt switch forwarding, do not confirm the operation.

Unexpected Failover Results

If you experience unexpected failover results, use the following command to help determine the reason:

```
show msm-failover
```

Information about the state of the new master and the current slave is displayed:

- Old states and old reasons of the new master—Helps determine the state of the current master (former slave) before the last failover when it became the master.
- Current states and current reasons of the new slave—Helps determine the current state of the slave.

The following three sample scenarios describe how you can use the `show msm-failover` command to help troubleshoot unexpected failover results.

Scenario 1. If you perform a conservative upgrade and later configure the switch, this causes an unexpected failover result. Issuing any configuration command (configuration, enable, disable, create, etc.) after a conservative upgrade prevents you from downgrading to the previous software release on the slave.

If this happens, you can change the image of the slave to the new software release and reboot the slave. By doing this, the new software image runs on both the master and the slave. After some time, the `show msm-failover` command indicates that the slave is in the ready state and failover to the same software release is possible.

Scenario 2. If a failover occurs and there was a network hit, use the `show-msm failover` command to view the output to find out why the hit occurred.

Scenario 3. If the slave is in the failed state, hitless failover cannot occur. Use the `show-msm failover` command to view the output to display the reason for the failure.

Table 2 describes the failover reason codes displayed with the `show msm-failover` command.

Table 2: Descriptions of failover reason codes

Reason Code	Description
none	No failure occurred.
rev(M) > rev(S)	An older version slave is present, but a hitless upgrade from the slave to the master was not performed.
hotswap	A slave was removed after performing a conservative upgrade.
config command	A configuration command was entered after a conservative upgrade. This prevents you from downgrading to the pervious software release on the slave.
memory	Contact Extreme Networks Technical Support.
brkt ovflow	Contact Extreme Networks Technical Support.
invalid brkt	Contact Extreme Networks Technical Support.
invalid subtype	Each subsystem classifies its messages between the master and slave. If the slave does not recognize a message classifier, an error occurs.
comm error	This error message is currently not in use. If you see this message, contact Extreme Networks Technical Support.
no xh support	This error message is currently not in use. If you see this message, contact Extreme Networks Technical Support.
L2,L3 failed	Contact Extreme Networks Technical Support.

Table 2: Descriptions of failover reason codes (continued)

Reason Code	Description
config failed	This error message is currently not in use. If you see this message, contact Extreme Networks Technical Support.
keepalive	Contact Extreme Networks Technical Support.
watchdog	Contact Extreme Networks Technical Support.