

Extreme Network OS Software Upgrade Guide

Supporting Network OS 7.1.0

© 2018, Extreme Networks, Inc. All Rights Reserved.

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries. All other names are the property of their respective owners. For additional information on Extreme Networks Trademarks please see www.extremenetworks.com/company/legal/trademarks. Specifications and product availability are subject to change without notice.

Contents

Preface	5
Document conventions.....	5
Notes, cautions, and warnings.....	5
Text formatting conventions.....	5
Command syntax conventions.....	6
Extreme resources.....	6
Document feedback.....	6
Contacting Extreme Technical Support.....	7
About this document	9
Supported hardware and software.....	9
Using the Network OS CLI	9
What's new in this document.....	9
Installing and Maintaining Firmware	11
Firmware management overview.....	11
Upgrading firmware on a modular chassis.....	11
Upgrading firmware on a modular chassis.....	12
Automatic firmware synchronization.....	12
Preparing for a Firmware Download	13
Prerequisites.....	13
Obtaining and decompressing firmware.....	13
Basic Firmware Upgrade	15
Upgrading firmware on a local device.....	15
Upgrading considerations and restrictions.....	15
IGMP snooping upgrade and downgrade considerations.....	16
Connecting to the device.....	16
Obtaining the firmware version.....	17
Using the firmware download command.....	17
Downloading firmware using ISSU.....	18
Downloading firmware using the coldboot option.....	19
Downloading firmware using the default-config option.....	20
Downloading firmware from a USB device.....	20
Downloading firmware using the noactivate option.....	21
Downloading firmware using the manual option.....	21
Verifying a firmware download session.....	22
Upgrading firmware in a logical-chassis.....	22
Overview of firmware download logical-chassis.....	22
Using ISSU with firmware download logical-chassis.....	23
Using auto-active with firmware download logical-chassis.....	24
Using coldboot with firmware download logical-chassis	25
Using default-config with firmware download logical-chassis.....	26
Updating the peripheral firmware.....	27
Advanced Upgrade Scenarios	29
Upgrading firmware within a VCS Fabric.....	29
Tested topology.....	29

Upgrading nodes by using an odd/even approach.....	31
Preparing for the maintenance window.....	31
Optimizing reconvergence in the VCS Fabric.....	35
Maintaining the VCS Fabric.....	35
Understanding traffic outages.....	37

Preface

- Document conventions..... 5
- Extreme resources..... 6
- Document feedback..... 6
- Contacting Extreme Technical Support..... 7

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Extreme technical documentation.

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

NOTE

A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION

An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.



CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.



DANGER

A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used to highlight specific words or phrases.

Format	Description
bold text	Identifies command names. Identifies keywords and operands. Identifies the names of GUI elements.
<i>italic text</i>	Identifies text to enter in the GUI. Identifies emphasis. Identifies variables.
Courier font	Identifies document titles. Identifies CLI output.

Format	Description
	Identifies command syntax examples.

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member[member...]</i> .
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Extreme resources

Visit the Extreme website to locate related documentation for your product and additional Extreme resources.

White papers, data sheets, and the most recent versions of Extreme software and hardware manuals are available at www.extremenetworks.com. Product documentation for all supported releases is available to registered users at www.extremenetworks.com/support/documentation.

Document feedback

Quality is our first concern at Extreme, and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you.

You can provide feedback in two ways:

- Use our short online feedback form at <http://www.extremenetworks.com/documentation-feedback-pdf/>
- Email us at internalinfodev@extremenetworks.com

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Contacting Extreme Technical Support

As an Extreme customer, you can contact Extreme Technical Support using one of the following methods: 24x7 online or by telephone. OEM customers should contact their OEM/solution provider.

If you require assistance, contact Extreme Networks using one of the following methods:

- [GTAC \(Global Technical Assistance Center\)](#) for immediate support
 - Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact.
 - Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- [GTAC Knowledge](#) - Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- [The Hub](#) - A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- [Support Portal](#) - Manage cases, downloads, service contracts, product licensing, and training and certifications.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

About this document

- Supported hardware and software.....9
- Using the Network OS CLI9
- What's new in this document.....9

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some devices but not to others, this guide identifies exactly which devices are supported and which are not.

Although many different software and hardware configurations are tested and supported by Extreme Networks, Inc. for Network OS, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by this release of Network OS:

- ExtremeSwitching VDX 2746
- ExtremeSwitching VDX 6740
 - ExtremeSwitching VDX 6740-48
 - ExtremeSwitching VDX 6740-64
- ExtremeSwitching VDX 6740T
 - ExtremeSwitching VDX 6740T-48
 - ExtremeSwitching VDX 6740T-64
 - ExtremeSwitching VDX 6740T-1G
- ExtremeSwitching VDX 6940-36Q
- ExtremeSwitching VDX 6940-144S
- ExtremeSwitching VDX 8770
 - ExtremeSwitching VDX 8770-4
 - ExtremeSwitching VDX 8770-8

To obtain information about a Network OS version other than this release, refer to the documentation specific to that version.

Using the Network OS CLI

For complete instructions and support for using the Extreme Network OS command line interface (CLI), refer to the *Extreme Network OS Command Reference*.

What's new in this document

This document describes the concepts and configuration of the upgrade and downgrade processes for Network OS.

On October 30, 2017, Extreme Networks, Inc. acquired the data center networking business from Brocade Communications Systems, Inc. This document has been updated to remove or replace references to Brocade Communications, Inc. with Extreme Networks, Inc., as appropriate.

The content has been updated with the following changes:

- Updated to support upgrading from previous versions to Network OS v7.1.0.
- New support for updating the peripheral firmware on some devices.

Installing and Maintaining Firmware

- [Firmware management overview](#)..... 11
- [Upgrading firmware on a modular chassis](#)..... 11
- [Upgrading firmware on a modular chassis](#)..... 12

Firmware management overview

Extreme firmware upgrades consist of multiple firmware packages listed in a **.plist** file. The **.plist** file contains specific firmware information (time stamp, platform code, version, and so on) and the names of the firmware packages to be downloaded. These packages are made available periodically to add features or to remedy defects in the firmware. In Network OS 4.0.0 and later, firmware upgrade is performed incrementally. The **firmware download** command compares the new firmware packages against the current installation and only downloads the packages that contain new features or have been modified.

Network OS provides a single command line interface (CLI) to download firmware to a Top-of-Rack (ToR) device with a single control processor or to a modular chassis with two management modules. You can download the firmware from a remote server by means of the File Transfer Protocol (FTP), SSH File Transfer Protocol (SFTP), or the Secure Copy Protocol (SCP), or you can download the firmware from an attached Extreme-branded USB device. If you want to download firmware from a remote server, you must connect the management Ethernet port of the device to the server. In a modular chassis, both management Ethernet ports need to be connected.

Refer to the respective NOS-version release notes for ISSU and upgrade-path information. An ISSU allows a dual management module system or Top of Rack devices to be upgraded non-disruptively and is invoked by entering the firmware download command from the active management module.

In Network OS v4.0.0 and later, the logical-chassis firmware download command allows you to upgrade a single device or multiple devices of your choice that are connected in logical chassis mode. This command can only be executed from the principal node (coordinator). The firmware can only be downloaded from the file server through the management Ethernet port, so all nodes must have the management Ethernet ports connected. Only one **logical-chassis firmware download** command instance can run at any given time.

