

Network OS Software Upgrade Guide

Supporting Network OS v6.0.1

© 2016, Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, Brocade Assurance, the B-wing symbol, ClearLink, DCX, Fabric OS, HyperEdge, ICX, MLX, MyBrocade, OpenScript, VCS, VDX, Vplane, and Vyatta are registered trademarks, and Fabric Vision is a trademark of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned may be trademarks of others.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. assume no liability or responsibility to any person or entity with respect to the accuracy of this document or any loss, cost, liability, or damages arising from the information contained herein or the computer programs that accompany it.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Contents

Preface	5
Document conventions.....	5
Text formatting conventions.....	5
Command syntax conventions.....	5
Notes, cautions, and warnings.....	6
Brocade resources.....	6
Contacting Brocade Technical Support.....	6
Brocade customers.....	6
Brocade OEM customers.....	7
Document feedback.....	7
About this document	9
Supported hardware and software.....	9
Using the Network OS CLI	10
What's new in this document.....	10
Installing and Maintaining Firmware	11
Firmware management overview.....	11
Upgrading firmware on a ToR switch.....	12
Upgrading firmware on a modular chassis.....	12
Automatic firmware synchronization.....	12
ISSU features and limitations.....	12
Supported platforms.....	12
Downgrading considerations and restrictions.....	13
Preparing for a Firmware Download	15
Prerequisites.....	15
Obtaining and decompressing firmware.....	15
Basic Firmware Upgrade	17
Upgrading firmware on a local switch.....	17
Upgrading considerations and restrictions.....	17
Connecting to the switch.....	17
Obtaining the firmware version.....	18
Using the firmware download command.....	19
Downloading firmware using ISSU.....	19
Downloading firmware using the noactivate option.....	20
Downloading firmware using the coldboot option.....	21
Downloading firmware using the default-config option.....	21
Downloading firmware from a USB device.....	21
Downloading firmware using the manual option.....	22
Verifying a firmware download session.....	23
Upgrading firmware in a logical-chassis.....	23
Overview of firmware download logical-chassis.....	23
Using ISSU with firmware download logical-chassis.....	24
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

Mixed-node fabric cluster support.....	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.....	36
Mixed-node fabric cluster support.....	37

Preface

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

Document conventions

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

Text formatting conventions

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

Format	Description
bold text	Identifies command names Identifies keywords and operands Identifies the names of user-manipulated GUI elements
<i>italic text</i>	Identifies text to enter at the GUI Identifies emphasis Identifies variables
Courier font	Identifies document titles Identifies CLI output 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.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, --show WWN.
[]	Syntax components displayed within square brackets are optional.
{ x y z }	Default responses to system prompts are enclosed in square brackets. A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	In Fibre Channel products, square brackets may be used instead for this purpose. A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.

Convention	Description
...	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.

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.

Brocade resources

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

You can download additional publications supporting your product at www.brocade.com. Select the Brocade Products tab to locate your product, then click the Brocade product name or image to open the individual product page. The user manuals are available in the resources module at the bottom of the page under the Documentation category.

To get up-to-the-minute information on Brocade products and resources, go to MyBrocade. You can register at no cost to obtain a user ID and password.

Release notes are available on MyBrocade under Product Downloads.

White papers, online demonstrations, and data sheets are available through the Brocade website.

Contacting Brocade Technical Support

As a Brocade customer, you can contact Brocade Technical Support 24x7 online, by telephone, or by e-mail. Brocade OEM customers contact their OEM/Solutions provider.

Brocade customers

For product support information and the latest information on contacting the Technical Assistance Center, go to <http://www.brocade.com/services-support/index.html>.

If you have purchased Brocade product support directly from Brocade, use one of the following methods to contact the Brocade Technical Assistance Center 24x7.

Online	Telephone	E-mail
<p>Preferred method of contact for non-urgent issues:</p> <ul style="list-style-type: none"> • My Cases through MyBrocade • Software downloads and licensing tools • Knowledge Base 	<p>Required for Sev 1-Critical and Sev 2-High issues:</p> <ul style="list-style-type: none"> • Continental US: 1-800-752-8061 • Europe, Middle East, Africa, and Asia Pacific: +800-AT FIBREE (+800 28 34 27 33) • For areas unable to access toll free number: +1-408-333-6061 • Toll-free numbers are available in many countries. 	<p>support@brocade.com</p> <p>Please include:</p> <ul style="list-style-type: none"> • Problem summary • Serial number • Installation details • Environment description

Brocade OEM customers

If you have purchased Brocade product support from a Brocade OEM/Solution Provider, contact your OEM/Solution Provider for all of your product support needs.

