

Network OS Troubleshooting Guide, 7.0.1b

Supporting Network OS 7.0.1b

Contents

Preface.....	7
Document conventions.....	7
Notes, cautions, and warnings.....	7
Text formatting conventions.....	7
Command syntax conventions.....	8
Text formatting conventions.....	8
Command syntax conventions.....	9
Notes, cautions, and warnings.....	9
Brocade resources.....	9
Contacting Brocade Technical Support.....	10
Brocade customers.....	10
Brocade OEM customers.....	10
Document feedback.....	10
About this document.....	13
Supported hardware and software.....	13
Using the Network OS CLI	13
What's new in this document.....	13
Using the Chassis ID (CID) Recovery Tool.....	15
CID overview.....	15
Critical SEEPROM data.....	15
Noncritical SEEPROM data.....	15
Automatic auditing and verification of CID card data.....	16
Enabling the CID recovery tool.....	16
Managing data corruption or mismatches.....	16
Understanding CID card failure.....	17
Troubleshooting procedures.....	19
Troubleshooting overview.....	19
Gathering troubleshooting information.....	19
Using a troubleshooting methodology.....	20
Understanding troubleshooting hotspots.....	21
Troubleshooting standard issues.....	28
AMPP is not working.....	29
CID card is corrupted.....	32
Clearing the Boot PROM password.....	34
CPU use is unexpectedly high.....	35
ECMP not load balancing as expected.....	35
ENS not working correctly	35
Fabric does not form correctly.....	36
FCoE devices unable to log in.....	38
Heavy disk utilization.....	40
ISL does not come up on some ports.....	40
License is not properly installed.....	43
Need to recover password for Brocade VDX switches.....	44
Packets are dropped in hardware.....	51
Ping fails.....	57

QoS configuration causes tail drops.....	57
QoS is not marking or treating packets correctly.....	57
RBridge ID is duplicated.....	57
SNMP MIBs report incorrect values.....	58
SNMP traps are missing.....	58
Telnet operation into the switch fails.....	58
Traffic is not being forwarded	60
Trunk member not used.....	60
Upgrade fails.....	62
VCS Fabric cannot be formed.....	63
vLAG cannot be formed.....	63
Zoning conflict needs resolution.....	65
Using troubleshooting and diagnostic tools.....	65
Using Layer 2 traceroute.....	66
Using show commands.....	69
Using debug and system diagnostic commands.....	71
Using SPAN port and traffic mirroring.....	72
Using hardware diagnostics.....	72
Viewing routing information	73
Using the packet capture utility.....	74
Debug Automation Suite for tracing data-path.....	75
TACACS+ Accounting Exceptions.....	85
TACACS+ command-accounting limitations.....	85
Unsupported Network OS command line interface commands.....	85
Supported NTP Regions and Time Zones.....	89
Africa.....	89
America.....	89
Antarctica.....	91
Arctic.....	91
Asia.....	91
Atlantic.....	92
Australia.....	92
Europe.....	92
Indian.....	93
Pacific.....	93

Copyright Statement

© 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>.

Preface

- Document conventions..... 7
- Brocade resources..... 9
- Contacting Brocade Technical Support..... 10
- Document feedback..... 10

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Brocade 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.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, <code>--show WWN</code> .
[]	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. In Fibre Channel products, square brackets may be used instead for this purpose.
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, <code>member[member...]</code> .
\	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.

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.
...	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.

White papers, data sheets, and the most recent versions of Brocade software and hardware manuals are available at www.brocade.com.

Product documentation for all supported releases is available to registered users at MyBrocade.

Click the **Support** tab and select **Document Library** to access documentation on MyBrocade or www.brocade.com. You can locate documentation by product or by operating system.

Release notes are bundled with software downloads on [MyBrocade](#). Links to software downloads are available on the MyBrocade landing page and in the Document Library.

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 should contact their OEM/solution provider.

Brocade customers

For product support information and the latest information on contacting the Technical Assistance Center, go to www.brocade.com and select **Support**.

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"> • Case management through the MyBrocade portal. • Quick Access links to Knowledge Base, Community, Document Library, Software Downloads and Licensing tools 	<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) • Toll-free numbers are available in many countries. • For areas unable to access a toll-free number: +1-408-333-6061 	<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

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.....](#) 13
- [Using the Network OS CLI](#) 13
- [What's new in this document.....](#) 13

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, 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 2741
- 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

This document supports Network OS 7.0.1b.

NOTE

The current version of this guide does not deal with issues specific to IP Fabrics.

What's new in this document

This guide has been updated with content on tracing the data path with a script.

Using the Chassis ID (CID) Recovery Tool

• CID overview.....	15
• Critical SEEPROM data.....	15
• Noncritical SEEPROM data.....	15
• Automatic auditing and verification of CID card data.....	16
• Enabling the CID recovery tool.....	16
• Managing data corruption or mismatches.....	16
• Understanding CID card failure.....	17

CID overview

Each Brocade VDX 8770-4 and VDX 8770-8 contains two chassis ID cards (CIDs) called *CD1* and *CD2*. Most data on each card is identical, and CID2 is used only as a backup if CID1 encounters an issue.

The data contained on the CID card is essential for correct operation of the switch and is accessed most frequently during system startup.

Each CID contains two serial electronically erasable programmable read-only memory (SEEPROM) devices:

- Critical SEEPROM. This SEEPROM is read-only.
- Noncritical SEEPROM. This SEEPROM can be written to by the software.

Critical SEEPROM data

The critical SEEPROM contains the following:

- A header with the CID part number, serial number, and other data about the CID card. If this data is corrupted or cannot be accessed, the card is identified as faulty in RASLogs:
 - [EM-1003], M1 | FFDC, CRITICAL, ..., CID 2 has unknown hardware identifier: FRU faulted.
 - [FW-1432], M1, WARNING, sw0, Switch status change contributing factor CID-Card: 1 bad.
- A chassis part number and serial number. Cluster configuration management uses the serial number to uniquely identify the chassis in the fabric.
- An eight-byte number that represents both the license ID and World Wide Name (WWN) base value for the chassis. The license ID is used to validate installed licenses. Licenses are invalid if the license ID is not available. The WWN is used to identify the switch in a fabric.

Noncritical SEEPROM data

The noncritical SEEPROM contains the following data sets.

- The FRU history table, which contains logs of insertions and removals of FRUs into and from the chassis. The content of this table is not audited or verified.
- The IP data table, which contains management module and chassis management IP addresses/masks, the IP default gateway, and the chassis name.
- A power-off list, which controls the order in which blades are automatically powered off if an impending power loss is detected.

- A set of Data Center Ethernet (DCE) data containing chassis MAC addresses, without which the switch will not function.

Automatic auditing and verification of CID card data

The contents of both CID cards are verified on a periodic basis and whenever an event indicates that an issue may exist.

Under normal circumstances, the CID card audit is run for about one hour after a system startup or restart, then is repeated every 24 hours. If no errors occur, no action is taken.

If CID card errors occur, if mismatches between data sets on the two CID cards are detected, or if a card is inserted, RASLogs are shown on the console:

- [EM-1020], ... M1, ERROR, ... A problem was found on one or both CID cards (x), please run the *cidrecov* tool to get more information and recovery options.
- [EM-1021], ... M1, INFO, ... A CID card has been inserted, a CID verification audit will be run to detect any mismatches or other problems.
- [EM-1022], ... M1, WARNING, ... A CID card access problem has been encountered, please run the *cidrecov* tool to get more information and recovery options.

Enabling the CID recovery tool

You should run the CID recovery tool when instructed by a RASLog, and you can also run the tool if you suspect an issue with one or both of the CID cards. To run the CID recovery tool, enter the **cidrecov** command in privileged EXEC mode on the command line:

```
sw0# cidrecov
```

Managing data corruption or mismatches

If **cidrecov** detects any CID 1 or CID 2 non-critical SEEPROM corruption or mismatches, the tool displays related data and the following data-recovery options as applicable for each data-set error:

- **Exit.** Select this option if you do not want to change any data values.
- **Recover with default values.** Select this option if you want to reset all data in the data set to the factory defaults. For IP data, dummy IP addresses and masks are written. DCE and chassis-configuration data are based on the chassis type.

A system restart repopulates IP addresses and chassis names that appear in the startup configuration file. If you want to manually change the IP data, you can use the **ip-address**, **chassis virtual-ip**, and **chassis-name** commands. For more information, refer to the *Network OS Command Reference*.

- **Recover BAD from GOOD.** This option is offered only if one CID card contains good data and the other card contains corrupt data. If you select this option, **cidrecov** copies the good data onto the affected card.
- **Recover CID 2 from CID 1** and **Recover CID 1 from CID 2.** These options are offered only if the data on both CID cards is good but there is a mismatch. You can select which card to use to overwrite data on the other card.

The following is an example of running the **cidrecov** tool, receiving errors that can be fixed, and selecting the **Recover BAD from GOOD** option (note that the example below contains only some of the actual output):

```
sw0# cidrecov
CID 1 Non-Critical Seeprom is Inaccessible or Corrupted.
  CID Non-Critical Seeprom Problem Details
CID 1 Non-Critical Seeprom IP address Control Data Checksum Bad !!!!
  CID Recovery Options
```



```
0. Exit
1. Recover with default values
2. Recover BAD from GOOD
Enter Selection > 2
Copy IP Data table...
  Copy 384 bytes from CID 2 to CID 1, num blks 1 resid 128
  Read block 1 from CID 2 succeeded
  Write block 1 to CID 1 succeeded
  Read last block from CID 2 succeeded
  Write last block to CID 1 succeeded
  copy successful
Copy succeeded for all data types attempted
IP Address CID Recovery completed.
```

Understanding CID card failure

If the critical SEEPROM of a CID card contains any errors, or if the noncritical SEEPROM cannot be read, then recovery is not possible and the following message is displayed:

```
Recovery is not possible. Please contact Brocade Technical Support for replacement of the inaccessible
CID(s).
```


Troubleshooting procedures

- [Troubleshooting overview.....](#) 19
- [Troubleshooting standard issues.....](#) 28
- [Using troubleshooting and diagnostic tools.....](#) 65

Troubleshooting overview

This chapter provides tips and procedures for troubleshooting issues that may occur while operating a Brocade switch running Network OS. It also introduces some of the common troubleshooting tools.

Gathering troubleshooting information

The first step in any successful troubleshooting is to gather the appropriate information (including *supportSave* data). For details refer to [Using a troubleshooting methodology](#) on page 20 and [Capturing supportSave data](#) on page 19.

Capturing supportSave data

Capturing *supportSave* data is key to successful troubleshooting. The **copy support** command not only runs diagnostic commands, but also gathers core dumps, trace files, and other relevant data. In the same action, the command also copies all this information to a remote host. Once on the remote host, your switch provider can proceed to analyze the problem. Meanwhile, your switch can be returned to production with minimal downtime.

To capture supportSave data, complete the following steps:

1. Log in to the switch.
2. In privileged EXEC mode, enter the **copy support** command to capture the supportSave data.

The **copy support** command has options to copy the supportSave files to a remote server using FTP or SCP, or you can save to a local USB device. You can use the command in a single command line, or in interactive mode.