In Network OS v4.1.0, the one-version upgrade and downgrade is no longer enforced, (for example, the need to downgrade from Network OS v4.1.0 to v4.0.0, in order to download Network OS v3.1.0), and you can skip versions when performing upgrades and downgrades, (for example, downgrading from Network OS v4.1.0 to v3.1.0); however, the previous configurations are not preserved after the upgrade or downgrade. This new capability is available using the **firmware download default-config** command.

If you are in logical chassis mode, after you perform a firmware upgrade, you may find that the device reverts to its default configurations. To preserve the configurations after an upgrade, back up the configuration using the **copy running-config** command before the firmware download. After the upgrade is completed, run the **copy running-config** command.

If a firmware download session is interrupted by an unexpected reboot, Network OS attempts to recover the previously installed firmware. Success depends on the state of the firmware download. You must wait for the recovery to complete before initiating another firmware download.

Upgrading firmware on a modular chassis

In Network OS v5.0.0 and later, in-service software upgrade (ISSU) is supported on a Top of Rack device (ToR). An ISSU allows a ToR to be upgraded non-disruptively and is invoked when you enter the **firmware download** command without any options. Refer to [Downloading firmware using ISSU](#) on page 18

Upgrading firmware on a modular chassis

In Network OS 4.0.0 and later in-service software upgrade (ISSU) is supported. An ISSU allows a dual management module (MM) system to be upgraded non-disruptively and is invoked when you enter the firmware download command from the active management module.

Network OS 7.2.0 supports upgrading from Network OS 4.1.x directly to Network OS 7.2.0. This feature is not supported on versions later than Network OS 4.1.x.

Automatic firmware synchronization

When you replace or insert a second management module into a chassis, the active management module automatically synchronizes the hot-plugged standby management module with the same firmware version. The standby management module reboots with the upgraded firmware. The automatic firmware synchronization takes place only if all of the following conditions are met:

- The standby management module is inserted while the chassis is already up (hot-plugged insert).
- There was no firmware download process running when the standby management module was inserted.
- The active and standby firmware versions must be different.

NOTE

Automatic firmware synchronization is intrinsic to Network OS v4.0.0 and later and no corresponding **enable** or **disable** commands are associated with automatic firmware synchronization. As a result, automatic firmware synchronization cannot be disabled.

Preparing for a Firmware Download

- Prerequisites..... 13
- Obtaining and decompressing firmware..... 13

Prerequisites

To prepare for a firmware download, perform the following tasks. In the unlikely event of a failure or timeout, you will be able to provide your device support provider the information required to troubleshoot the firmware download.

1. Verify the current firmware version. Refer to [Obtaining the firmware version](#) on page 17.
2. Download the firmware package from the Extreme website to an FTP server.
3. Decompress the firmware archive. Refer to [Obtaining and decompressing firmware](#) on page 13.
4. Decide on a migration path. Check the connected devices to ensure firmware compatibility and that any older versions are supported. Refer to the “Network OS Compatibility” section of the *Network OS Release Notes* for the recommended firmware version.
5. In a modular system, if you are to download firmware from a file server, verify that the management ports on both MMs are connected to the firmware file server.
6. Back up your device configuration prior to the firmware download. Refer to [Installing and Maintaining Firmware](#) on page 11 for details.
7. For additional support, connect the device to a computer with a serial console cable. Ensure that all serial consoles and any open network connection sessions, such as Telnet, are logged and included with any trouble reports.
8. Enter the **copy support** command to collect all current core files prior to executing the firmware download. This information helps to troubleshoot the firmware download process in the event of a problem. Once the **copy support** command collects the files, you can use the **clear support** command to remove the files from the list.
9. Enter the **clear logging raslog** command to erase all existing messages in addition to internal messages.

Obtaining and decompressing firmware

Firmware upgrades are available for customers with support service contracts and for partners on the Extreme website.

You must download the firmware package to an FTP-variant server and decompress the package *before* you can use the **firmware download** command to upgrade the firmware on your equipment.

You may also download the firmware from a USB drive using the **firmware download usb** command.

When you unpack the downloaded firmware, it expands into a directory that is named according to the firmware version. When issued with the path to the directory where the firmware is stored, the **firmware download** command performs an automatic search for the correct package file type associated with the device.

Five **firmware download** command options are available:

- **coldboot**: Downloads the firmware to the system and reboots the device.
- **default-config**: Removes all configuration and is similar to an initial installation and configuration.
- **usb**: Downloads the firmware to the system without activating it, so the device is not automatically rebooted.
- **noactivate**: Downloads the firmware to the system without activating it, so the device is not automatically rebooted.

- `nocommit`: Disables auto-commit mode. When auto-commit mode is disabled, firmware is downloaded only to the primary partition.
- `noceboot`: Disables auto-reboot mode. When auto-reboot mode is disabled, you must reboot the device manually.

Refer to the *Extreme Network OS Command Reference* for complete information on all of the available options for the **firmware download** command.

NOTE

To be able to address the FTP server by its name, ensure that a Domain Name System (DNS) entry is established for the server.

NOTE

Network OS does not support the use of special characters (such as `&`, `!`, `%`, or `#`) in FTP, TFTP, SFTP, or SCP passwords. If your password contains special characters, the download fails.

Basic Firmware Upgrade

- Upgrading firmware on a local device..... 15
- Upgrading considerations and restrictions..... 15
- Connecting to the device..... 16
- Obtaining the firmware version..... 17
- Using the firmware download command..... 17
- Upgrading firmware in a logical-chassis..... 22
- Updating the peripheral firmware..... 27

Upgrading firmware on a local device

A basic firmware download upgrades the local device only. This section explains how to use the **firmware download** command and its various options in a local device firmware upgrade.

Upgrading considerations and restrictions

Consider the following when upgrading your firmware version:

- UpWhen upgrading to Network OS v7.0.1, three-tuple support for SNMP server is not active by default. You must execute the **snmp-server three-tuple-if enable** command to enable this feature.
- Prior to upgrading from Network OS v6.0.2x, the IGMP snooping static mrouter feature must be deactivated from all snoop or non-snoop enabled VLANs.
- Hardware profile support for TCAM and Routing Table is supported when upgrading from Network OS v5.x to v6.x. During the upgrade process from Network OS v5.x to Network OS v6.x, the DB conversion and startup file replay on hardware profile configuration is fully supported. During Network OS v6.0 to Network OS v5.x downgrade, DB conversion and startup file replay on profile configuration is also fully supported on the existing platforms which are supported by Network OS v5.x.
- If you are upgrading directly from Network OS v5.0.0 to Network OS v6.0.1, your device configuration information will be deleted.
- The interface configurations, **no fabric isl enable** and **no fabric trunk enable** can revert back to the default values of "fabric isl enable" and "fabric trunk enable" after upgrading from Network OS v4.1x to v5.0.0 or later. In order to keep these features disabled, Extreme recommends configuring each interface using the fabric neighbor-discovery disable command instead.
- Upgrading to Network OS v4.0 or later is automatically allowed because the Telnet server and SSH server status are enabled by default.
- Upgrading from Network OS v5.0.0 or Network OS v4.x to Network OS v5.0.2 or Network OS v5.0.1d with any SNMP V3 users with auth/priv security level associated with "snmpadmin" group are moved to new group created as "admin". The read-write-notify permissions are given for that "admin" group after upgrade. Any V3 users with auth/priv security level associated with the "snmpuser" group are moved to new group created as "user" and the read-notify permissions are given for that user group after the upgrade is complete.