- OEM/Solution Providers are trained and certified by Brocade to support Brocade® products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/Solution Provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/Solution Provider.

Document feedback

To send feedback and report errors in the documentation you can use the feedback form posted with the document or you can e-mail the documentation team.

Quality is our first concern at Brocade 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:

- Through the online feedback form in the HTML documents posted on www.brocade.com.
- By sending your feedback to documentation@brocade.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.

About this document

- Supported hardware and software.....9
- Using the Network OS CLI10
- What's new in this document.....10

Supported hardware and software

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

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for Network OS 6.0.1, 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:

- Brocade VDX 2740

NOTE

The Brocade VDX 2740 is the equivalent of the Lenovo Flex System EN4023 10Gb Scalable Switch. This platform is identified in the system as EN4023.

- Brocade VDX 2746
- Brocade VDX 6740
 - Brocade VDX 6740-48
 - Brocade VDX 6740-64
- Brocade VDX 6740T
 - Brocade VDX 6740T-48
 - Brocade VDX 6740T-64
 - Brocade VDX 6740T-1G
- Brocade VDX 6940-36Q
- Brocade VDX 6940-144S
- Brocade VDX 8770
 - Brocade VDX 8770-4
 - Brocade 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 Network OS command line interface (CLI), refer to the *Network OS Command Reference*.

What's new in this document

The *Network OS Upgrade Guide* is a new publication. It replaces the "Installing and Maintaining Firmware" chapter in the *Network OS Administration Guide*.

The content has been updated with the following changes:

- Updated to support upgrading from previous versions to Network OS v6.0.1.
- Updated the VCS Fabric upgrade process.

Installing and Maintaining Firmware

- [Firmware management overview](#).....11
- [Upgrading firmware on a ToR switch](#).....12
- [Upgrading firmware on a modular chassis](#).....12
- [ISSU features and limitations](#).....12
- [Supported platforms](#).....12
- [Downgrading considerations and restrictions](#).....13

Firmware management overview

Brocade 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 forth) 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) switch 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 Brocade-branded USB device. If you want to download firmware from a remote server, you must connect the management Ethernet port of the switch 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 switches to be upgraded non-disruptively and is invoked by entering the **firmware download** command from the active management module. Refer to [Downloading firmware using ISSU](#) on page 19.

In Network OS v4.0.0 and later, the **logical-chassis firmware download** command allows you to upgrade a single switch or multiple switches of your choice that are connected in logical chassis cluster 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. Refer to [Downloading firmware using the default-config option](#) on page 21.

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

In Network OS v5.0.0 and later, in-service software upgrade (ISSU) is supported on a Top of Rack switch (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 19.

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.

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 the feature. As a result, the feature cannot be disabled.

ISSU features and limitations

ISSU is used automatically by the system for firmware upgrades when hardware and version requirements are met. ISSU is an advanced firmware upgrade method that is designed to be hitless. ISSU includes the following characteristics and limitations:

- No data path disruption occurs for Layer 2, Layer 3, and FCoE traffic.
- All Layer 2 control protocol states are retained.
- Topology state and interface state are retained.
- Last accepted user configuration is retained.
- For Release 4.0.x and 4.1.x, data path and control path disruption occurs for IP. The IP configuration is replayed after the failover.
- For Release 5 and later no Layer 3 (IP) disruption will occur when upgrading to a later patch or maintenance release.

Supported platforms

In most cases, you will be upgrading firmware by installing a more recent firmware version than the one you are currently running. However, some circumstances may require that you downgrade the firmware to an earlier version. The procedures described in the following section assume that you are upgrading firmware, but they work for downgrading as well, provided that the firmware version you are downgrading to is compatible with the version you are currently running. The following lists supported firmware versions by platform.