The following example uses the single command line mode to copy the supportSave files to a remote host using FTP.

```
switch# copy support ftp host 10.38.33.131 user admin directory 108
Password: *****
```

The following example uses the interactive form of the command and FTP:

```
switch# copy support-interactive
Server Name or IP Address: 10.38.33.131
Protocol (ftp, scp): ftp
User: admin
Password: *****
Directory:/home/admin/support
VCS support [y/n]? (y): y
```

Using information resources

The following information is helpful for incident investigation and resolution when you contact your switch-support provider:

- A network diagram and topology information
- A record of the steps and events leading to the incident

- Lists of applications, management agents, and scripts running at the time of the incident
- supportSave files
- Output from the **show media** command if the issue is related to SFP transceivers
- Outputs from any commands run while attempting to troubleshoot the problem yourself
- Any network traces captured using Wireshark software or other network analyzer.
- Terminal Access Controller Access-Control System (TACACS) server version if the issue is related to TACACS.

Using a troubleshooting methodology

Once all relevant information is collected, success is improved significantly with a sound troubleshooting approach.

This section outlines a methodology for troubleshooting issues. It introduces steps that you might consider using, depending on the issue in question.

1. Check whether the switch has all of the relevant licenses:
 - Available licenses are: POD1, POD2, 10G Port Upgrade, 40G Port Upgrade, FCoE, Layer 3, and Advanced Features.
 - After adding or modifying a POD or port upgrade license, re-enable the ports.
2. Verify the topology and switch configuration as conveyed by the switch
3. Enter the **copy support** command.
4. Run other relevant show commands (for example, **show logging raslog**) to look for clues or triggers of the reported failure.
5. Check the utilization of various resources.
 - a) Enter the **show process cpu** command to determine CPU use.
 - b) Enter the **show process me** command to determine memory use.
 - c) Enter the **show mac-address-table count** command to determine the number of MAC addresses used.
 - d) Enter the **show fabric route topology** command to determine the number of routes.
 - e) Enter the **show fabric all** command to determine the number of VCS Fabric nodes.
 - f) Enter the **show media** command to investigate any optics issues.

6. Conduct data-path fabric continuity tests:
 - a) Issue pings from and to the end-stations or devices.
 - b) Check the counters in the output of the **show interface** command to detect if packets are coming in or are being dropped as errors.
 - c) Verify that optics used are Brocade-certified. Enter the **show media interface** command and verify that the Vendor name field shows "Brocade." Check also that the Tx and Rx Power fields are not zero.
 - d) Verify that the MAC address table learns the MAC addresses.
 - e) If the switch is part of a VCS Fabric cluster, verify that the MAC address tables are synchronized properly across all Brocade VDX switches in the cluster.
 - f) Check whether LLDP reports neighbors.
 - g) Check the Ethernet Name Server (ENS) functionality by ensuring that the MAC address table reports MAC addresses learned from other VCS Fabric switches.
 - h) Use the **l2tracert** command for validating the data-path fabric continuity. This command helps identify where the packets are being dropped within the fabric.

The command prompts for some basic and allows you to choose to enter some extended parameters. Currently supported basic parameters include:

- Source Address (SA) and Destination Address (DA) of dynamically learned MAC addresses
- VLAN
- Edge routing bridge (RBridge) ID

Currently supported extended parameters include:

- Protocol type (IP)
- Source and destination IP addresses
- IP protocol type (recommend TCP)
- Source and destination port numbers

The purpose of IP parameters is to provide a way to make the traceroute packet traverse a specific ECMP link.



CAUTION

The following step affects configuration and should be used with care.

7. To track certain flows within the fabric, use permit ACLs and monitor the hit increments.

Understanding troubleshooting hotspots

This section provides relevant background information and best practices guidance related to features of Network OS where problems have been reported. With this guidance, you should be able to avoid many potential problems.

Licensing

When a licensed feature does not work, one likely cause is that the license has not been installed correctly. Follow the guidelines and procedures in the *Network OS Software Licensing Guide* to ensure your features are licensed properly and those licenses installed correctly.

For license recovery procedures, refer to the *Network OS Software Licensing Guide*.

STP interoperability with Brocade MLX or other switches

- To use the Spanning Tree Protocol (STP) in a network with Brocade MLX switches, or switches from other vendors such as Juniper or Cisco, you may have to configure the interface to send BPDUs to the shared spanning tree MAC address 0100.0ccc.cccd. Without this setting, the RPVST/PVST root bridge is not recognized on VLANs other than VLAN 1.

To interoperate with MLX switches or other vendors' switches, enter the following command in interface configuration mode:

```
switch(config-if-te-0/1)# spanning-tree bpdu-mac 0100.0ccc.cccd
```

- If a Brocade IP switch has a VLAN is configured with tagged ports and Rapid Spanning Protocol (RSTP) is enabled under the VLAN (PVST), then BPDUs from the tagged ports received by the Brocade VDX switch will be dropped if pvst-mode is not configured under the ports that are in the VLAN and connected to the Brocade VDX switches.

The following example shows a configuration on a Brocade IP switch with tagged ports and RSTP enabled under the VLAN:

```
vlan 2
tagged ethe 1/24 ethe 2/1 to 2/2
router-interface ve 2
rstp priority 100
```

If the conditions are met, then all the ports should have pvst-mode configured so that tagged BPDUs pass through the Brocade VDX switch. If pvst-mode is not enabled, enable it as follows:

```
Brocade(config)# interface ethernet 2/1
Brocade(config-if-2/1)# pvst-mode
```

Load balancing distribution

Understanding issues related to load balancing requires some basic knowledge of the criteria used by load balancing algorithms. The table below provides details for each feature that provides load balancing.

TABLE 1 Load balancing algorithms

Feature	Algorithm
ECMP IP	<p>Paths are selected on the basis of a hash derived from the following parameters:</p> <ul style="list-style-type: none"> Source MAC address Destination MAC address VID IP protocol Source IP address Destination IP address Layer 4 source port Layer 4 destination port <p>You can configure the hashing fields using the fabric-ecm load-balance and fabric-ecmp load-balance-hash-swap commands.</p> <p>For related recovery procedures, refer to ECMP not load balancing as expected on page 35.</p>
ECMP FCoE	<p>Paths are selected on the basis of a hash derived from the following parameters:</p> <ul style="list-style-type: none"> Input Port ID Source MAC address Destination MAC address VID FID SID

TABLE 1 Load balancing algorithms (continued)

Feature	Algorithm
	<ul style="list-style-type: none"> • DID • OXID
LACP	Provides adaptive load balancing based on up to seven criteria (7-tuple), depending upon what fields are available in the frame.
Brocade trunk	Provides equal packet load balancing (round-robin) among member links.

Static assignment of the routing bridge ID

Duplicate routing bridge (RBridge) IDs are a common source of error when a switch is added to an Ethernet fabric. Before adding a switch to an Ethernet fabric, you must assign it a unique RBridge ID. If the new switch is to be added to an existing VCS Fabric cluster, it must be assigned the same VCS ID as other switches in the cluster. Once the switch is added, the principal routing bridge performs the negotiation in the control plane to include the new switch and rebuild the fabric. The data plane remains unaffected.

Procedures for recovering from duplicate routing IDs are provided in [RBridge ID is duplicated](#) on page 57.

FSPF route change

When the Fabric Shortest Path First (FSPF) algorithms select a new route, a temporary disruption of traffic can occur. This behavior is normal as the old path is first deleted and then the new path is programmed. Such path changes can occur when FSPF calculates a new shortest route, or when the current path is down.

vLAG overview

You should be aware of the following aspects of the vLAG feature before troubleshooting vLAG problems:

- Multicast (BUM) traffic in vLAG
- Edge-port feature requirements
- Failover

Multicast traffic in vLAG

Flooding traffic always goes through a primary link of the vLAG. You should consider this restriction when provisioning bandwidth for most traffic. This link is marked with an asterisk (*) in the output of the **show port-channel** command.

```
switch# show port-channel 38
LACP Aggregator: Po 38
Aggregator type: Standard
Admin Key: 0038 - Oper Key 0038
Partner System ID - 0x8000,01-e0-52-00-20-00
Partner Oper Key 0038
Member ports:
Link: Te 0/13 (0x180D0102) sync: 1
Link: Te 0/14 (0x180E0103) sync: 1 *
```

Edge-port feature requirements for vLAG

LACP can be configured on edge ports only with either Brocade or Standard types. If Brocade is chosen, so that Link Reset (LR) primitives are exchanged properly, make sure that the edge peering device is a Brocade Converged Network Adapter (CNA), a standalone Brocade VDX switch, or a Brocade VDX 8000 series switch.

Failover and vLAG

For the fast failover convergence requirements, Brocade recommends using the **vlag ignore-split** command, which enables sub-second failover times. This command is included in all port-channel configurations.

When planning to deploy this feature in production, use care to prevent a "split-brain" scenario, in which vLAG members detach from each other. Brocade recommends having more than one interswitch link (ISL) between the vLAG member switches and physically routing them through separate conduits and cable trays. Secondly, Brocade strongly recommends using topologies that are certified by Brocade.

NOTE

Brocade does not recommend using vLAG failover in a network with Cisco or Juniper switches that are connected using copper. Brocade has observed greater than one-second failover times in networks with this hardware.

vLAG and split-brain

The following topics discuss the split-brain scenario and how to mitigate it.

Understanding "split-brain"

A split-brain can occur when the end-hosts or edge switches are connected to two separate cluster switches by way of a vLAG (using LACP). The end-devices perceive those two cluster switches as one switch because they have the same system ID advertised in LACP.

Under rare conditions, when all the ISLs between the two cluster switches are broken and both the cluster switches continue to advertise the same system ID to their LACP partner, a "segmented fabric" or "split-brain" condition exists, where the end-host or edge switch might not detect this segmentation and could continue to treat both the vLAG switches as one switch.

ATTENTION

This condition can cause packet duplication or unexpected packet loss.

Traffic protection during split-brain conditions

By default, Network OS has a capability to recover gracefully from the split-brain scenario. When all the ISLs between the VDX cluster switches go down, the switch with the lower RBridge ID uses LACP to inform the edge-switch partner that it has segmented out of the port-channel. It does this by changing its advertised system ID. When the edge switch learns a different system ID on one of its members, it removes this member from that port-channel, and continues to function with only one vLAG member — the switch with the higher RBridge ID. The other vLAG member switch still has the link up, but remains segmented out of the original port-channel (sync: 0). This capability prevents duplication of packets or potential packet drops resulting from a split-brain scenario.

When a member switch is reloaded

Reloading the switch with the lower RBridge ID has no impact.

When the switch with the higher RBridge ID is reloaded, the other vLAG member perceives all of its ISLs as down. Though this is *not* a real split-brain scenario, the switch with the lower RBridge ID may not be able to differentiate, and thus would inform the partner about a changed system ID. The partner edge switch would detect two events:

- The system ID on one link changes.
- The other interface goes down.

In such a case, LACP will renegotiate and reform the port-channel, which could flap the port-channel, impacting traffic momentarily. The same effect could occur when the switch boots up and joins the fabric again.

Thus, if the switch with the higher RBridge ID is reloaded, the potential impact could be a port-channel flap that can momentarily disrupt traffic. Notice that this effect does not occur when the switch with the lower RBridge ID is reloaded.

Avoiding traffic disruption during switch reload

Network OS switches offer flexibility to the user by providing a special vLAG **ignore-split** option that you can configure for the logical port-channel. This option should be configured on both vLAG member ports.