SNMP V3 users with noauth security level associated with snmpadmin/snmpuser group are not be moved to admin/user groups after upgrade. After upgrading to Network OS v5.0.2, the snmpadmin/snmpuser group still has only notify permission, in order to block read/write access for noauth users and maintain backward compatibility.

IGMP snooping upgrade and downgrade considerations

The following table lists the IGMP snooping upgrade and downgrade considerations on VLANs across Network OS 6.0.1 and Network OS 7.x.x.

TABLE 1 IGMP snooping upgrade and downgrade considerations

Network OS 6.0.1	Network OS 7.x.x
Fewer than 512 VLANs are configured with global snooping enabled.	IGMP snooping configuration will be present at the global and VLAN levels after upgrade.
Fewer than 512 snooping-enabled VLANs configured. Global snooping is not enabled. Only VLAN-level snooping is enabled.	Global snooping is automatically enabled after upgrade.
More than 512 snooping-enabled VLANs, irrespective of whether global snooping is enabled or not.	<p>If the number of snooping-enabled VLANs is more than 512 in Network OS 6.0.x, snooping will be disabled on all VLANs.</p> <ul style="list-style-type: none"> If global snooping is enabled and more than 512 VLANs are configured, after the upgrade, you must enable snooping on the VLANs. If global snooping is not enabled, and if there are more than 512 snooping-enabled VLANs, after upgrading, you must enable VLAN-level and global-level IGMP snooping. This is because only 512 snooping-enabled VLANs are supported across the fabric. If this number was more than 512 in Network OS 6.0.x, snooping on VLANs will be disabled during upgrade.
Downgrading from 7.x.x to 6.0.x	<p>Downgrade is not allowed if IGMP snooping is enabled either at the global or VLAN levels. The following messages display when the downgrade command is executed:</p> <p>Only global snooping enabled</p> <p>Message: - Downgrade is not allowed because global IGMP snooping is enabled. Please disable global IGMP snooping.</p> <p>Only VLAN level snooping enabled</p> <p>Message: - Downgrade is not allowed because IGMP snooping is enabled on one or more VLANs. Please disable IGMP snooping on VLANs.</p> <p>Global and VLAN-level snooping enabled</p> <p>Message: - Downgrade is not allowed because IGMP snooping is enabled globally and on one or more VLANs. Please disable global IGMP snooping and VLAN-level IGMP snooping.</p>

Connecting to the device

When you upgrade firmware in default mode, you connect to the device through the management IP address. Modular devices have one management IP address for the chassis and separate IP addresses for each management module. To upgrade both management modules, you can either connect to the chassis management IP address or to the IP address of the active management module. If you want to upgrade a single management module only, you must connect to the IP address of that management module and run the **firmware download** command in manual mode. In manual mode, only the local management module is upgraded.

Use the **show system** command to display the management IP address for the chassis.

```
device# show system
Stack MAC                : 00:05:33:15:FA:70
  -- UNIT 0 --
Unit Name                 : sw0
Switch Status             : Online
Hardware Rev              : 1000.0
TengigabitEthernet Port(s) : 56
Up Time                   : up 8:38
Current Time              : 16:39:56 GMT
```



```

NOS Version           : 5.0.0
Jumbo Capable        : yes
Burned In MAC        : 00:05:33:15:FA:70
Management IP        : 10.24.73.131 <- Chassis Management IP address
Management Port Status : UP

```

Use the **show interface management** command to display the IP addresses for the management modules.

```

device# show interface management
interface Management 10/1
 ip address 10.24.73.130/20
 ip gateway-address 10.24.64.1
 ipv6 ipv6-address [ ]
 ipv6 ipv6-gateways [ ]
 line-speed actual "1000baseT, Duplex: Full"
 line-speed configured Auto
interface Management 10/2
 ip address 10.24.74.23/20
 ip gateway-address 10.24.64.1
 ipv6 ipv6-address [ ]
 ipv6 ipv6-gateways [ ]
 line-speed actual "1000baseT, Duplex: Full"
 line-speed configured Auto

```

NOTE

You must configure the gateway and default route that is pointing to the management interface within the management VRF and address-family unicast context.

Obtaining the firmware version

Enter the **show version** command with the **all-partitions** option to obtain the firmware version for both primary and secondary partitions of each module.

```

device# show version all-partitions

Network Operating System Software
Network Operating System Version: 5.0.0
Copyright (c) 1995-2014 Brocade Communications Systems, Inc.
Firmware name:      5.0.0_radius
Build Time:         14:06:41 Apr 16, 2014
Install Time:       05:19:44 Apr 29, 2014
Kernel:             2.6.34.6

BootProm:           2.2.0
Control Processor:  e500v2 with 2048 MB of memory

Appl      Primary/Secondary Versions
-----
NOS       5.0.0_radius
          5.0.0_radius

```

Using the firmware download command

Four upgrade options are available:

- ISSU—In-Service Software Upgrade (ISSU) for dual-MM nodes (since Release 4.0) and TOR devices, is a non-disruptive upgrade option, for upgrading to a patch or maintenance release within the same major and minor release. ISSU provides the only upgrade option that is nondisruptive for active FC/FCoE sessions. FC/FCoE sessions are affected by other upgrade options. Use ISSU for qualifying nodes in combination with the Procedure for Upgrading Odd/Even Nodes, which is still required for single-MM and other nodes that do not qualify for ISSU.

- Coldboot—Coldboot is the standard option for dual-MM system for all firmware upgrades between major releases.
- Default-config—Removes all configuration and is similar to an initial installation and configuration.
- Manual - This option is not recommended. It uses the noactivate and nocommit options to perform firmware upgrades and downgrades with multiple steps.

NOTE

To be able to address the FTP server by its name, ensure that a Domain Name System (DNS) entry is established for the server.

NOTE

Network OS does not support the use of special characters (such as & ! % #) in FTP and SCP passwords. If your SCP or FTP password contains special characters, the download fails.

Downloading firmware using ISSU

If you enter the **firmware download** command without any options, the command invokes ISSU to upgrade the entire system. On a modular chassis, if you enter the **firmware download** command on the active MM without any options, the command invokes the ISSU process to upgrade the entire system.

NOTE

ISSU is not supported when downloading to versions previous to Network OS 7.0.0. Please specify the **coldboot** option in the command-line for downloads.

If you decide to invoke other **firmware download** command options, refer to the following:

- [Downloading firmware using the noactivate option](#) on page 21
- [Downloading firmware using the manual option](#) on page 21
- [Downloading firmware using the coldboot option](#) on page 19
- [Downloading firmware using the default-config option](#) on page 20

To download firmware from an attached USB device, refer to [Downloading firmware from a USB device](#) on page 20.

When upgrading multiple devices, complete the following steps on each device before you upgrade the next one.