TABLE 1 Network OS firmware support by platform

Platform	4.0.0	4.1.0	5.0.0	5.0.1	6.0.0	6.0.1
Brocade VDX 2740	no	no	yes	yes	yes	yes
Brocade VDX 6710	yes	yes	no	no	no	no
Brocade VDX 6720	yes	yes	no	no	no	no
Brocade VDX 6730	yes	yes	no	no	no	no
Brocade VDX 6740	yes	yes	yes	yes	yes	yes
Brocade VDX 6740T						
Brocade VDX 6740T-1G	no	yes	yes	yes	yes	yes
Brocade VDX 6940	no	no	no	no	yes	yes
Brocade VDX 8770	yes	yes	yes	yes	yes	yes

Downgrading considerations and restrictions

Consider the following when downgrading your firmware version:

- If a feature is new for the current version of your firmware, it will not function if you downgrade your firmware version.
- If you want to downgrade the firmware version in logical chassis mode, use the *coldboot* option only. This option is disruptive.
- Firmware downgrades from Network OS v6.0.0 to previous versions are prohibited when security parameters are configured for HTTPS support.
- SNMP trap functionality is not available in pre-Network OS v6.0.0 releases.
- If any SNMP v3 users are associated with the "admin"/"user" group, after downgrading to a version earlier than Network OS v5.x.x, the same user is automatically associated with the "snmpadmin"/"snmpuser" group for backward compatibility.
- Firmware downgrade from Network OS v6.0.0 will be blocked if switched or routed ACLs are configured. You must manually remove the ACL configurations in order for the firmware downgrade to take place.
- Firmware downgrade from Network OS v6.0.1 will be blocked if Monitoring and Policy Alert Suite (MAPS) is configured. You must manually remove the MAPS configuration in order for the firmware downgrade to take place.
- On the Brocade VDX 6740 and Brocade VDX 6940 hardware series, downgrading the firmware when the HA is not in sync and the SW1 partition is active may lead to loss of configuration. To avoid configuration loss, the HA must be synchronized before downgrading the firmware.
- Hardware profile support for TCAM and Routing Table is supported when downgrading from Network OS v6.x to Network OS v5.x.
- 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.
- Prior to downgrading to Network OS v6.0.2x, the IGMP snooping static mrouter feature must be deactivated from all snoop or non-snoop enabled VLANs.

The VCS fabric is segmented after ISSU downgrade from Network OS v5.0.1b to Network OS v5.0.1a with the *noactivate* option. Please follow these steps while downgrading from NOS 5.0.1b to 5.0.1a:

1. Trigger the **firmware download int** command, either with or without *firmware activate* option. It will install the firmware on secondary partitions of each GOS without forming the cluster.
2. Use the **firmware activate rbridge-id xx** command to activate the firmware on each node individually. It forms the cluster with the new version automatically.

Or you may:

1. Use the "firmware download coldboot" command instead of the "firmware download int" command.

If the secondary node in a cluster has its firmware downgraded to a lower version, then all commands are blocked in the cluster and the secondary node will not be able to join the cluster until the secondary node is rebooted.

If the primary node in a cluster has its firmware downgraded to a lower version, then the cluster is not usable (no commands and no new cluster formations are allowed) until the primary node is rebooted.

Do not downgrade the firmware on the Brocade VDX 6740 and Brocade VDX 6940 series when HA is not in sync and the SW1 partition is active. This may lead to loss of configurations. To avoid this situation, HA should be synchronized before downgrading

Downgrading the firmware on a switch to a Network OS version earlier than 4.0 is not allowed when either the Telnet server or the SSH server on the switch is disabled. To downgrade to a lower version, both the Telnet Server and SSH Server must be enabled.

When downgrading to lower firmware version which does not support MLD Snooping, before downgrading, you must disable the feature and then proceed with the operation.

User names with a leading underscore block a firmware downgrade. You must modify the user name to remove the leading underscore.

For the following features, learned routes are not preserved in control and forwarding planes when downgrading. All routes are lost and traffic disruption occurs.

- PIM Routes
- Layer 3 MCACHE
- PIM VIF
- Static RP
- BSR learned RP

Always refer to the release notes for compatibility information and take note of restrictions that may exist regarding upgrades and downgrades under particular circumstances.