Configuring this option prevents the switch with the lower RBridge ID from changing its system ID, so both switches will continue to advertise the same system ID. This action prevents the partner edge switch from detecting a change when one of the member switches is reloaded and the traffic is handled gracefully.

Using the vLAG ignore-split option

To use the vLAG **ignore-split** option, redundancy should be built around ISLs to prevent a situation in which all ISLs are broken at the same time. Brocade recommends using multiple ISLs, and routing those ISLs through different physical paths or conduits to eliminate the possibility of accidental damage to all links at the same time.

Principal routing bridge availability

If a new principal routing bridge is introduced into a working VCS Fabric cluster, or if the principal routing bridge is lost and a new switch must be elected, the fabric is rebuilt from the control-plane viewpoint, whereas the data plane continues to forward traffic without disruption. The primary responsibilities of the principal routing bridge in a VCS Fabric are:

- RBridge ID allocation
- Ownership of virtual management IP address
- Keeping the configuration database synchronized

Brocade trunks

Brocade trunking is the only aggregation method that works using ISLs. Brocade ISL trunks are formed automatically with other switches using Line Reset (LR) primitives signaling with the peer switch. For a successful trunk formation, all ports must be part of the same trunk group and must be configured at the same speed. The trunk is turned on by default.

The table below shows the allocation of port numbers to trunk groups for Brocade VDX switches.

TABLE 2 Trunk groups

Brocade platform	Trunk groups	No. of trunk groups per platform
VDX 6740 series	<ul style="list-style-type: none"> • 1-15 • 16-32 • 33-40; 49-50 • 41-48; 51-52 	4
VDX 6940-36Q	<ul style="list-style-type: none"> • 1-9 • 10-18 • 19-27 • 28-36 	4
VDX 6940-144S (Middle of Row [MoR])	<ul style="list-style-type: none"> • 1-24; 61-72 • 25-48; 73-84 • 85-96 • 103-108 	6 (considering TG-3/TG-3A; and TG-4/TG-4A as distinct trunk groups)

TABLE 2 Trunk groups (continued)

Brocade platform	Trunk groups	No. of trunk groups per platform
	<ul style="list-style-type: none"> • 49-60 • 97-102 	
VDX 8770 (with VDX LC48x1G line card)	<ul style="list-style-type: none"> • 1-8 • 9-16 • 17-24 • 25-32 • 33-40 • 41-48 	6 per 1G blade
VDX 8770 (with VDX LC48x10G line card)	<ul style="list-style-type: none"> • 1-8 • 9-16 • 17-24 • 25-32 • 33-40 • 41-48 	6 per 10G blade
VDX 8770 (with VDX LC12x40G line card)	<ul style="list-style-type: none"> • 1-2 • 3-4 • 5-6 • 7-8 • 9-10 • 11-12 	6 per 40G blade
VDX 2741	<ul style="list-style-type: none"> • 29-36 • 37-44 • 45-46 • 47-48 	3
VDX 2746	<ul style="list-style-type: none"> • 43-50 • 51-56 • 57-58 	3

NOTE

Brocade trunks are not supported over 1-Gbps links.

To utilize the advantages of Brocade trunking between VDX switches, Brocade recommends having at least a two-member trunk and multiple ECMP paths. Brocade also recommends routing the cables in a trunk through separate conduits to ensure connectivity in case a conduit is accidentally cut.

NIC teaming with vLAG

NIC teaming permits link aggregation between server and switch. It can be one of two types: active/passive model or active/active model. For the active/passive model, you may not need to configure a LAG on the switch side, as unique MAC addresses will be seen on only one link.

For the active/active model, the same MAC address may appear on both the links terminating on a switch (or pair of switches). In such a case, you must configure a LAG on the switch side.

Selecting the MTU

Always set the switch MTU to the maximum host MTU plus 100 bytes. This method is recommended because the definition of MTU sometimes varies among different vendors. If the switch MTU is set to the same as the connected host MTU, packets could be dropped.

Avoiding oversubscription

Under certain congestion conditions, you may observe incrementing packet drops representing "tail-drops" in the output of the **show qos rcv-queue interface tengigabitethernet** command, as shown underlined in the following example:

```
switch# show qos rcv-queue interface tengigabitethernet 5/0/1
Interface TenGigabitEthernet TenGigabitEthernet 5/0/1
In-use 0 bytes, Total buffer 144144 bytes
0 packets dropped
-----
CoS      In-use   Max
      Bytes   Bytes
-----
0         0       18018
1         0       18018
2         0       18018
3         0       18018
4         0       18018
5         0       18018
6         0       18018
7         0       18018
```

In such conditions, you must first identify the bottleneck, and then take action to mitigate the congestion.

Identifying the congestion bottleneck

To identify the bottleneck in the Brocade VDX network, enter the **show interface** command at various locations, and identify interfaces with incrementing TX and RX discards. Depending upon the TX or RX discards, the congestion could be anywhere downstream.

Mitigating the congestion

Try the following actions to mitigate congestion:

- Increase bottleneck bandwidth.
 - Add more links to the LAG and ECMP paths.
 - Use higher-speed interfaces.
- Implement flow control on the bottleneck and on neighboring devices.
- Implement QoS congestion management schemes.
 - Classify, mark, and prioritize critical traffic.
 - Modify scheduling schemes. Consider and compare the effects of using strict priority or deficit weighted round-robin (DWRR) scheduling schemes.

For the flow control solution, enable flow control either on the ports receiving the traffic from end-devices (servers or personal computers) and the connected end-device itself, or enable flow control on the port-channel as shown in the following example.

```
switch(conf-if-te-1/0/24)# interface port-channel 100
switch(config-Port-channel-100)# qos flowcontrol tx on rx on
```

Once flow control is enabled, enter the **show qos rcv-queue interface tengigabitethernet** command again and check the output. It should no longer be reporting packet drops. If the packet drops continue or the ingress rate is considerably lower than expected, contact your switch support provider for further investigation.

We recommend enabling asymmetric flow control with Brocade VDX switches. For any two adjacent devices, one device should have Rx ON and Tx OFF, while the other device should have Rx OFF and Tx ON.

Refer to the “Congestion control and queuing” section of the *Network OS Layer 2 Switching Configuration Guide* for further details about congestion control.

ACL limits issues

If you keep within the supported limits of ACL usage as shown in the table below, you are unlikely to run into system limits issues. ACLs should instantiate quickly and correctly.

TABLE 3 ACL limits per switch in VCS mode

System resource	Brocade VDX 6740 series Brocade VDX 6940 series Brocade VDX 2741 Brocade VDX 2746	Brocade VDX 8770
MAC (L2) ACLs created	512	2048
IPv4 (L3) ACLs created	512	2048
IPv6 (L3) ACLs created	512	2048
Number of rules per ACL	256	2048
Total number of rules in all ACLs	30K	30K
ACL name length	63 characters	63 characters
Range of ACL sequence numbers	From 0 through 4294967290	From 0 through 4294967290

As you approach or exceed combinations of these limits, you might encounter slow instantiation of ACL rules.

Delays of several minutes can occur in the instantiation of ACL rules and counters if the number of ACLs or VLANs is excessive.

To display the hardware instantiation status (Active/Partial/In Progress/Inactive), run the **show access-list** command .

Troubleshooting standard issues

This section describes some potential problems you may encounter and suggestions on how to investigate or resolve each issue. If these steps do not lead to resolution of the problem, prepare a case for your switch provider, as described in [Contacting Brocade Technical Support](#) on page 10.

- [AMPP is not working](#) on page 29
- [CID card is corrupted](#) on page 32
- [CPU use is unexpectedly high](#) on page 35
- [ECMP not load balancing as expected](#) on page 35
- [ENS not working correctly](#) on page 35
- [FCoE devices unable to log in](#) on page 38
- [Traffic is not being forwarded](#) on page 60
- [Heavy disk utilization](#) on page 40
- [ISL does not come up on some ports](#) on page 40
- [License is not properly installed](#) on page 43

- [Packets are dropped in hardware](#) on page 51
- [Need to recover password for Brocade VDX switches](#) on page 44
- [Ping fails](#) on page 57
- [QoS configuration causes tail drops](#) on page 57
- [QoS is not marking or treating packets correctly](#) on page 57
- [RBridge ID is duplicated](#) on page 57
- [SNMP MIBs report incorrect values](#) on page 58
- [SNMP traps are missing](#) on page 58
- [Telnet operation into the switch fails](#) on page 58
- [Trunk member not used](#) on page 60
- [Upgrade fails](#) on page 62
- [VCS Fabric cannot be formed](#) on page 63
- [vLAG cannot be formed](#) on page 63
- [Zoning conflict needs resolution](#) on page 65
- [Fabric does not form correctly](#) on page 36

AMPP is not working

Configuring Brocade Automatic Migration of Port Profiles (AMPP) is complex. For details on configuring AMPP, refer to the “Configuring AMPP” section of the *Network OS Layer 2 Switching Configuration Guide*.

Problems encountered while using AMPP are usually the result of configuration errors in the port-profile itself, errors in the associated virtual machine (VM) configuration, or compatibility problems between the host adapters and AMPP. Specifically, AMPP problems can be caused by the following conditions:

- A port-profile configuration does not exist on the target switch or does not contain a basic switchport and VLAN configuration. Refer to [Verifying the port-profile configuration](#) on page 30.
- The VM MAC address does not appear in the MAC address table. Refer to [Verifying the VM MAC address](#) on page 30.
- The port-profile is not activated or is not associated with the correct MAC address. Refer to [Verifying the port-profile state](#) on page 30.
- The VM kernel MAC addresses are not associated correctly with the port-profile on the respective switches. Refer to [Verifying the VM kernel MAC addresses](#) on page 31.
- The VM and its associated hosts do not share a common storage device. Refer to [Verifying a shared storage device](#) on page 31.
- The port-profile was learned on a nonprofiled VLAN. Refer to [Verifying the status of a learned profiled MAC address](#) on page 31.
- A conflicting port-profile is applied to the same interface. Refer to [Verifying that port profiles do not conflict](#) on page 32.
- The Ethernet Name Server is not functioning correctly. Refer to [Verifying the Ethernet Name Server](#) on page 32.
- An ESX host has an incompatible network adapter or driver installed. Refer to [Verifying an ESX host](#) on page 32.

Verifying the port-profile configuration

A valid port-profile must exist on the target switch. It must contain a basic switchport and VLAN configuration.

1. In the privileged EXEC mode, enter the **show running-config port-profile** command to verify that the port-profile configuration exists on the target switch, and that it contains a basic switchport and VLAN configuration.

```
switch# show running-config port-profile
port-profile default
vlan-profile
switchport
switchport mode trunk
switchport trunk allowed vlan all
switchport trunk native-vlan 1
!
!
port-profile pp1
vlan-profile
!
!
port-profile pp2
vlan-profile
!
!
```

2. If the port-profile configuration does not exist or is missing the required switchport or VLAN configuration, create the port-profile as described in the “Configuring AMPP profiles” section of the *Network OS Layer 2 Switching Configuration Guide*.

Verifying the VM MAC address

For the correct functioning of AMPP, the MAC address for the VM and its associated hosts must appear in the MAC address table.