1. Perform the steps described in [Prerequisites](#) on page 13.
2. Verify that the FTP or SSH server is running on the remote server and that you have a valid user ID and password on that server.
Network OS does not support the use of special characters (such as & ! % #) in FTP and SCP passwords. If your SCP or FTP password contains special characters, the download fails.
3. Connect to the device or management module you are upgrading.
Refer to [Connecting to the device](#) on page 16 for more information.
4. Issue the **show version** command to determine the current firmware version.
5. Enter the **firmware download interactive** command to download the firmware interactively. When prompted for input, choose the defaults whenever possible.

6. If you invoked the **firmware download** command using the **interactive** option, at the Do you want to continue? [y/n]: prompt, enter y.

```
device# firmware download interactive
Download to multiple nodes in the cluster? [n]:
Server name or IP address: 10.70.4.106
File name: dist
Protocol (ftp, scp, sftp, tftp) [ftp]: scp
User: fvt
Password: *****
Enter VRF name[mgmt-vrf]:
Select procedure (1=ISSU, 2=coldboot, 3=default-config) [1]:1
```

Performing system sanity check...

This command will use the ISSU protocol to upgrade the system. It will cause a WARM reboot and will require that existing telnet, secure telnet or SSH sessions be restarted

Do you want to continue? [y/n]y

Downloading firmware using the coldboot option

The coldboot option of the firmware download command allows you to download new firmware onto a device and forces the device to perform a cold reboot.

In a chassis system, this option downloads the firmware on both the active and standby MMs and reboots both of the MMs at the same time. After the firmware completes downloading on both MMs, they are rebooted at the same time. This ensures that both MMs reboot with the same firmware, and prevents any firmware compatibility issues that may exist between the old and the new firmware.



WARNING

This option causes traffic disruption.

1. Download the firmware from the source directory with the coldboot option. The following example is generic and not related to any specific Extreme software platform.

```
device# firmware download scp host 10.70.4.109 user fvt directory /buildsjc/sre/SQA/nos/os1.0.0/
osl.0.0_bld01 password pray4green coldboot
Performing system sanity check...
```

This command will cause a cold/disruptive reboot and will require that existing telnet, secure telnet or SSH sessions be restarted.

Do you want to continue? [y/n]:y

2. After the device completes the reboot sequence, you may log in to the device and operate normally.
3. Log back into the device. Enter the **firmware commit** command to commit the new firmware. If you entered y after the prompt, the device will commit the firmware automatically upon booting up.

```
device# firmware commit
Validating primary partition...
Doing firmwarecommit now.
Please wait ...
Replicating kernel image
.....
FirmwareCommit completes successfully.
```

4. Enter the **show version** command. Both partitions on the device or on the modules should contain the new firmware.

Downloading firmware using the default-config option

The **firmware download default-config** command allows you to download new firmware onto the device, clean up the configuration, and then force the device to perform a cold reboot.

This option is useful to prevent issues caused by incompatible configurations between the old and the new firmware.



CAUTION

When you use **firmware download default-config**, traffic is disrupted and the configuration is lost. You must save the configuration information before you execute the command and then restore it afterwards.

You can download firmware on a local device and optionally change the VCS mode, the VCS ID, and the RBridge ID before rebooting the device with the new firmware.

To download firmware by using the **default-config** option with VCS mode 1, VCS ID 7, and RBridge 10, use the following command :

```
device# firmware download default-config ftp host 10.20.1.3 user fvt password pray4green directory dist
file release.plist vcs-mode 1 vcs-id 7 rbridge-id 10

Performing system sanity check...
This command will set the configuration to default and set the following parameters: vcs-mode, vcs-id and
rbridge-id.
This command will cause Cold reboot on both MMs at the same time and will require that existing telnet,
secure telnet or SSH sessions be restarted.

Do you want to continue? [y/n]: y
```

Downloading firmware from a USB device

Extreme devices support firmware download from a Extreme-branded USB device. You cannot use a third-party USB device. Before you can access the USB device, you must enable the device and mount it as a file system. The firmware images to be downloaded must be stored in the factory-configured `firmware` directory. Multiple images can be stored under this directory.

1. Ensure that the USB device is connected to the device.
2. Enter the **usb on** command in privileged EXEC mode.

```
device# usb on
Trying to enable USB device. Please wait...
USB storage enabled
```

3. Enter the **usb dir** command. In this sample output, the "X" refers to the current version number.

```
device# usb dir
firmwarekey\ 0B 2013 Jun 15 15:13
support\ 106MB 2013 Jun 24 05:36
config\ 0B 2013 Jun 15 15:13
firmware\ 380MB 2013 Jun 15 15:13
NOS_vX.X.X\ 379MB 2013 Jun 15 15:31
Available space on usbstorage 74%
```

4. Enter the **firmware download usb** command followed by the relative path to the firmware directory, where the "X" refers to the current version number.

```
device# firmwaredownload usb directory NOS_vX.X.X
```

5. Unmount the USB storage device.

```
device# usb off
Trying to disable USB device. Please wait...
USB storage disabled.
```

Downloading firmware using the noactivate option

The **noactivate** option in the **firmware download** command allows you to download new firmware onto a device without rebooting the system. In a chassis system, you use this option on the active MM only; the firmware is then downloaded onto both the active and standby MMs.

This option is normally used for the ISSU method.

The following example shows the results of the **firmware download noactivate** command.

```
device# firmware download scp noactivate host 10.70.12.110 directory /users/home24/
smith/smith500 user fvt password pray4green
Performing system sanity check...
You are running firmware download without Activating the downloaded firmware. Please
use firmware activate to activate the firmware.
Do you want to continue? [y/n]: y
```

After the new firmware is downloaded, you can later execute the **firmware activate** command on the device to reboot the device and activate the new firmware.



CAUTION

Do not execute the reboot command to activate the new firmware. Doing so causes the old firmware to be restored.

The following example shows a request to activate the node after running **firmware activate** command.

```
device# firmware activate
This command will use the ISSU protocol to upgrade the system. It will cause a WARM reboot and will require
that existing telnet, secure telnet or SSH sessions be restarted.
Do you want to continue? [y/n]: y
2010/01/29-23:48:35, [HAM-1004], 226, switchid 1, CHASSIS | VCS, INFO, Extreme_Elara2, Switch will be
rebooted with the new firmware.
```

Downloading firmware using the manual option

On a Top-of-Rack (ToR) device, this manual option allows you to specify the **noreboot** and **nocommit** options so that you have exact control over the firmware download sequence.

In a dual management-module (MM) system, the manual mode allows you to upgrade only the MM on which the **firmware download** command is issued. Furthermore, the manual option allows you to specify the **noreboot** or **nocommit** options. Therefore, you need to invoke the **firmware download** command with this option on both MMs.

NOTE

Network OS does not support the use of special characters (such as &!%#) in FTP and SCP passwords. If your SCP or FTP password contains special characters, the download fails.



CAUTION

Using the manual option causes disruption to the traffic. Do not use this option unless instructed to do so by Extreme Technical Support.

The following procedure applies to a ToR device or a single MM.