NOTE

Refer to the Network OS documentation for details on the features listed.

Preparing for a Firmware Download

- Prerequisites.....15
- Obtaining and decompressing firmware.....15

Prerequisites

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

1. Verify the current firmware version. Refer to [Obtaining the firmware version](#) on page 18 for details.
2. Download the firmware package from the Brocade website to an FTP server.
3. Decompress the firmware archive. Refer to [Obtaining and decompressing firmware](#) on page 15.
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 switch configuration prior to the firmware download. Refer to [Installing and Maintaining Firmware](#) on page 11 for details.
7. For additional support, connect the switch 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 is issued and the files collected, the **clear support** command can be issued 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 Brocade website at www.mybrocade.com.

You must download the firmware package either to an FTP server or to a USB device and decompress the package *before* you can use the **firmware download** command or **firmware download logical-chassis** command (if you are in VCS mode) to upgrade the firmware on your equipment. Use the UNIX **tar** command for .tar files, the **gunzip** command for all .gz files, or a Windows unzip program for all .zip files.

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 or **firmware download logical-chassis** command (if you are in logical chassis cluster mode) performs an automatic search for the correct package file type associated with the device.

Basic Firmware Upgrade

- Upgrading firmware on a local switch.....17
- Upgrading considerations and restrictions.....17
- Connecting to the switch.....17
- Obtaining the firmware version.....18
- Using the firmware download command.....19
- Upgrading firmware in a logical-chassis.....23
- Mixed-node fabric cluster support.....27

Upgrading firmware on a local switch

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

Upgrading considerations and restrictions

Consider the following when upgrading your firmware version:

- 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 switch 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, Brocade 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.

Connecting to the switch

When you upgrade firmware in default mode, you connect to the switch through the management IP address. Modular switches 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.

```
switch# 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.

```
switch# 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 mgmt-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.

```
switch# 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 switches (since Release 5.0), 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 - Simplifies upgrading dual-MM systems between releases that do not support ISSU. Coldboot is the standard option for both dual-MM system and TOR switches for all firmware upgrades and downgrades between major releases when ISSU is not supported.
- Default-config - Removes all configuration and is similar to an initial installation and configuration. This is the standard option for most of upgrades and downgrades between two major releases when ISSU and ColdBoot are not supported. User may choose to perform upgrade/downgrade with this option between N<-> (N +/- 1) releases as well if there is no requirement to preserve the configuration.
- Manual - This option is not recommended in releases since Network OS v4.1. 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.

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

- [Downloading firmware using the noactivate option](#) on page 20
- [Downloading firmware using the manual option](#) on page 22
- [Downloading firmware using the coldboot option](#) on page 21
- [Downloading firmware using the default-config option](#) on page 21

To download firmware from an attached USB device, refer to [Downloading firmware from a USB device](#) on page 21.

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

1. Perform the steps described in [Prerequisites](#) on page 15.
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 switch or management module you are upgrading.

Refer to [Connecting to the switch](#) on page 17 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 noactivate option

The **noactivate** option in the **firmware download** command allows you to download new firmware onto a switch 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.

```
switch# 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 switch to reboot the switch 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.

```
switch# 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, Brocade_Elara2, Switch will be
rebooted with the new firmware.
```

Downloading firmware using the coldboot option

The **coldboot** option in the **firmware download** command allows you to download new firmware onto a switch and forces the switch 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.

NOTE

This option causes traffic disruption.

1. Download the firmware from the source directory with the coldboot option.

```
switch# firmware download scp host 10.70.4.109 user fvt directory /buildsjc/sre/SQA/nos/nos5.0.0/
nos5.0.0_bld30 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 switch completes the reboot sequence, you may log in to the switch and operate normally.

Downloading firmware using the default-config option

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

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



CAUTION

When you invoke **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 switch and optionally change the VCS mode, the VCS ID, and the RBridge ID before rebooting the switch 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:

```
switch# 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

The Brocade switches support firmware download from a Brocade-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 switch.

2. Enter the **usb on** command in privileged EXEC mode.

```
switch# 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.

```
switch# 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.

```
switch# firmware download usb directory NOS_vX.X.X
```

5. Unmount the USB storage device.

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

Downloading firmware using the manual option

On a Top-of-Rack (ToR) switch, 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 Brocade Technical Support.