1. Enter the **show mac-address-table** command to verify that the VM MAC addresses appear in the switch MAC address table.

```
switch# show mac-address-table
VlanId  Mac-address      Type      State      Ports
1        0000.0010.0001    Static    Inactive    Te 4/0/3
1        0000.0010.0002    Static    Inactive    Te 4/0/3
Total MAC addresses : 2
```

2. If a VM MAC address is not present, contact your switch support provider for further investigation and provide this data.

Verifying the port-profile state

For the correct functioning of AMPP, the port-profile must be active and must be associated with the correct MAC address.

1. Enter the **show port-profile status** command to verify that the port-profile is activated and is associated with the correct MAC address.

```
switch# show port-profile status
Port-Profile  PPID  Activated  Associated MAC  Interface
pp1           1     No         None            None
pp2           2     No         None            None
```

2. Correct any misconfigurations as follows:

- If the port-profile is not activated, enter the **port-profile *profile-name* activate** command to activate it.
- If the port-profile is not associated with a MAC address, enter the **port-profile *port-profile-name* static** command to perform the association.

```
switch(config)# port-profile PP3 static 0050.5600.10030
```

- If the port-profile is associated with the wrong MAC address, enter the **no port-profile *port-profile-name* static** command to break the association with the incorrect MAC address, and then reassociate the port with the correct MAC address.

```
switch(config)# no port-profile PP3 static 0050.5600.10020
switch(config)# port-profile PP3 static 0050.5600.10030
```

Refer to the “Configuring a new port-profile” section of the *Network OS Layer 2 Switching Configuration Guide* for details about activating a port-profile and associating a port-profile with a MAC address.

Verifying the VM kernel MAC addresses

Confirm that the virtual machine (VM) kernel MAC addresses are also associated with the port-profile on the respective switches. If not, perform the association as described in [Verifying the port-profile configuration](#) on page 30.

Verifying a shared storage device

Confirm that the VM and its associated hosts are sharing a storage device. If not, then reconfigure the VM and hosts to share a storage device.

Verifying the status of a learned profiled MAC address

For correct functioning of AMPP, the MAC address must be learned from a valid source— a profiled VLAN. This procedure determines whether a MAC address was learned from a valid source.

Enter the **show mac-address-table port-profile** command to check the status on learned profiled MAC addresses.

```
switch# show mac-address-table port-profile
Legend: Untagged(U), Tagged (T), Not Forwardable(NF) and Conflict(C)
VlanId  Mac-address      Type      State      Port-Profile      Ports
1       0050.5679.5351   Dynamic   Active     Profiled(U)       Te 111/0/10
1       0050.567b.7030   Dynamic   Active     Profiled(U)       Te 111/0/12
1       005a.8402.0000   Dynamic   Active     Profiled(T)       Te 111/0/24
1       005a.8402.0001   Dynamic   Active     Profiled(NF)      Te 111/0/24
1       005a.8402.0002   Dynamic   Active     Not Profiled      Te 111/0/24
1       005a.8402.0003   Dynamic   Active     Not Profiled      Te 111/0/24
1       005a.8402.0004   Dynamic   Active     Not Profiled      Te 111/0/24
(output truncated)
Total MAC addresses : 17
```

Check for and investigate MAC addresses identified in the output as “Not Profiled.”

Verifying that port profiles do not conflict

1. Enter the **show port-profile name pp1_name name pp2_name validate** command to validate whether multiple port-profiles applied on an interface can co-exist without conflict.

```
switch# show port-profile name pp1 name pp2 validate
Port-Profile          Port-Profile          Conflicts
-----
pp1
vlan-profile          vlan-profile          No
qos-profile           qos-profile           No
security-profile      security-profile      No
```

2. If a conflict exists, reconfigure one of the port-profiles to avoid the conflict.

Refer to the e “Configuring AMPP” section of the *Network OS Layer 2 Switching Configuration Guide* for information about the rules for co-existence.

Verifying the Ethernet Name Server

AMPP requires each VCS Fabric switch in the cluster have the same view of the MAC address table. Any differences in the view indicate a failure of the Ethernet Name Server (ENS). Refer to [ENS not working correctly](#) on page 35 for details.

Verifying an ESX host

Verify that each ESX host has the correct Converged Network Adapter (CNA) installed with appropriate drivers, and does not use the Cisco Nexus 1000V software switch, as that switch might send out specially crafted packets.

CID card is corrupted

In the case of a corrupted CID card, perform the following steps.

1. Link the **wncardshow** command to survey the extent of the damage. (This does not have to be done for single boards.)

```
switch# ln -s /fabos/cliexec/em /fabos/bin/wncardshow
```


2. Display the wwncardshow data.

```
switch# wwncardshow ipdata
packet count is 2
++ Wwn Card IP Data ++
Type Num Field Address Mask Cfg/Zone
-----
CP 0 Eth IP: 255.255.255.255 255.255.255.255
CP 1 Eth IP: 255.255.255.255 255.255.255.255
Chassis GW IP: 255.255.255.255
LicID: 10:00:00:ff:ff:ff:ff:ff enet cfg
Name: VDX 6710-54 Gen# : -1/0
Sw 0 Eth IP: 10.17.10.84 255.255.240.0
FC IP: 0.0.0.0 0.0.0.0
GW IP: 10.17.0.1
WWN: 10:00:00: 05:33:14:b2:70
Name: swd77 Gen# : 0/0
Sw 1 Eth IP:
FC IP:
GW IP:
WWN: 10:00:00:05:33:14:b2:71
Name: Gen# : 0/0
```

Items that are FFs, 255s, or zeros are unacceptable. Only the first two groups count, and the items that must be correct are the following:

- - The CP Eth IP entries. They need valid data only if that CP/MM is present.
- The chassis LicID entry.
- The Sw 0 Eth IP entry.
- The Sw 0 GW IP entry.
- The Sw 0 WWN entry.

3. To correct the CP Eth IP entries, run **ipaddrset -cp x**, where x is 0 for MM1 and 1 for MM2, and put in correct data at the prompts. Then run **ipaddrset -chassis** and enter the correct data as needed.

Sometimes, if the entries have enough 255/0xFFs in them, running **ipaddrset** does not update the values properly, in which case you have to run **test_sysmod** to clear a couple of entries.

4. To correct Sw 0 WWN, enter **wwn -d626 xx:xx:xx:xx:xx:xx:xx** with the correct wwn value. The system must be rebooted for the change to take effect (at the prompt or manually).
5. To correct chassis LicID, you need the test_sysmod tool. Mount a filesystem (if necessary get eth0 up manually with ifconfig, or set the gateway first).

```
switch# run test_sysmod
test_sysmod
```

6. At the first menu, enter 11 for WWN testing, then 2 for copy WWN to LID, and then enter 1 to confirm. Perform a **Ctrl+C** to exit.
7. The system must be rebooted for the change to take effect. Exit test_sysmod with **Ctrl+C**.

If you have lost both the WWN and license ID, then you must perform step 4 first. If you do not know the value, it is available in the MAC address in the boot environment variables (for pizza boxes only).

This value can be entered in the **wwn** command by inserting 10:00: before the MAC value).

8. Finally, if you can't correct the IP addresses, there is one more option in test_sysmod that can help. At the main menu, enter 11 for WWN testing and then 1 for clear WWN IP data entry, then 0, 1, 2, or 3 for entries that had a lot of FFs. If you clear all of the entries that are corrupted with FFs, you should be able to run **ipaddrset** to restore the real addresses.
9. Reboot the switch in order for the change to take effect and make the **ipaddrset** command available.

Verifying SEEPROM data

1. To verify the SEEPROM, copy the `test_symod` file to `/fabos/bin` as `test_sysmod`, and select option **10** for i2c and option **27** to Verify FRU Seeprom. The test begins automatically.
2. Use the offset of `0x6a4c`, as that is where the IP table starts (size 256), but any offset (and size less than or equal to 256) will access that device.

Clearing the Boot PROM password

After you complete the procedure [Obtaining the Boot PROM recovery password](#) on page 50, the BootPROM password is set. To avoid needing the Boot PROM password during future password-recovery operations, you can reset the Boot PROM password.

To reset the Boot PROM password, perform the following steps:

1. Connect to the serial console port of the switch.
2. Manually reboot the switch.
3. When prompted to stop test or stop AutoBoot, press **ESC**.

NOTE

If the **ESC** key is not effective during reboot, turn the power off and back on, and then try again. If the **ESC** key is still not effective, check the serial console cable. If the cable is connected correctly, then the unit must be returned for service or repair.

The Boot PROM menu is displayed with the following options:

Start system.	Used to reboot the system.
Recover password.	Used to generate a character string for your support provider to recover the Boot PROM password. ATTENTION Use this feature only when directed by technical support personnel.
Enter command shell.	Used to enter the command shell to reset all passwords on the system.

```
Checking system RAM - press any key to stop test
Checking memory address: 00100000
System RAM test terminated by keyboard
set_bootstatus: BS_LOAD_OS, platform_idx = 6
Hit ESC to stop autoboot: 0
```

- 1) Start system.
- 2) Recover password.
- 3) Enter command shell.

Option?

4. Enter 3 at the prompt to open the command shell.
5. At the prompt, enter the Boot PROM password.

```
password: *****
=>
```

6. To reset the password, enter the **resetpw** command.

```
=> resetpw
.
.
Done
```

- To allow the switch to continue booting up, enter the **reset** command.

```
=> reset
do_reset: PERFORM HARD RESETi
The system is coming up, please wait...
```

When the boot-up process is finished, the Boot PROM password is gone.

CPU use is unexpectedly high

Unexpectedly high CPU use is usually the result of a process consuming a large percentage of available CPU cycles. It can prevent access to the switch by Telnet or make an ISL nonfunctional.

If you suspect high CPU use, complete the following steps.

- In privileged EXEC mode, enter the **show process cpu** command to determine which process is causing the high CPU reading.
- Shut down the corresponding interface or delete the configuration suspected of causing the high CPU use.

ECMP not load balancing as expected

Equal cost multipath (ECMP) routing increases throughput by balancing traffic across multiple routes that tie for best cost. If you suspect that traffic is not being balanced as expected, complete the following steps.

- In privileged EXEC mode, enter the **show fabric route topology** command to display whether ECMP routes are expected.

```
switch# show fabric route topology
Total Path Count: 1
Src  Dst  Out  Out          Nbr  Nbr
RB-ID RB-ID Index Interface    Hops Cost Index Interface  BW  Trunk
-----
66   1    124  Fi 66/0/4    1   500  129  Fi 1/-1/-1  32G Yes
```

If the output shows multiple equal-cost paths between the source and destination switches, then ECMP load balancing is expected.

- Check the interface utilization to verify whether it matches with the expected number of flows.
- Enter the **I2tracert** command to investigate whether Layer 2, Layer 3, and Layer 4 flows hash to separate ECMP links.

To avoid disruption of operation inherent in ECMP, the correctly functioning Brocade routing strategy routes a specific flow along one deterministic route. Additional flows take available equal-cost routes. This step verifies whether this flow hashing strategy is functioning correctly.

For details about using the **I2tracert** command, refer to [Using Layer 2 traceroute](#) on page 66.