1. Enter the **firmware download interactive** command and respond to the prompts.

```
device# firmware download ftp host 10.20.1.3 user fvt password pray4green directory dist file
release.plist manual nocommit noreboot
Do you want to continue [y/n]: y
[output truncated]
```

2. After download completes, enter the **show version all-partitions** command to confirm that the primary partitions of the device contain the new firmware.
3. If you entered the `noreboot` option, enter the **reload** command to reboot the device. If you entered `y` after the prompt, the device will reboot automatically. The device performs a reboot and comes up with the new firmware. Your current CLI session will automatically disconnect.
4. Log back into the device. If you entered the `nocommit` option, enter the **firmware commit** command to commit the new firmware. If you entered `y` after the prompt, the device will commit the firmware automatically upon booting up.

```
device# firmware commit
Validating primary partition...
Doing firmwarecommit now.
Please wait ...
Replicating kernel image
.....
FirmwareCommit completes successfully.
```

5. Enter the **show version** command with the **all-partitions** option. Both partitions on the device or on the modules should contain the new firmware.

Verifying a firmware download session

After the firmware download completes, you can verify that the download has completed properly.

The following steps verify that the download completed correctly:

1. Execute the **show version all-partitions** command to verify that the MMs and all line-card partitions have the correct firmware.
2. Execute the **show ha all-partitions** command to verify that the MMs and all line-card partitions are in HA sync.
3. Execute the **show slots** command to verify that the MMs and all line cards are in the "enabled" state.
If the MMs are running different firmware, you need to execute the **firmware download** command with the **manual** option to update the standby MM to the same level as the active MM. If a line card is in the faulty state or the line-card partitions are not in sync, you must execute the **poweroff linecard** and **power-on linecard** commands to recover the line card.

Upgrading firmware in a logical-chassis

To upgrade a Logical Chassis fabric, you must log onto individual nodes in the fabric and issue the **firmware download** command. Alternatively, you can issuing the **firmware download logical-chassis** command on the principle node to upgrade multiple nodes at the same time.

NOTE

During the upgrade process there is a small chance that the invalid IP address 255.0.0.0/8 may display on the screen when you execute the **show run** command. This is normal and is not a cause for concern.

Overview of firmware download logical-chassis

In a logical-chassis device, you can run the **firmware download logical-chassis** command on the principle node to upgrade multiple nodes at the same time. The firmware is downloaded to the specified nodes simultaneously through their respective management ports. The number of nodes does not change the download time.

The following options are available with the **firmware download logical-chassis** command:

- Coldboot - download the new firmware to the nodes and reboot them automatically.

- Default-config - download the new firmware to the nodes, remove the configuration, and reboot them automatically.
- Auto-activate - download the new firmware to the nodes at the same time and activate the new firmware automatically. This option should be used for ISSU installations only.

If none of the above options are specified, the command defaults to downloading the firmware to the nodes. You must run **firmware activate rbridge-id <rid>** to activate the new firmware on the nodes. This scenario should be used for ISSU installations only.

Using ISSU with firmware download logical-chassis

When no option is specified, the **firmware download logical-chassis** command downloads the new firmware to all of the nodes at the same time. If the process fails on any of the nodes during the download stage, the command abort and the old firmware is restored on the nodes.

1. Run the **firmware download logical-chassis** command to download the firmware to the principal node.

```
device# firmware download logical-chassis protocol ftp host 10.10.10.10 user fvt password buzz
directory /dist/nos/6.0.1 file release.plist rbridge-id 1-3
Rbridge-id Sanity Result          Current Version
-----
1          Non-disruptive (ISSU)          6.0.1
2          Non-disruptive (ISSU)          6.0.1
3          Non-disruptive (ISSU)          6.0.1
```

2. Run the show **firmwaredownloadstatus summary rbridge-id <rid>** command to verify whether the nodes are ready for activation. It should show "Ready for activation" for the nodes.

```
device# show firmwaredownloadstatus summary bridge-id
Rid 1: INSTALLED (Ready for activation)
Rid 2: INSTALLED (Ready for activation)
Rid 3: INSTALLED (Ready for activation)
```

3. Run the **firmware activate** command to activate the firmware on the RBridge IDs.

```
device# firmware activate rbridge-id 1-2,3
This command will activate the firmware on the following nodes.
rbridge-id 1 : uses ISSU protocol, non-disruptive.
rbridge-id 2 : uses ISSU protocol, non-disruptive.
rbridge-id 3 : uses ISSU protocol, non-disruptive.

Do you want to continue? [y/n]:y
```

- Verify the firmware has been updated by running the **show version brief rbridge-id all** command.

```

device# show version brief rbridge-id all
rbridge-id 1
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1                          ACTIVE*
          NOS       6.0.1
SW/1     NOS       6.0.1                          STANDBY
          NOS       6.0.1
rbridge-id 2
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1                          ACTIVE*
          NOS       6.0.1
SW/1     NOS       6.0.1                          STANDBY
          NOS       6.0.1
rbridge-id 3
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1                          ACTIVE*
          NOS       6.0.1
SW/1     NOS       6.0.1                          STANDBY
          NOS       6.0.1

```

Using auto-active with firmware download logical-chassis

When the *auto-active* option is specified, the **firmware download logical-chassis** command downloads the new firmware to the nodes at the same time. After it completes downloading the firmware to all of the specified nodes, it causes the nodes to warm boot to initiate the firmware activation. If the process fails on any of the nodes during the download stage, the command abort and the old firmware is restored on the nodes. This option is for ISSU only.

- Run the **firmware download logical-chassis** command with the *auto-active* option to download the firmware to the principal node.

```

device# firmware download logical-chassis protocol ftp host 10.10.10.10 user fvt password buzz
directory /dist/nos/6.0.1 file release.plist rbridge-id 1-3 auto-activate
Rbridge-id Sanity Result      Current Version
-----
1          Non-disruptive (ISSU)      6.0.1
2          Non-disruptive (ISSU)      6.0.1
3          Non-disruptive (ISSU)      6.0.1

This command will download firmware to the specified nodes, and cause warm reboot on the nodes
automatically.
Do you want to continue? [y/n]:y

```


- Verify the firmware has been updated by running the **show version brief rbridge-id all** command.

```

device# show version brief rbridge-id all
rbridge-id 1
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0      NOS       6.0.1                          ACTIVE*
          NOS       6.0.1
SW/1      NOS       6.0.1                          STANDBY
          NOS       6.0.1
rbridge-id 2
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0      NOS       6.0.1                          ACTIVE*
          NOS       6.0.1
SW/1      NOS       6.0.1                          STANDBY
          NOS       6.0.1
rbridge-id 3
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0      NOS       6.0.1                          ACTIVE*
          NOS       6.0.1
SW/1      NOS       6.0.1                          STANDBY
          NOS       6.0.1

```

Using coldboot with firmware download logical-chassis

When the *coldboot* option is specified, the command **firmware download logical-chassis** downloads the new firmware to the nodes at the same time. After it completes downloading the firmware to all of the specified nodes, it causes the nodes to cold boot to initiate the firmware activation. If the process fails on any of the nodes during the download stage, the command abort and the old firmware is restored on the nodes.

NOTE

The *coldboot* option causes traffic disruption.