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

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

```
switch# 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 switch contain the new firmware.
3. If you entered the **noreboot** option, enter the **reload** command to reboot the switch. If you entered **y** after the prompt, the switch will reboot automatically. The switch performs a reboot and comes up with the new firmware. Your current CLI session will automatically disconnect.

- Log back into the switch. If you entered the `nocommit` option, enter the **firmware commit** command to commit the new firmware. If you entered `y` after the prompt, the switch will commit the firmware automatically upon booting up.

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

- Enter the **show version** command with the **all-partitions** option. Both partitions on the switch 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 by doing the following:

- Execute the **show version all-partitions** command to verify that the MMs and all line-card partitions have the correct firmware.
- Execute the **show ha all-partitions** command to verify that the MMs and all line-card partitions are in HA sync.
- 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 cluster, you must log onto individual nodes in the cluster and issue the **firmware download** command. Alternatively, you can issue the **firmware download logical-chassis** command on the principle node to upgrade multiple nodes at the same time.

This command works on a Logical Chassis only. In a Fabric Cluster, you must log on to individual nodes and run the **firmware download** command.

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
```

4. 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
```

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.

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

```
switch# 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

2. 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
```

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.

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

```
switch# 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                         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
```

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

```

Mixed-node fabric cluster support

A mixed-node fabric cluster is a group of fabric cluster nodes running Network OS v4.1.3 and Network OS v5.0.0. This allows retired hardware to continue to function with hardware running Network OS v5.0.0.

The following limitations and considerations apply to mixed node support:

- Mixed node fabric clusters are supported only in fabric cluster mode and are not supported in logical cluster mode.
- For a fabric cluster running Network OS v4.1.3 and Network OS v5.0.0, the cluster supports the Network OS v4.1.3 feature set. For a list of features that the nodes running Network OS v5.0.0 support, refer to the Network OS v5.0.0 release notes.

Because only Network OS v4.1.3 and Network OS v5.0.0 are supported in a mixed-node fabric cluster, you should:

- Upgrade all Brocade VDX 2740, 6740, 6740T, and 8770 units to Network OS v5.0.0.
- Update all other Brocade hardware to Network OS v4.1.3.

When an fabric cluster is loaded with mixed-node fabric clusters, the cluster only works with the older feature set. Any cluster-wide features that were introduced in Network OS v5.0.0 are not supported in the mixed-node fabric cluster. Most fabric cluster features work properly in a mixed-node cluster. You can perform cluster management normally, as in a normal fabric cluster with all nodes running the same Network OS releases.

During a VCS cluster upgrade, all nodes are not necessarily upgraded at the same time. During this time, the cluster is in a mixed-version state with the VCS cluster in an incomplete state, but the fabric connections remain intact.

TABLE 2 Feature support for mixed-node fabric clusters

Network OS feature	Description of support
Modular HA	Yes, but only on the Network OS v5.0.0 switch.
Modular ISSU	Yes, but only on the Network OS v5.0.0 switch.
Fixed-port ISSU	Yes, but only on the Network OS v5.0.0 switch.
AutoFabric (Pre-provisioning)	No
Flexports	Yes, but only on the Network OS v5.0.0 switch.
IPv6	Yes, but only on the Network OS v5.0.0 switch.
CML	No
REST API	Yes, but only on the Network OS v5.0.0 switch.
OpenStack	No, this feature is only available in Local Chassis mode.
Access Gateway/NPIV	No
Virtual Fabric (Basic + Enhancements)	No

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 Brocade 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 switches after the installation, the following benefits are achieved:

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

This procedure is supported on the following Brocade platforms:


- VDX 2740
- VDX 6740 and VDX 6740T
- VDX 8770-4 and VDX 8770-8

The example approach presented here, tested in a Brocade 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 Brocade VDX cluster VCS 8192. The cluster consists of two spine nodes and two leaf nodes, connected to the core through a vLAG. Two servers, running Brocade 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 3 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 Brocade 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 cluster 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 4 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 Brocade 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.

```
switch# 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 switches.
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.

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

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
 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.