ENS not working correctly

The Ethernet Name Server (ENS) is working correctly when the content of MAC address tables is the same among switches in the same VCS Fabric cluster. Perform the following checks to ensure that ENS is working correctly:

- Check that the fabric membership information is what you expect. Refer to [Verifying the fabric](#) on page 36.
- Ensure that MAC addresses are not moving among ports. Refer to [Checking for MAC address movement among ports](#) on page 36.
- Ensure that no edge port has an external loopback. Refer to [Verifying edge ports have no external loopback](#) on page 36.

Verifying the fabric

Enter the **show fabric all** command and ensure that information about all switches in the VCS Fabric cluster is displayed.

```
switch# show fabric all
VCS Id: 1
Config Mode: Local-Only
Rbridge-id WWN IP Address Name
-----
 1 50:00:51:E4:44:40:0E:04 0.0.0.0 "fcr_fd_1"
 2 50:00:51:E4:44:50:0F:09 0.0.0.0 "fcr_xd_2_128"
60 10:00:00:05:33:5F:EA:A4 10.24.81.65 "switch"
66 10:00:00:05:33:67:26:78 10.24.81.66 >"switch"
The Fabric has 4 Rbridge(s)
```

Checking for MAC address movement among ports

MAC address movement from port to port occurs when the same source address is detected on multiple ports. This condition is sometimes known as "MAC address flapping."

To check for MAC address flapping, enter the **show mac-address-table** command multiple times and check the output.

Verifying edge ports have no external loopback

Physically check for extended loopback.

Fabric does not form correctly

Some problems you might encounter when configuring zones include potential Fibre Channel router issues. For a more detailed discussion of possible Fibre Channel issues, refer to the *Fabric OS Troubleshooting and Diagnostics Guide*.

Some of the following problems may contribute to the zone not forming correctly:

- A Brocade VDX switch gets isolated when an RBridge ID matches the front domain ID or translate domain ID in a mixed network. Refer to [Recovering an isolated switch in a mixed FCoE fabric](#) on page 36.
- A "FID over-subscribed" message occurs during attempts to connect a backbone fabric to an edge fabric. Refer to [Recovering from FID oversubscription](#) on page 37.
- A "FID conflict" message occurs during attempts to connect a backbone fabric to an edge fabric. Refer to [Recovering from Fabric ID conflict](#) on page 37.
- Interfabric link (IFL) traffic does not flow over the intended link. Refer to [Rebalancing traffic over multiple IFLs](#) on page 37.
- Zone merge was expected to be blocked following reboot, but was not blocked. Refer to [Blocking zone merge after reboot](#) on page 38.
- Stale translate domains exist in an edge fabric. Refer to [Removing stale translate domains](#) on page 38.

Recovering an isolated switch in a mixed FCoE fabric

In an FCoE fabric that spans Network OS switches and Fabric OS switches, a Network OS switch with an RBridge ID that matches a front phantom domain ID or translate phantom domain ID of a connecting Fibre Channel router can become isolated.

FCoE connectivity across the Fibre Channel link between VCS Fabric clusters and Fibre Channel routers uses domain IDs to identify switches. Within a VCS Fabric cluster, a domain ID is the same as an RBridge ID. When you connect to a Fibre Channel router, the Fibre Channel router service in the Fibre Channel fabric emulates virtual phantom Fibre Channel domains in the FCoE fabric. Each Fibre Channel router-enabled switch emulates a single front phantom domain and each FC fabric is represented by a translate phantom domain.

To recover an isolated Network OS switch, complete the following steps.

1. Disable all FC routers that connect to the VCS Fabric cluster.
2. Reboot the isolated Network OS switch.
3. Re-enable all disabled FC routers.
4. To prevent switch isolation, follow these steps on each FC router that attaches to a VCS Fabric cluster.
 - a) Enter the **portCfgExPort -d** Fabric OS command to set a unique front phantom domain ID.
 - b) Enter the **fcrXlateConfig importedFID exportedFID preferredDomainID** command to set a unique translate phantom domain ID.

Refer to the *Fabric OS Command Reference* for details about the **portCfgExPort** and **fcrXlateConfig** commands.

Recovering from FID oversubscription

A "FID over-subscribed" message occurs when different Fibre Channel backbones attempt to connect to the same edge fabric using different Fabric IDs (FIDs). When you assign a FID to the edge fabric (**portCfgExPort -f** command), you must use the same FID as any other Fibre Channel backbone that connects to the edge fabric.

To resolve this problem, complete the following steps.

1. On the Fibre Channel router on the backbone with the errant FID configured, disable the EX_Port.
2. Enter the **portCfgExPort -f** command to configure the EX_Port with the same FID as the EX_Port on the other Fibre Channel router that connects to the same edge fabric.
3. Re-enable the EX_Port.

Refer to "Configuring an IFL for both edge and backbone connections" in the *Fabric OS Administrator's Guide* for details.

Recovering from Fabric ID conflict

The "FID conflict" message occurs when a backbone fabric connects to two or more edge fabrics that have the same Fabric ID (FID). Every edge fabric that a Fibre Channel router connects to must have a Fabric ID configured for the EX_Port that is unique on that Fibre Channel backbone. This error is most likely to occur when an edge fabric temporarily splits, causing it to appear as two edge fabrics with the same Fabric ID. This symptom might occur during VCS Fabric or Fibre Channel fabric upgrade, or as a result of a Brocade VDX or Fibre Channel switch reboot or crash.

Problem resolution depends on the cause of the problem. If the error is due to a temporary split, the problem will go away when the fabrics merge again.

If the problem is not due to a temporary fabric split, the most likely cause is misconfiguration. In this case, enter the **portCfgExPort -f** command to reconfigure one of the EX_Ports with a unique fabric ID.

Rebalancing traffic over multiple IFLs

If traffic across multiple interfabric links (IFLs) between a Fibre Channel router and an edge fabric is not balanced as you intended, it may be because the Fibre Channel router cannot determine an FSFP path from the Fibre Channel backbone to the target in the edge fabric. It uses all paths.

To direct the traffic the way you intend, on the FC router, use the **fcrRouterPortCost** command to configure a cost for each IFL. Traffic will flow across the lowest-cost IFL.

1. Connect to the FC router and log in using an account with admin permissions.
2. Disable the EX_Port.

3. Enter the **fcrRouterPortCost** command to configure the link cost.

Set the cost to 1000 if you want the link to carry traffic during normal operation. If you want the link to not carry traffic under normal operation, set the cost to 10000 (ten thousand) and set the cost of at least one other link to 1000. The default value is 1000, which you get when you enter a value of 0.

4. Re-enable the port.

For details about the **fcrRouterPortCost** command, refer to the *Fabric OS Command Reference*.

Blocking zone merge after reboot

To be sure of blocking zone merge following a switch reboot, enter the **no fabric isl enable** command to disable the ISL between neighboring Brocade VDX switches.



CAUTION

Brocade recommends that you do *not* use the shutdown command. If you use the shutdown command, then following switch reboot, the zone merge could happen before the shutdown command is replayed by the running configuration.

To block zone merge following reboot, follow these steps on each ISL port.

1. In global configuration mode, enter the **interface tengigabitethernet** (or **interface gigabitethernet**) command to enter interface configuration mode.
2. Enter the **no fabric isl enable** command.

Removing stale translate domains

A translate domain becomes stale when the edge fabric it represents becomes unreachable. By default, the stale translate domain is not deleted until the local edge fabric is rebuilt.

To delete a stale translate domain and avoid the disruption caused by rebuilding the local edge fabric, complete the following steps.

1. Connect to the Fibre Channel router and log in using an account with admin permissions.
2. On the FC router, enter the **fcrXlateConfig --show stalexd** command to list any stale translate domains.
3. Enter the **fcrXlateConfig --delete stalexd** command to delete the stale translate domain.

Refer to the *Fabric OS Command Reference* for details about the **fcrXlateConfig** command.

FCoE devices unable to log in

The inability to log in from a device connected through FCoE is usually because either the port or LLDP has been incorrectly configured. Potential reasons include:

- The default profile map has not been applied correctly. Refer to [Verifying the default profile map](#) on page 39.
- Required TLVs have not been advertised under LLDP. Refer to [Verifying TLVs](#) on page 40.

Verifying CNA login

If CNAs are not logging into the switch, perform the following procedure.

1. Check that the physical port is provisioned for FCOE.

```
switch# show fcoe interface ethernet | include "1/0/5"
TenGigaBitEthernet 1/0/5      default
```

- If the physical port is not provisioned, provision the interface for FCOE.
- If the CNA is still not logging in, check that the logical FCOE interface is online by using the **show running-config interface fcoe** command, as in the following example:

```
switch# show running-config interface fcoe
interface Fcoe 1/11/1
no shutdown
!
interface Fcoe 1/11/2
no shutdown
!
interface Fcoe 1/11/3
no shutdown
!
interface Fcoe 1/11/4
no shutdown
!
interface Fcoe 1/11/5
no shutdown
!
interface Fcoe 1/11/6
no shutdown
!
interface Fcoe 1/11/7
no shutdown
```

- Remove the FCOE provisioning and re-provision the physical interface.
- If that does not work, execute the **shut** command, and then the **no shut** command on the FCOE logical interface.
- If it still fails, collect the *supportSave* information and contact support. Refer also to [Gathering troubleshooting information](#) on page 19, which provides information about Network OS supportSave files.

Verifying the default profile map

- In privileged EXEC mode, enter **show running-config interface tengigabitethernet** followed by the interface ID to determine whether the default profile map has been applied to the interface.

```
switch# show running-config interface tengigabitethernet 5/0/1
interface TenGigabitEthernet 5/0/1
 fcoeport default
 shutdown
```

- If the default profile map has not been applied to the interface, or the initiator and target do not share the same VLAN ID, in interface configuration mode, enter **fcoeport default** to apply it.

```
switch(conf-if-te-0/1)# fcoeport default
```

This command not only applies the default profile map, but also associates the initiator and target with the same VLAN ID.

Verifying TLVs

The following TLVs —**dcbx-fcoe-app-tlv**, **dcbx-fcoe-logical-link-tlv**, and **dcbx-tlv**— must be advertised under LLDP or FCoE devices will not be able to log in.

1. In the privileged EXEC mode, enter the **show running-config protocol lldp** command to verify that the required TLVs are advertised.

```
switch# show running-config protocol lldp
protocol lldp
advertise dcbx-fcoe-app-tlv
advertise dcbx-fcoe-logical-link-tlv
advertise dcbx-tlv
```

2. If any of the required TLVs is missing, in protocol configuration mode, enter the corresponding **advertise** command.

```
switch# configure terminal
switch(config)# protocol lldp
switch(config-lldp)# advertise dcbx-fcoe-app-tlv
switch(config-lldp)# advertise dcbx-fcoe-logical-link-tlv
switch(config-lldp)# advertise dcbx-tlv
```

Heavy disk utilization

If FWDL and reboot taking longer than expected it may be due to heavy disk utilization.

It is recommended for the customers to check their disk utilization using the **dir** command and determine if their disk utilization exceeds 60% of total disk space.

If it exceeds 60% of the total disk space, Brocade recommends that you reduce the disk utilization by using the **capture copy-support** and **clear support** commands to delete the existing core files.

If the disk utilization does not reduce to less than 60%, prepare a case for your switch provider, as described in [Contacting Brocade Technical Support](#) on page 10.