- Run the **firmware download logical-chassis** with the *coldboot* option command to download the firmware to the principal node.

```

device# firmware download logical-chassis protocol ftp host 10.10.10.10 user fvt password buzz
directory /dist/nos/6.0.1 file release.plist rbridge-id 1-3 coldboot
Rbridge-id Sanity Result Current Version
-----
1          Disruptive          6.0.1
2          Disruptive          6.0.1
3          Disruptive          6.0.1
You are invoking firmware download with the coldboot option. This command will download the new
firmware to the specified nodes, and cause cold reboot.
Do you want to continue? [y/n]: y

```

- Verify the firmware has been updated by running the **show version brief rbridge-id all** command.

```
device# show version brief rbridge-id all
rbridge-id 1
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1
          6.0.1
          ACTIVE*
SW/1     NOS       6.0.1
          6.0.1
          STANDBY

rbridge-id 2
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1
          6.0.1
          ACTIVE*
SW/1     NOS       6.0.1
          6.0.1
          STANDBY

rbridge-id 3
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1
          6.0.1
          ACTIVE*
SW/1     NOS       6.0.1
          6.0.1
          STANDBY
```

Using default-config with firmware download logical-chassis

When the *default-config* option is specified, the **firmware download logical-chassis** command downloads the new firmware to the nodes at the same time. After it completes downloading the firmware to all of the specified nodes, it removes the current configuration and causes the nodes to cold boot to initiate the firmware activation. If the process fails on any of the nodes during the download stage, the command abort and the old firmware is restored on the nodes.

NOTE

The *default-config* option causes traffic disruption and configuration loss to the nodes.

- Run the **firmware download logical-chassis** with the *default-config* option command to download the firmware to the principal node.

```
device# firmware download logical-chassis default-config protocol ftp host 10.10.10.10 user fvt
password buzz directory /dist/nos/6.0.1 file release.plist rbridge-id 1-3
Rbridge-id Sanity Result Current Version
-----
1           Disruptive          6.0.1
2           Disruptive          6.0.1
3           Disruptive          6.0.1
```

You are invoking firmware download with the default-config option. This command will download the new firmware to the specified nodes and default their configuration.
Do you want to continue? [y/n]: y

- Verify the firmware has been updated by running the **show version brief rbridge-id all** command.

```

device# show version brief rbridge-id all
rbridge-id 1
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1                          ACTIVE*
          6.0.1
SW/1     NOS       6.0.1                          STANDBY
          6.0.1

rbridge-id 2
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1                          ACTIVE*
          6.0.1
SW/1     NOS       6.0.1                          STANDBY
          6.0.1

rbridge-id 3
Slot      Name      Primary/Secondary Versions      Status
-----
SW/0     NOS       6.0.1                          ACTIVE*
          6.0.1
SW/1     NOS       6.0.1                          STANDBY
          6.0.1

```

Updating the peripheral firmware

Updates the peripheral firmware on the VDX 6740T/VDX 6740T devices.

The VDX 6740T/VDX 6740T devices have twelve AQ1402 PHY chips accessible through the MDIO bus from the CPU Complex. Each of these chips have a on-board IRAM that needs to be loaded with a firmware Image from Aquantia using an external flash chip connected to one of the AQ1402 Chips (which is the master chip). The rest of the PHY Chips are daisy-chained and are gang-loaded with the firmware through the previous chip's daisy chain I/F. The gang-loading happens when the chips are released out of reset. The AQ1402 firmware source is a binary file which is a few hundred KB's in size.

- Execute the **firmware peripheral-update** command to update the peripheral devices.

```
device# firmware peripheral-update microcode phy usb Firmware_1.38.c1_Extreme.Castor.cld
```

- Execute the **show version peripheral phy** command to confirm the update.

```

device# show version peripheral phy
1.38.c1
device#

```

Example of typical execution using a flash source file.

```

device# firmware peripheral-update microcode phy flash Firmware_1.38.c1_Extreme.Castor.cld
Starting to Program AQ Flash....
AQ1402[0xbb300000] phy handle
CRC check good on image file
Taking Control Of the Flash Interface..
Determining Flash Type...
bootLoad=0x2
primary=0x2
FLASH type = Atmel AT45DB041D
Erasing Flash.....
Atmel AT45DB041D Erase Started
Writing Image to Flash .....
Bytes in file 0x2E000

Starting read of FLASH data
OK
AQ Flash Upgrade Successful

```

Updating the peripheral firmware

```
2016/07/19-05:53:24, [SULB-1000], 13875, SW/0 | Active, WARNING, VDX6740T-1G, The firmware download command has been started.
```

```
AQ Firmware_1.38.c1_Extreme.Castor.cld upgrade success...PowerCycle/Reboot Switch.
```

```
device# reload
```

```
Are you sure you want to reload the switch? [y/n]:y
```

```
The system is going down for reload NOW !!
```

```
When upgrading the AQ1402 PHY from flash, it is expected that the firmware has been copied to the flash:// folder.
```

Advanced Upgrade Scenarios

- [Upgrading firmware within a VCS Fabric.....](#)29

Upgrading firmware within a VCS Fabric

Use this procedure, illustrated by an example topology, to upgrade or downgrade firmware within a VCS Fabric.

It is important to reduce the downtime incurred by planned software upgrades. This section describes how to upgrade and downgrade Extreme Network OS firmware images efficiently and safely onto a variety of platforms in a VCS Fabric.

ATTENTION

This procedure illustrates an upgrade for minor releases only, such as Network OS 5.0.0 to Network OS 5.1.0. For the limitations and caveats related to a specific Network OS release, refer to the release notes for that release.

Although it is necessary to reboot the devices after the installation, the following benefits are achieved:

- An optimal upgrade cycle for the entire VCS fabric
- Minimal loss of traffic
- No loss of configuration status

This procedure is supported on the following Extreme platforms:


- VDX 2741
- VDX 2746
- VDX 6740 and VDX 6740T
- VDX 8770-4 and VDX 8770-8

The example approach presented here, tested in a Extreme lab topology, is intended as a best-practices model that is to be modified for existing customer deployments. Software release versions will vary.

Tested topology

The tested topology, illustrated in the following figure, is a four-node Extreme VDX fabric on VCS 8192. The fabric consists of two spine nodes and two leaf nodes, connected to the core through a vLAG. Two servers, running Extreme CNA, are dual-homed through two vLAGs to both the leaf nodes of VCS 8192.

FIGURE 1 Tested topology



The following table summarizes the tested components.

TABLE 2 Tested components and roles

Position	VCS name	Chassis type	Description
Leaf	8192	VDX 6740-48	Dual-homed TOR VDX
Spine	8192	VDX 6740-48	Dual-homed
Core	NA	VDX 8770	Connected to spine nodes of VCS 8192 through 16-port vLAG
Servers	NA	ESX 5.1 with Extreme CNA	Dual-homed to both leaf nodes through a vLAG

Upgrading nodes by using an odd/even approach

To reduce downtimes during planned software upgrades, most networks have been provisioned with redundancy in all layers. Once such in-built redundancy is in place, an "odd/even" approach is used, whereby the fabric is split equally into odd and even nodes that represent both sides of the redundant traffic path. Therefore, both groups (either all odd or all even) have traffic connectivity to all hosts and end devices during the upgrade process. As a result, reloading any one group results in minimal traffic loss.

The following table summarizes the classification of the odd and even nodes that were tested.

TABLE 3 Classification of odd and even nodes

Position	Chassis type	Description	Odd group	Even group
Leaf	VDX 6740-48	Dual-homed TOR VDX	sw81	sw79
Spine	VDX 6740-48	Dual-homed	sw87	sw86
Servers	ESX 5.1 with Extreme CNA	Dual-homed to both leaf nodes through a vLAG		

Preparing for the maintenance window

Do the following before the start of the maintenance window.