```
switch# 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

switch#
```

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

```
switch# 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)

switch#
```

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

```
switch# 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"

switch#
```

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

```
switch# 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

```

```
switch#
```

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

```
switch# 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 cluster and other details, by using the **show mac-address-table count** and **show mac-address-table** commands.

```

switch# show mac-address-table count
Dynamic Address Count : 11
Static Address Count  : 0
Internal Address Count : 3
Total MAC addresses   : 14
switch#
switch# 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
switch#

```

- 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.

```
switch# 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
```

- h) Copy the running configuration to the startup configuration on all nodes, as well as to an external FTP or SCP server (by means of the **ftp://** or **scp://** options, not shown here).

```
switch# copy running-config startup-config
This operation will modify your startup configuration. Do you want to continue? [y/n]:y
2013/12/05-21:30:52, [DCM-1101], 8374,, INFO, VDX6740, Copy running-config to startup-config
operation successful on this node.
switch#
```

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

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

NOTE

In logical chassis cluster mode, the **copy running-config startup-config** command is not applicable. Use **copy running-config ftp** or **copy running-config scp**.

```
switch# 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)

```
switch#
```

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

```
switch# 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

switch#
```

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

Once the cluster is dual-homed and redundancy in all layers is established, split the cluster 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 cluster 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 cluster 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 Brocade 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. Refer to [Restoring firmware in the VCS Fabric](#).

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 cluster 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 switches (in the example sw81 and sw87) to come back up.

NOTE

The time required for the switches 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 36 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 cluster, 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 switch 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 36 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. Refer to [Restoring firmware in the VCS Fabric](#).

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" switches, sw87 and sw81, were reloaded.

TABLE 5 Traffic-outage times: "Odd" switches, upgrading from 4.0.1 to 5.0.0

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 6 Traffic-outage times: "Odd" switches, reloading within 4.0.1

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" switches, sw79 and sw86, were reloaded.

TABLE 7 Traffic-outage times: "Even" switches, upgrading from 4.0.1 to 5.0.0

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 8 Traffic-outage times: "Even" switches, reloading within 4.0.1

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

Mixed-node fabric cluster support

A mixed-node fabric cluster is a group of fabric cluster nodes running Network OS v4.1.3 and Network OS v5.0.0. This allows retired hardware to continue to function with hardware running Network OS v5.0.0.

The following limitations and considerations apply to mixed node support:

- Mixed node fabric clusters are supported only in fabric cluster mode and are not supported in logical cluster mode.
- For a fabric cluster running Network OS v4.1.3 and Network OS v5.0.0, the cluster supports the Network OS v4.1.3 feature set. For a list of features that the nodes running Network OS v5.0.0 support, refer to the Network OS v5.0.0 release notes.

Because only Network OS v4.1.3 and Network OS v5.0.0 are supported in a mixed-node fabric cluster, you should:

- Upgrade all Brocade VDX 2740, 6740, 6740T, and 8770 units to Network OS v5.0.0.
- Update all other Brocade hardware to Network OS v4.1.3.

When an fabric cluster is loaded with mixed-node fabric clusters, the cluster only works with the older feature set. Any cluster-wide features that were introduced in Network OS v5.0.0 are not supported in the mixed-node fabric cluster. Most fabric cluster features work properly in a mixed-node cluster. You can perform cluster management normally, as in a normal fabric cluster with all nodes running the same Network OS releases.

During a VCS cluster upgrade, all nodes are not necessarily upgraded at the same time. During this time, the cluster is in a mixed-version state with the VCS cluster in an incomplete state, but the fabric connections remain intact.

TABLE 9 Feature support for mixed-node fabric clusters

Network OS feature	Description of support
Modular HA	Yes, but only on the Network OS v5.0.0 switch.
Modular ISSU	Yes, but only on the Network OS v5.0.0 switch.
Fixed-port ISSU	Yes, but only on the Network OS v5.0.0 switch.
AutoFabric (Pre-provisioning)	No
Flexports	Yes, but only on the Network OS v5.0.0 switch.
IPv6	Yes, but only on the Network OS v5.0.0 switch.
CML	No
REST API	Yes, but only on the Network OS v5.0.0 switch.
OpenStack	No, this feature is only available in Local Chassis mode.
Access Gateway/NPIV	No
Virtual Fabric (Basic + Enhancements)	No