ISL does not come up on some ports

The failure of an Inter-SwitchLink (ISL) between two switches in a VCS Fabric cluster can occur for various reasons:

- The ISL configuration is segmented or disabled. Refer to [Verifying the status of ISLs](#) on page 41.
- There are duplicate switch IDs or inconsistency in specifying the VCS ID. Refer to [Verifying VCS ID and RBridge IDs](#) on page 42.
- LLDP is not reporting its neighbors. Refer to [Verifying LLDP](#) on page 43.
- An overloaded CPU fails to generate keepalive packets. Refer to [Checking for CPU overload](#) on page 43.

Verifying the status of ISLs

If any port looks suspicious, begin by checking the status of ISLs.

1. On the switches at each end of the broken link, in privileged EXEC mode, enter the **show fabric isl** command to view the status of ISL connections. Here we look at switch1.

```
switch1# show fabric isl
Rbridge-id: 2 #ISLs: 2
Src      Src      Nbr      Nbr
Index Interface Index Interface Nbr-WWN          BW      Trunk  Nbr-Name
-----
1      Te 2/0/1  1      Te 3/0/1  10:00:00:05:1E:CD:7A:7A 10G     Yes   "switch1"
2      Te 2/0/2  ?      Te ?/?/?  ??:?:?:?:?:?:?:?:?:?? (segmented - incompatible)
26     Te 2/0/26 56     Te 25/0/56 10:00:00:05:33:40:2F:C9 60G     Yes   "Edget12r31_25"
34     Te 2/0/34 58     Te 26/0/58 10:00:00:05:33:41:1E:B7 40G     Yes   "Edget12r32_26"
```

Ports on which the ISL link is broken appear with the text "(segmented - incompatible)." Ports for which the ISL configuration is disabled do not appear in the output.

2. To examine the peer, switch2, enter the **show fabric islports** command to gather more information about the status of suspect ports.

```
switch2# show fabric islports
Name:      switch2
Type:      107.4
State:     Online
Role:      Fabric Subordinate
VCS Id:    10
Config Mode: Local-Only
Rbridge-id: 11
WWN:      10:00:00:05:33:6d:7f:77
FCF MAC:   00:05:33:6d:7f:77
Index Interface State Operational State
-----
1      Te 11/0/1  Up    ISL 10:00:00:05:33:00:77:80 "switch" (upstream) (Trunk Primary)
2      Te 11/0/2  Down
3      Te 11/0/3  Down
4      Te 11/0/4  Up    ISL (Trunk port, Primary is Te 11/0/1)
5      Te 11/0/5  Down
6      Te 11/0/6  Down
7      Te 11/0/7  Down
8      Te 11/0/8  Down
9      Te 11/0/9  Down
10     Te 11/0/10 Down
11     Te 11/0/11 Up    ISL 10:00:00:05:1e:00:50:00 "switch" (Trunk Primary)
121    Fi 11/0/1  Up    LS ISL 50:00:53:37:b6:93:5e:02 "fcr_fd_160" (downstream) (Trunk Primary)
122    Fi 11/0/2  Up    LS ISL (Trunk port, Primary is Fi 11/0/1)
123    Fi 11/0/3  Down
124    Fi 11/0/4  Down
125    Fi 11/0/5  Down
126    Fi 11/0/6  Down
127    Fi 11/0/7  Down
```

3. If the state of the corresponding port is "Down" (your network will vary), enable the port with the **no shutdown** command.

```
switch2# configure terminal
Entering configuration mode terminal

switch(config)# interface tengigabitethernet 11/0/9
switch(conf-if-te-11/0/9)# no shutdown
```

4. If the port state is "Up", but the ISL is segmented (see Step 1), examine the Operational State string for further clues to the reason for the segmentation.

Refer to the *Network OS Command Reference* for details about the **show fabric islports** command and help in interpreting the Operational State string for a segmented ISL.

Verifying VCS ID and RBridge IDs

For the ISL to function correctly, the following criteria must be true:

- Both switches must have the same VCS ID.
- Each switch must have a unique RBridge ID.

To check the criteria, complete the following steps.

1. Enter the **show vcs** command on each switch.
2. Depending on the output, proceed as follows:
 - If the **show vcs** command indicates that the VCS ID is not the same on each switch, enter the **vcs vcsid** command to correct the VCS ID on the switch that is in error.

```
switch1# show vcs
Config Mode : Local-Only
VCS ID      : 1
Total Number of Nodes      : 1
Rbridge-Id   WWN           Management IP   VCS Status   Fabric Status
-----
66           >10:00:00:05:33:67:26:78* 10.24.81.66   Online       Online
cz41-h06-m-r2
```

```
switch2# show vcs
Config Mode : Local-Only
VCS ID      : 2
Total Number of Nodes      : 1
Rbridge-Id   WWN           Management IP   VCS Status   Fabric Status
-----
66           >10:00:00:05:33:67:26:78* 10.24.81.77   Online       Online
cz41-h06-m-r2
```

```
switch2# vcs vcsid 1
```

- If both switches have the same RBridge ID, enter the **vcs rbridge-id** command to change the RBridge ID to a unique value.

```
switch1# show vcs
Config Mode : Local-Only
VCS ID      : 1
Total Number of Nodes      : 1
Rbridge-Id   WWN           Management IP   VCS Status   Fabric Status
-----
66           >10:00:00:05:33:67:26:78* 10.24.81.66   Online       Online
cz41-h06-m-r2
```

```
switch2# show vcs
Config Mode : Local-Only
VCS ID      : 1
Total Number of Nodes      : 1
Rbridge-Id   WWN           Management IP   VCS Status   Fabric Status
-----
66           >10:00:00:05:33:67:26:78* 10.24.81.66   Online       Online
cz41-h06-m-r2
```

```
switch2# vcs rbridge-id 77
```

Verifying LLDP

When ISLs are functioning correctly, the **show lldp neighbors** command reports on each neighbor switch in the VCS Fabric cluster.

1. Enter the **show lldp neighbors** command to verify that LLDP reports on all of its neighbors.

```
switch1# show lldp neighbors
Local Intf   Dead Interval   Remaining Life   Remote Intf   Chassis ID       Tx           Rx
Te 66/0/55   120             106             port1         0005.1e78.f004   20300       19914
Te 66/0/60   120             108             port0         0005.1e55.16c8   20300       19911
```

2. If neighbors are missing, you will need to perform further debugging or contact your switch support provider.

Checking for CPU overload

An abnormally high CPU load can cause an ISL to malfunction. Use the **show process cpu** command as described in [CPU use is unexpectedly high](#) on page 35 to troubleshoot an overloaded CPU.

License is not properly installed

If a licensed feature is not functioning, a probable reason is that the license for that feature has not been installed correctly. Either the license was not installed, or it was installed and a required system reboot was not performed.

If you are unable to connect an FCoE device or unable to use Fibre Channel ports on a Brocade VDX switch, it is likely that the FCoE license is not installed.

If you suspect a license is not properly installed, complete the following steps.

1. In privileged EXEC mode, enter the **show license** command to display the currently installed licenses.

```
switch# show license
rbridge-id: 66
xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
FCoE Base license
Feature name:FCOE_BASE
License is valid
```

2. If the FCoE or DPOD license appears in the **show license** command output, but the feature does not work for the expected ports, the probable cause is that the affected ports were not re-enabled after installing the license.

NOTE

After adding an FCoE or DPOD license, you must disable and re-enable all affected ports.

You can disable and then enable each affected port, or you can enter the **chassis disable** command followed by the **chassis enable** command to re-enable the entire chassis.

```
switch# chassis disable
switch# chassis enable
```

- If the license does not appear in the **show license** command output, then it was not installed. In privileged EXEC mode, enter the **license add licstr** command to install the license. For FCoE and DPOD licenses, you must also disable and enable the switch or port.

```
switch# license add licstr "*B
s1SETgzTgeVGUDeQR4WIFRx7mmXODdSwENoRGENAmX3Ca3uHeZgXK0b, jzxyzfzKLrMsPN8C1SxvDQRRT8VyuULyyKTO0ryU6qm4s
1jjiSAeV, COoedzCx1v6ycQgnYMeSVp#"
License Added [*B
s1SETgzTgeVGUDeQR4WIFRx7mmXODdSwENoRGENAmX3Ca3uHeZgXK0b, jzxyzfzKLrMsPN8C1SxvDQRRT8VyuULyyKTO0ryU6qm4s
1jjiSAeV, COoedzCx1v6ycQgnYMeSVp# ]
For license change to take effect, please disable/enable port or switch...
switch# chassis disable
switch# chassis enable
```

Need to recover password for Brocade VDX switches

Use these procedures to recover access to your switch when normal access to the admin account has been lost.



CAUTION

Because of the complexity of these procedures, we highly recommend that you contact support for guidance, especially for recovering the root password. The recovery steps must be followed exactly as presented below. Any variation in the procedure might cause unpredictable results.

There are several methods for recovering passwords on a Brocade Network OS switch. The correct approach depends on the accounts to which you have access. The table below lists the procedures and conditions under which you would use each procedure to recover passwords. These procedures apply to all versions of Network OS firmware across all Network OS platforms, except where noted.

Account access availability	Use these procedures
<ul style="list-style-type: none"> Access to admin account 	Use normal password management procedures.
<ul style="list-style-type: none"> No access to admin account Access to root account 	Recovering the admin password by using the root account on page 49
<ul style="list-style-type: none"> No access to admin account No access to root account Access to boot PROM interface 	Recovering the root password for Brocade VDX switches on page 45
<ul style="list-style-type: none"> No access to admin account No access to root account No access to boot PROM interface 	Obtaining the Boot PROM recovery password on page 50 and Recovering the root password for Brocade VDX switches on page 45

If you still have access to the admin account, you can change the admin account password or change passwords on user accounts by using normal password-management procedures. Refer to the “Managing User Accounts” section of the *Network OS Security Configuration Guide*.

Even if you have lost access to the admin account but you do have access to the root account, you can use the root account to reset passwords for the root, admin, and user accounts. Refer to [Recovering the admin password by using the root account](#) on page 49. Once you have reset the admin account password, you can use that account to set user login passwords.

If you do not have access to the root account, you can use the Boot PROM method. Refer to [Recovering the root password for Brocade VDX switches](#) on page 45 in this section. If the password is set on the boot PROM and is unknown, contact your switch service provider for a boot PROM recovery string to regain access to the switch. Refer to [Obtaining the Boot PROM recovery password](#) on page 50 in this section.

Try the factory default passwords before proceeding in case any are still in effect. The following table lists the default passwords for Network OS switches.

Account	Default password
root	fibranne
admin	password
user	password

NOTE

When connected through a serial cable to the console, always save the output by using the capture functionality under Windows, or the script functionality for UNIX or Linux.

Recovering the root password for Brocade VDX switches

This procedure must be performed from the serial interface console.



CAUTION

Enter commands at the Boot PROM interface exactly as shown. Commands entered incorrectly at the Boot PROM interface can render your switch unstable or unusable. To recover, you would need to seek help from your switch service provider or return your switch to the factory for repair.

You can use this procedure if you need to recover passwords on a device for which the root account is not accessible. If the root account is accessible, use [Recovering the admin password by using the root account](#) on page 49 instead.

To use this procedure, you must have access to the Boot PROM interface; that is, its password must be available or not set. If you do not have access to the Boot PROM interface, use [Obtaining the Boot PROM recovery password](#) on page 50 before using this procedure.