1. Take a "golden" snapshot of the running configuration, by copying the running configuration onto flash or an external FTP or SCP server. The following command copies the running configuration to flash memory.


```
device# copy running-config flash://running.Config.master
```
2. Establish Telnet connections and console connections to all VDX Fabric nodes. Telnet sessions are used to perform configurations, while console sessions are used to monitor the devices.
3. View the running configuration (as illustrated in the following example) to ensure that all the port-channel interfaces have been configured by means of the **vlag ignore-split** command on all VDX nodes.

NOTE

By default, port-channel configurations have the **vlag ignore-split** command enabled. However, if this default has been changed, it must be re-established.

```
device# configure terminal
Entering configuration mode terminal
device(config)# int po 6144
device(config-Port-channel-6144)# vlag ignore-split
device(config-Port-channel-6144)# exit
device(config)# exit
device#

device# show running-config interface Port-channel 6144
interface Port-channel 6144
 vlag ignore-split
 switchport
 switchport mode trunk
 switchport trunk allowed vlan all
 switchport trunk tag native-vlan
 no shutdown
!
```

4. Check the state of the system by using the following **show** commands.

- a) Verify that all the nodes to be upgraded are running the same version, by using the **show version** command.

```
device# show version

Network Operating System Software
Network Operating System Version: 3.0.1
Copyright (c) 1995-2012 Brocade Communications Systems, Inc.
Firmware name:      3.0.1c
Build Time:         23:15:48 Oct 18, 2013
Install Time:       04:27:13 Dec  3, 2013
Kernel:             2.6.34.6

BootProm:           2.2.0
Control Processor:  e500v2 with 2048 MB of memory

Appl   Primary/Secondary Versions
-----
NOS    3.0.1c
       3.0.1c

device#
```

- b) Verify the state of the fabric, by using the **show fabric all** command.

```
device# show fabric all

VCS Id: 8192
Config Mode: Local-Only

Rbridge-id      WWN                IP Address          Name
-----
79              10:00:00:27:F8:44:50:C2  10.20.53.79        "sw79"
81              10:00:00:05:33:FA:44:08  10.20.53.81        "sw81"
86              10:00:00:27:F8:0C:D1:BD  10.20.53.86        >"sw86"
87              10:00:00:05:33:FA:4A:48  10.20.53.87        "switch"*

The Fabric has 4 Rbridge(s)

device#
```

- c) Verify that all fabric ISLs are up, by using the **show fabric isl** command.

```
device# show fabric isl

Rbridge-id: 87  #ISLs: 3

  Src      Src      Nbr      Nbr
Index  Interface  Index  Interface      Nbr-WWN          BW  Trunk  Nbr-Name
-----
0      Te 87/0/1  14      Te 79/0/15     10:00:00:27:F8:44:50:C2  10G  Yes   "sw79"
3      Te 87/0/4  7       Te 86/0/8      10:00:00:27:F8:0C:D1:BD  10G  Yes   "sw86"
4      Te 87/0/5  3       Te 81/0/4      10:00:00:05:33:FA:44:08  10G  Yes   "sw81"

device#
```

- d) Verify the state of the port-channels, by using the **show port-channel summary** command.

```
device# show port-channel summary
LACP Aggregator: Po 6144 (vLAG)
Aggregator type: Standard
Ignore-split is enabled
Member rbridges:
  rbridge-id: 86 (16)
  rbridge-id: 87 (16)
Admin Key: 6144 - Oper Key 6144
Member ports on rbridge-id 87:
  Link: Te 87/0/9 (0x5718050008) sync: 1
```



```

Link: Te 87/0/10 (0x5718050009) sync: 1
Link: Te 87/0/11 (0x571805800A) sync: 1
Link: Te 87/0/12 (0x571806000B) sync: 1
Link: Te 87/0/13 (0x571806800C) sync: 1
Link: Te 87/0/14 (0x571807000D) sync: 1
Link: Te 87/0/15 (0x571807800E) sync: 1
Link: Te 87/0/16 (0x571808000F) sync: 1
Link: Te 87/0/17 (0x5718088010) sync: 1
Link: Te 87/0/18 (0x5718090011) sync: 1
Link: Te 87/0/19 (0x5718098012) sync: 1
Link: Te 87/0/20 (0x57180A0013) sync: 1
Link: Te 87/0/21 (0x57180A8014) sync: 1
Link: Te 87/0/22 (0x57180B0015) sync: 1
Link: Te 87/0/23 (0x57180B8016) sync: 1
Link: Te 87/0/24 (0x57180C0017) sync: 1

```

```
device#
```

- e) Verify that the required ports and port-channels are up, by using the **show ip interface brief** command.

```
device# show ip interface brief
```

```

Interface                               IP-Address Status                Protocol
=====                               =
Port-channel 6144                        unassigned up                        up
TenGigabitEthernet 87/0/1                unassigned up                        up (ISL)
TenGigabitEthernet 87/0/2                unassigned up                        down
TenGigabitEthernet 87/0/3                unassigned administratively down    down
TenGigabitEthernet 87/0/4                unassigned up                        up (ISL)
TenGigabitEthernet 87/0/5                unassigned up                        up (ISL)
TenGigabitEthernet 87/0/6                unassigned up                        down
TenGigabitEthernet 87/0/7                unassigned up                        down
TenGigabitEthernet 87/0/8                unassigned up                        up
TenGigabitEthernet 87/0/9                unassigned administratively down    down
TenGigabitEthernet 87/0/10              unassigned up                        up
<output truncated>

```

- f) Verify the number of MAC addresses in the fabric and other details, by using the **show mac-address-table count** and **show mac-address-table** commands.

```

device# show mac-address-table count
Dynamic Address Count : 11
Static Address Count  : 0
Internal Address Count : 3
Total MAC addresses   : 14
device#
device# show mac-address-table
VlanId  Mac-address      Type      State      Ports
99      0005.3378.442a   Dynamic  Active    Po 6144
99      0005.3378.5242   Dynamic  Active    Po 6144
99      0005.33fa.4429   System   Remote    XX 81/X/X
99      0027.f80c.d1de   System   Remote    XX 86/X/X
99      0027.f844.50e3   System   Remote    XX 79/X/X
208     0009.8a06.6cbd   Dynamic  Active    Po 6144
208     0050.5656.3d02   Dynamic  Remote    Po 100
208     0050.5656.3d03   Dynamic  Remote    Po 100
208     0050.5656.3f42   Dynamic  Remote    Po 400
208     0050.5656.3f43   Dynamic  Remote    Po 400
208     0050.5661.2b00   Dynamic  Remote    Po 300
208     0050.566c.c929   Dynamic  Remote    Po 200
208     0050.56b3.18e3   Dynamic  Remote    Po 300
208     0050.56b3.2801   Dynamic  Remote    Po 400
Total MAC addresses   : 14
device#

```

- g) Check the traffic rate on all the ports and port-channels along the traffic path, by using the **show interface** command with the following output option.

```
device# show interface port-channel 6144 | inc rate
Queueing strategy: fifo
  Input 86.487040 Mbits/sec, 84460 packets/sec, 8.65% of line-rate
  Output 86.487040 Mbits/sec, 84460 packets/sec, 8.65% of line-rate
```