This section provides detailed procedures for performing password recovery in addition to a quick reference for advanced users who need only a reminder of the basic steps:

- [Recovering the root password for Brocade VDX switches: Quick reference](#) on page 48
- [Recovering the root password for Brocade VDX switches: Detailed procedure](#) on page 45

Recovering the root password for Brocade VDX switches: Detailed procedure

Use this procedure if you do not have access to the root account.

To reset the root password to its factory default value on a Brocade VDX switch, set a password for the admin account, and then restore nondefault user accounts, follow these steps:

You may need to disable the root account using the **no root enable** command.

1. Connect to the serial console port of the switch.
2. Manually reboot the switch.

- When prompted to stop test or stop AutoBoot, press **ESC**.

NOTE

If the **ESC** key is not effective during reboot, turn the power off and back on, and then try again. If the **ESC** key is still not effective, check the serial console cable. If the cable is connected correctly, then the unit must be returned for service or repair.

The Boot PROM menu is displayed with the following options:

Start system	Used to reboot the system.
Recover password	Used to generate a character string for your support provider to recover the Boot PROM password. ATTENTION Use this feature only when directed by technical support personnel.
Enter command shell	Used to enter the command shell to reset all passwords on the system.

```
Checking system RAM - press any key to stop test
Checking memory address: 00100000
System RAM test terminated by keyboard
set_bootstatus: BS_LOAD_OS, platform_idx = 6
Hit ESC to stop autoboot: 0
```

- Start system.
- Recover password.
- Enter command shell.

Option?

- Enter 3 at the prompt to open the command shell.
- If prompted, enter the Boot PROM password and press **Enter**.

NOTE

The Boot PROM has a password only if one has been defined. If you are prompted to enter a new Boot PROM password, make sure that it is at least 8 characters long. Do not select this option unless specifically instructed to do so by support personnel.

- Append "S" to the boot arguments so that the switch boots into single-user mode, by entering the following at the prompt:
=> setenv bootargs "root=/dev/sda1 rootfstype=ext4 quiet S"
- Enter the **printenv** command to verify the change.

```
=> printenv
AutoLoad=yes
LoadIdentifiers=Fabric Operating System;Fabric Operating System
OSLoadOptions=quiet
OSRootPartition=sda2;sda1
SkipWatchdog=yes
autoreset_mac=true
baudrate=9600
bootargs=root=/dev/sda1 rootfstype=ext4 quiet S
bootcmd=execute_internal_bootcmd
(output truncated)
```

- Enter the **saveenv** command to save the changes.

```
=> saveenv
Saving Environment to Flash.....Done
```

9. Enter the **reset** command to bring up the device in single-user mode

```
=> reset
BootROM version: 1.0.48
Copyright (C) 2011 Brocade Communication.

CPU0: P4080E, Version: 2.0, (0x82080020)
(output truncated)
```

10. Enter the **mount** command with the following parameters to remount the root partition as read/write capable.

```
sh-2.04# mount -vo remount,rw,noatime /
/dev/root on / type ext4 (rw,noatime)
```

11. Mount the secondary partition.

Examine the output of the **printenv** command in Step 7, to check which partition the root points to in the boot arguments (bootargs = root setting). If the root partition is sda2, then use **sda1** in this command. If the root partition is sda1, then use **sda2**.

```
sh-2.04# mount /dev/sda2 /mnt
```

12. Enter the **passwddefault** command to reset the root password to the factory default value or the **/sbin/passwddefault -f** (to reset the root password and forcefully enable the root account).

```
sh-2.04# /sbin/passwddefault -f
```

NOTE

For Network OS, the **passwddefault -f** command restores the passwords of factory default accounts to their default values, removes non-default user accounts that are present, and enables the root account (if it is disabled). Error messages seen during the execution of that command (applicable to Network OS 3.0.0) should be ignored.

In a dual management-module (MM) chassis, enter the **passwddefault** command on the standby MM for password recovery.

13. Reset the boot arguments by removing the "S".

```
sh-2.04# bootenv bootargs "root=/dev/sda1 rootfstype=ext4 quiet"
```

14. Reboot the switch by using the **partman -r** command.

```
sh-2.04# partman -r
```

15. Log in to the switch by using the serial interface or Telnet. Use the factory default accounts (root/admin/user).

16. Start the Network OS command line.

```
switch:root> noscli

SECURITY WARNING: The default password for at least
one default account (root, admin and user) have not been changed.

Welcome to the Brocade Network Operating System Software
admin connected from 127.0.0.1 using console on switch
```

17. Enter global configuration mode.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)#
```

18. Use the following syntax of the **username** command to reset passwords for the admin or user accounts, or for any other nondefault users.

```
username account-name password new-password
```

The following example resets the admin password to the default value of "password."

```
switch(config)# username admin password password
```

19. To restore the non-default user accounts, perform the following steps:

- a) Copy the running-config to a file.

```
switch: copy running-config flash://running-config.cfg  
2012/07/09-11:51:21, [DCM-1108], 4930, M2, INFO, VDX8770-4, Running  
configuration file has been uploaded successfully to the remote location.
```

- b) Copy the default-config to the startup-config, to reset the startup-config.

```
switch# copy default-config startup-config
```

- c) Reboot the switch.

```
switch# reload  
Warning: Unsaved configuration will be lost. Please run 'copy  
running-config startup-config' to save the  
current configuration if not done already.  
  
Are you sure you want to reload the switch? [y/n]:y  
The system is going down for reload NOW !!
```

- d) Copy the file saved in Step 19a to the running-config.

```
switch# copy flash://running-config.cfg running-config  
Loading.  
2012/07/09-12:08:13, [DCM-1105], 5456, M2, INFO, VDX8770-4, Copy of the  
downloaded config file to the current running-config has completed  
successfully on this node.
```

- e) Copy the running-config to the startup-config.

```
switch# copy running-config startup-config
```

The password recovery procedure is now complete. You can now use normal password-management procedures from the admin account.

Recovering the root password for Brocade VDX switches: Quick reference

Advanced users who need only a reminder of the basic steps can use this quick reference to recover passwords.

You may need to disable the root account using the **no root enable** command.

1. Press **ESC** during reboot.

NOTE

If the **ESC** key is not effective during reboot, turn the power off and back on, and then try again. If the **ESC** key is still not effective, check the serial console cable. If the cable is not connected correctly, then the unit must be returned for service or repair.

2. Choose option **3**.

3. Enter the following commands in sequence:
 - a) **setenv bootargs "root=/dev/sda1 rootfstype=ext4 quiet S"**
 - b) **saveenv**
 - c) **reset**
 - d) **mount -vo remount,rw,noatime /**
 - e) **mount /dev/sda1 /mnt**
For Step 3e, choose the root partition that was not set as root in Step 3a.
 - f) **/sbin/passwddefault** (to recover the default accounts password) or **/sbin/passwddefault -f** (to recover the default accounts password and forcefully enable the root account)
 - g) **bootenv bootargs "root=/dev/sda1 rootfstype=ext4 quiet"**
 - h) **partman -r**
4. Log in as root and enter the following commands in sequence:
 - a) **noscli**
 - b) **configure**
 - c) **username name password new-password**
5. Restore nondefault user accounts.

Recovering the admin password by using the root account

Use this procedure if you have lost access to the admin account, but you do have access to the root account.

To reset any account password from the root account, follow these steps:

NOTE

For a non-secured system, you can use the serial interface or Telnet. For a secure system, you can use the serial interface or secure Telnet.

1. Open a CLI session to the switch.
2. Log in as root.
3. At the prompt, enter the **noscli** command to start the Network OS command line.

```
switch:root> noscli
```

4. Enter global configuration mode.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)#
```

5. Use the following syntax of the **username** command to reset passwords for the admin or user accounts, or for any other nondefault users.

```
username account-name password new-password
```

The following example resets the admin password to the default value of "password."

```
switch(config)# username admin password password
```

You can now use the admin account to manage the admin and user passwords by using normal password-management procedures.



IMPORTANT

Keep a hard copy of your switch passwords in a secure location.

Obtaining the Boot PROM recovery password

Use this procedure when you do not have the Boot PROM password.

This procedure obtains a Boot PROM recovery password. It does not reset the Network OS passwords on the switch. Once the Boot PROM password has been recovered, you must go through the Boot PROM command shell to reset the Network OS passwords on the switch.

This procedure explains how to gather the information you need to send to your switch support provider in order to get a Boot PROM recovery password. Once you have received the Boot PROM recovery password, and gained access to the Boot PROM, you must reset the passwords by using [Recovering the root password for Brocade VDX switches](#) on page 45.

After completing this procedure, the Boot PROM password will be set. To avoid having to obtain a Boot PROM recovery password for future password-recovery operations, you can choose to reset the Boot PROM password as described in [Clearing the Boot PROM password](#) on page 34.

To obtain the Boot PROM recovery password from your switch support provider, perform the following steps:

1. Connect to the serial console port of the switch.
2. Manually reboot the switch.
3. When prompted to stop test or stop AutoBoot, press **ESC**.

NOTE

If the **ESC** key is not effective during reboot, turn the power off and back on, and then try again. If the **ESC** key is still not effective, check the serial console cable. If the cable is connected correctly, then the unit must be returned for service or repair.

The Boot PROM menu is displayed with the following options:

Start system.	Used to reboot the system.
Recover password.	Used to generate a character string for your support provider to recover the Boot PROM password. ATTENTION Use this feature only when directed by technical support personnel.
Enter command shell.	Used to enter the command shell to reset all passwords on the system.

```
Checking system RAM - press any key to stop test
Checking memory address: 00100000
System RAM test terminated by keyboard
set_bootstatus: BS_LOAD_OS, platform_idx = 6
Hit ESC to stop autoboot: 0
```

- 1) Start system.
- 2) Recover password.
- 3) Enter command shell.

Option?

4. Enter 2 at the prompt. A character string is displayed, highlighted in the following example.

```

1) Start system.
2) Recover password.
3) Enter command shell.

Option? 2
Send the following string to Customer Support for password recovery:
/uaSLR1raCqT3FToqy0ZjA==

```

5. Send the character string to your switch support provider to obtain a Boot PROM recovery password.



CAUTION

Do not reboot the switch at this point. Doing so will cause the password recovery string to change.

6. As prompted, perform the appropriate steps to set the Boot PROM password if it was not set.

```
Recovery password is NOT set. Please set it now.
```

7. When prompted, enter the Recovery Password that is generated from your support provider, and reenter it when prompted.

```

Enter the supplied recovery password.
Recovery Password: YnfG9DDr1FMDVknM0RkPtg== < Supplied by your support provider
Re-enter Recovery Password: YnfG9DDr1FMDVknM0RkPtg==

```

8. When prompted with "New password:", enter a new Boot PROM password, and reenter it when prompted.

```

New password: xxx
Re-enter new password: xxx

```

The switch reboots.

ATTENTION

Record the new password for future reference.

You are now ready to record passwords as described in [Recovering the root password for Brocade VDX switches](#) on page 45.

Packets are dropped in hardware

This section discusses how to troubleshoot problems in which loss of packets occurs in all traffic, on specific traffic flows, in specific types of traffic, consistently, or intermittently. Dropped packets could occur for many reasons, including the following:

- High latency in an end device. Refer to [Verifying packets dropped because of high-latency end device](#) on page 51.
- Broken data path. Refer to [Verifying the data path](#) on page 54.
- Noise on an optical line caused by too many CRC errors, packet errors, or NIC interoperability errors. Refer to [Checking for noise on an optical line](#) on page 56.