5. Identify the principal and multicast root nodes of the fabric.

- a) Identify the principal node (RBridge), by using the **show fabric all** command.

```
device# show fabric all

VCS Id: 8192
Config Mode: Local-Only
```

Rbridge-id	WWN	IP Address	Name
79	10:00:00:27:F8:44:50:C2	10.20.53.79	"sw79"
81	10:00:00:05:33:FA:44:08	10.20.53.81	"sw81"
86	10:00:00:27:F8:0C:D1:BD	10.20.53.86	>"sw86"
87	10:00:00:05:33:FA:4A:48	10.20.53.87	"switch"*

```
The Fabric has 4 Rbridge(s)

device#
```

- b) Identify the multicast root node (RBridge), by using the **show fabric route multicast** command.

```
device# show fabric route multicast

Root of the Multicast-Tree
=====
Rbridge-id: 79
Mcast Priority: 1
Enet IP Addr: 10.20.53.79
WWN: 10:00:00:27:f8:44:50:c2
Name: sw79
Rbridge-id: 87
```

Src-Index	Src-Port	Nbr-Index	Nbr-Port	BW	Trunk
0	Te 87/0/1	14	Te 79/0/15	10G	Yes

```
device#
```

6. Identify "odd" and "even" nodes in the fabric, as defined previously in [Upgrading nodes by using an odd/even approach](#) on page 31.

Once the fabric is dual-homed and redundancy in all layers is established, split the fabric into "odd" and "even" nodes so that all nodes, either all odd or all even, have traffic connectivity to all hosts and end devices, resulting in minimal traffic loss.

7. Terminate any Fibre Channel sessions that traverse the fabric.

NOTE

For fabrics that do not support HA, FC and FCoE logins are affected during reloads.

8. Check memory utilization by using the **show process memory** command, and ensure that the 70 percent threshold is not exceeded.

Optimizing reconvergence in the VCS Fabric

While the VCS Fabric is reconverging after odd/even groups are reloaded or coming back into the fabric, there may be momentary spikes in traffic that can result in traffic loss. Before upgrading or reloading VDX nodes, it is crucial to ensure that flow control is enabled on the following interfaces to minimize the impact of reconvergence:

- Access ports that face servers or hosts. These can be port-channel or physical interfaces, depending upon the host or server configuration.
- Uplink interfaces that connect to the core (port-channel 6144 in the example topology).
- Interfaces supporting the VCS Fabric topology on all core and end devices.

NOTE

ISL interfaces have flow control enabled by default.

Do the following to optimize reconvergence.

1. Confirm whether flow control is enabled, by using the **show running-config interface port-channel 6144** command on the previously listed interfaces, as in the following example.

```
switch# show running-config interface Port-channel 6144
interface Port-channel 6144
 vlag ignore-split
 switchport
 switchport mode trunk
 switchport trunk allowed vlan all
 switchport trunk tag native-vlan
 qos flowcontrol tx on rx on
 no shutdown
!
```

2. To enable flow control on an interface that does not have it enabled, use the **qos flowcontrol tx rx** command.
3. Where servers are connected to leaf nodes, it is recommended that Link Aggregation Control Protocol (LACP) vLAGs be used to minimize traffic loss.

Maintaining the VCS Fabric

During a VCS Fabric maintenance window, it is recommended that you do the following.

1. Download the required firmware on all VDX fabric nodes by using the **nocommit**, **noreboot**, and **coldboot** options as appropriate. The last option applies to non-ISSU firmware downloads.



CAUTION

Do not use the **coldboot** option unless directed to do so by Extreme Technical Support. If you do not select the **nocommit** option, firmware is downloaded automatically. This prevents you from backing out of an upgrade should that become necessary.

NOTE

In this test topology, because the upgrade is from 4.1.1 to 5.0.0, 5.0.0 is loaded on all VDX fabric nodes. Your versions will vary accordingly.

2. Verify that there is no control or data traffic outage during the firmware download process.

3. Reboot the "odd" nodes. Wait at least five minutes for the devices (in the example sw81 and sw87) to come back up.

NOTE

The time required for the devices to come back up with the configuration replay complete depends on the number of configuration lines that must be read. Refer to [Understanding traffic outages](#) on page 37 for example traffic-outage times.

4. Wait at least ten minutes for the fabric to converge and all back-end processes to be completed.

NOTE

At this point in the process, sw79 and sw86 are at 4.1.1, while sw87 and sw81 are at 5.0.0.

5. Verify that all four nodes are part of the same fabric, by using the **show fabric all** command.
6. At this point, reboot all the "even" nodes, including the principal and the multicast root node.

NOTE

After the firmware completes downloading on all nodes, in large-scale deployments you must reload the fabric "even" nodes, which includes the principal and multicast root nodes. It is imperative that the principal device be in the last batch of nodes to be activated with the new firmware, or they will be locked out of activating the other nodes in a logical chassis configuration.

NOTE

Because the fabric principal and multicast root nodes have already been identified previously as "even" nodes, Network OS reloads the "odd" nodes first, and then the "even" nodes. In our example, the "odd" nodes sw87 and sw81 are reloaded at the same time. Refer to [Understanding traffic outages](#) on page 37 for details of the traffic outages that occurred at different phases of the process in the tested topology.

7. Wait at least ten minutes for the fabric to converge and all back-end processes to be completed.

NOTE

At this point, all nodes are at 5.0.0.

8. Verify that all nodes have joined the fabric, by using the **show fabric all** and **show vcs** commands.

ATTENTION

Ensure that there are no traffic outages.

9. Verify that system health is as discussed in step 4 of [Preparing for the maintenance window](#) on page 31.
10. Evaluate the state of the upgrade process at this point. The upgraded firmware should have been downloaded onto the primary partition, while the original firmware should be present in the second partition.
11. To complete the process, commit the firmware, by using the **firmware commit** command.

If an unexpected development occurs, you can roll back to the original firmware.

Understanding traffic outages

For example traffic-outage times as measured by various traffic-analysis tools, note the following.

1. Note the following traffic-outage times that occurred when the "odd" devices, sw87 and sw81, were reloaded.

TABLE 4 Traffic-outage times: "Odd" devices, upgrading from 6.0.1 to 7.0.0

Tool	Traffic path 2	Traffic path 1
Layer 2 traffic	0 ms (within same rack)	~118 ms (ping from server to sw137)
Layer 3 traffic-generator traffic	N/A	0 ms

TABLE 5 Traffic-outage times: "Odd" devices, reloading within 6.0.1

Tool	Traffic path 2	Traffic path 1
Layer 2 traffic	0 ms (within same rack)	~122 ms (ping from server to sw137)
Layer 3 traffic-generator traffic	N/A	0 ms

2. Note the following traffic-outage times that occurred when the "even" devices, sw79 and sw86, were reloaded.

TABLE 6 Traffic-outage times: "Even" devices, upgrading from 6.0.1 to 7.0.0

Tool	Traffic path 2	Traffic path 1
Layer 2 traffic	0 ms (within same rack)	~129 (ping from server to sw137)
Layer 3 traffic-generator traffic	N/A	0 ms

TABLE 7 Traffic-outage times: "Even" devices, reloading within 6.0.1

Tool	Traffic path 2	Traffic path 1
Layer 2 traffic	0 ms (within same rack)	~123 (ping from server to sw137)
Layer 3 traffic-generator traffic	N/A	0 ms