Verifying packets dropped because of high-latency end device

Packets can sometimes be dropped because of buffer overrun within the fabric caused by end devices taking longer to respond than expected. For example, an overloaded disk array can cause such latency, as can a host that does not process data as quickly as expected. Devices that stop receiving data for an extended period of time can cause excessive latency.

The ultimate solution to these problems is to fix the end device itself. However, some adjustments to the switch and fabric configuration can help to reduce the problem.

To detect and relieve congestion and dropped packets resulting from latency in end devices, complete the following steps:

1. Enter the **show lldp neighbors detail** command to check under "DCBX TLVs" that the end device is DCB-ready and confirm that the end device is also advertising its DCB capabilities.

```
switch# show lldp neighbors detail
Neighbors for Interface Te 66/0/55
MANDATORY TLVs
=====
Local Interface: Te 66/0/55 (Local Interface MAC: 0005.3367.26d3)
Remote Interface: port1 (Remote Interface MAC: 0005.1e78.f004)
Dead Interval: 120 secs
Remaining Life : 104 secs
Chassis ID: 0005.1e78.f004
LLDP PDU Transmitted: 2412 Received: 2372
OPTIONAL TLVs
=====
DCBX TLVs
=====
Version : CEE
DCBX Ctrl OperVersion: 0 MaxVersion: 0 SeqNo: 1 AckNo: 4
DCBX ETS OperVersion: 0 MaxVersion: 0 Enabled: 1 Willing: 1 Error: 0
Enhanced Transmission Selection (ETS)
  Priority-Group ID Map:
    Priority : 0 1 2 3 4 5 6 7
    Group ID : 2 2 2 1 2 2 2 15
  Group ID Bandwidth Map:
    Group ID : 0 1 2 3 4 5 6 7
    Percentage: 0 40 60 0 0 0 0 0
  Number of Traffic Classes supported: 8
DCBX PFC OperVersion: 0 MaxVersion: 0 Enabled: 1 Willing: 1 Error: 0
Priority-based Flow Control (PFC)
  Enabled Priorities: 3
  Number of Traffic Class PFC supported: 8
FCoE App OperVersion: 0 MaxVersion: 0 Enabled: 1 Willing: 1 Error: 0
FCoE Application Protocol
  User Priorities: 3
```

2. Enter the **show qos flowcontrol interface** command to check for pause frames.

```
switch# show qos flowcontrol interface tengigabitethernet 66/0/55
Interface TenGigabitEthernet 66/0/55
Mode PFC
DCBX enabled for PFC negotiation
TX 4926331124 frames
  TX      TX      RX      RX      Output Paused
CoS Admin Oper  Admin Oper   512    BitTimes
-----
0 Off    Off    Off    Off    0
1 Off    Off    Off    Off    0
2 Off    Off    Off    Off    0
3 On     On     On     On     0
4 Off    Off    Off    Off    0
5 Off    Off    Off    Off    0
6 Off    Off    Off    Off    0
7 Off    Off    Off    Off    0
```

3. Enter the **show qos queue interface** command to check the CoS statistics.

```
switch# show qos queue interface tengigabitethernet 66/0/60
Interface TenGigabitEthernet 66/0/60
      RX      RX      TX      TX
CoS   Packets Bytes TC   Packets Bytes
-----
0     1600   354184  0     0       0
1       0     0       1    7962   636960
2       0     0       2     0       0
3     8508   544832  3     18     6048
4       0     0       4     0       0
5       0     0       5     0       0
6       0     0       6     0       0
7       0     0       7    2123   282360
untag 2082   216528
```

4. Enter the **show qos rcv-queue interface** command to check for indicators of congestion, including dropped packets, buffer consumption, and real-time queue statistics.

```
switch# show qos rcv-queue interface tengigabitethernet 66/0/55
Interface TenGigabitEthernet TenGigabitEthernet 66/0/55
In-use 27216 bytes, Total buffer 144144 bytes
0 packets dropped
      In-use   Max
TC    Bytes   Bytes
-----
0     0       252
1     0       252
2     0       252
3  27216   75284
4     0       252
5     0       252
6     0     57456
7     0     9576
```

5. Enter the **show qos interface** command to check the QoS configuration.

```
switch# show qos interface tengigabitethernet 66/0/55
Interface TenGigabitEthernet 66/0/55
Provisioning mode cee
Priority Tag disable
CEE Map default
FCoE Provisioned
Default CoS 0
Interface trust cos
      In-CoS: 0  1  2  3  4  5  6  7
-----
Out-CoS/TrafficClass: 0/6 1/6 2/6 3/3 4/6 5/6 6/6 0/7
Per-Traffic Class Tail Drop Threshold (bytes)
TC: 0  1  2  3  4  5  6  7
-----
Threshold: 252 252 252 75284 252 252 57456 9576
Flow control mode PFC
CoS3 TX on, RX on
Multicast Packet Expansion Rate Limit 3000000 pkt/s, max burst 4096 pkts
Multicast Packet Expansion Tail Drop Threshold (packets)
TrafficClass: 0  1  2  3  4  5  6  7
-----
Threshold: 64 64 64 64 64 64 64 64
Traffic Class Scheduler configured for 1 Strict Priority queues
TrafficClass: 0  1  2  3  4  5  6  7
-----
DWRRWeight: 0  0  0  40  0  0  60  --
Multicast Packet Expansion Traffic Class Scheduler
TrafficClass: 0  1  2  3  4  5  6  7
-----
DWRRWeight: 12 13 12 13 12 13 12 13
```

6. Reconfigure QoS. Refer to the “Configuring QoS” section of the *Network OS Layer 2 Switching Configuration Guide* for detailed information.

Verifying the data path

This procedure checks whether fabric continuity might be the reason for dropped packets.

NOTE

The E1MG-SX-OM and E1MG-LX-OM modules are not supported by Network OS. Despite being Brocade products, these modules will return the error ‘Optic is not Brocade qualified, optical monitoring is not supported’ and must be replaced with a supported module.

1. Enter the **ping** command to test for a complete path to the end device

```
switch# ping dest-address 10.24.81.2
PING 10.24.81.2 (10.24.81.2): 56 octets data
64 octets from 10.24.81.2: icmp_seq=0 ttl=128 time=9.4 ms
64 octets from 10.24.81.2: icmp_seq=1 ttl=128 time=0.3 ms
64 octets from 10.24.81.2: icmp_seq=2 ttl=128 time=0.3 ms
64 octets from 10.24.81.2: icmp_seq=3 ttl=128 time=0.3 ms
64 octets from 10.24.81.2: icmp_seq=4 ttl=128 time=0.3 ms
--- 10.24.81.2 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0.3/2.1/9.4 ms
```

2. Enter the **show interface** command to display whether packets are coming in or are dropped as errors. Specifically, examine the output fields shown underlined in the following example.

```
switch# show interface tengigabitethernet 66/0/60
TenGigabitEthernet 66/0/60 is up, line protocol is up (connected)
Hardware is Ethernet, address is 0005.3367.26d8
  Current address is 0005.3367.26d8
Pluggable media present
Interface index (ifindex) is 283874428169
MTU 2500 bytes
LineSpeed Actual    : 10000 Mbit
LineSpeed Configured : Auto, Duplex: Full
Flowcontrol rx: off, tx: off
Last clearing of show interface counters: 22:07:59
Queueing strategy: fifo
Receive Statistics:
  15254 packets, 1395269 bytes
  Unicasts: 10641, Multicasts: 2637, Broadcasts: 1976
  64-byte pkts: 10874, Over 64-byte pkts: 3294, Over 127-byte pkts: 117
  Over 255-byte pkts: 969, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
  Over 1518-byte pkts(Jumbo): 0
  Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
  Errors: 0, Discards: 0
Transmit Statistics:
  12633 packets, 1155963 bytes
  Unicasts: 18, Multicasts: 12615, Broadcasts: 0
  Underruns: 0
  Errors: 0, Discards: 0
Rate info:
  Input 0.000128 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 1d00h40m
```

- Enter the **show media interface** command to check that the optics used are Brocade-certified. Verify the Vendor Name field, shown underlined in the following example, reads BROCADE. If the Vendor Name field shows anything other than BROCADE, replace the optics with Brocade-certified optics.

Check also the TX Power and RX Power fields to ensure they are not zero.

```
switch# show media interface tengigabitethernet 66/0/60
Interface      TenGigabitEthernet 66/0/60
Identifier     13   QSFP
Connector     7    LC
Transceiver   0000000000000002 40_GB/s Long_dist
Name          lw
Encoding      5    IEEE 802.3ab
Baud Rate     103 (units 100 megabaud)
Length 9u     10   (units km)
Length E-50u  0    (units 2 meters)
Length 50u    0    (units 1 meters)
Length 62.5u  0    (units 1 meters)
Length Cu     0    (units 1 meter)
Vendor Name   BROCADE
Vendor OUI    00:05:1e
Vendor PN     57-1000263-01
Vendor Rev    A
Wavelength   26020(units 0.05 nm)
Options       0001
BR Max       12
BR Min       216
Serial No    LDF113390001CBS
Date Code    130928
Optical Monitor yes
Temperature   35 Centigrade
Voltage       3304.6 (mVolts)
Current       38.544 (mAmps)
TX Power      N/A
RX Power      2.7 (uWatts)
```

- Enter the **show mac-address-table** command to verify that the MAC address table learns new values.

The new MAC address should appear here.

```
switch# show mac-address-table
VlanId  Mac-address  Type  State  Ports
1002    0efc.0042.7300  FPMA  Active  Te 66/0/55
1002    0efc.0042.7302  FPMA  Active  Te 66/0/55
1002    0efc.0042.7800  FPMA  Active  Te 66/0/60
Total MAC addresses : 3
```

- Enter the **show lldp neighbors** command to verify that LLDP reports all neighbors.

```
switch# show lldp neighbors
Local Intf  Dead Interval  Remaining Life  Remote Intf  Chassis ID  Tx  Rx
Te 66/0/55  120           101             port1        0005.1e78.f004  3000 2948
Te 66/0/60  120           117             port0        0005.1e55.16c8  2999 2945
```

If the output does not show all neighbors, contact your switch support provider.

- Enter the **show mac-address-table** command to verify the Ethernet Name Service functionality and to detect whether MAC addresses learned from other VCS Fabric switches are present.

Enter this command on other switches in the fabric to ensure that those switches can detect this MAC address.

```
switch# show mac-address-table
VlanId  Mac-address  Type  State  Ports
1002    0efc.0042.7300  FPMA  Active  Te 66/0/55
1002    0efc.0042.7302  FPMA  Active  Te 66/0/55
1002    0efc.0042.7800  FPMA  Active  Te 66/0/60
Total MAC addresses : 3
```

7. Enter the **l2tracert** command to validate the data-path fabric continuity.
 - Enter dynamically learned source MAC address and destination MAC address for the data path.
 - Among the extended commands, use IP, SIP, DIP, TCP, Src Port, and Dest Port commands.
 - Enter the IP command parameters to ensure that the traceroute packet traverses a specific ECMP link.

For details on using the **l2tracert** command, refer to [Using Layer 2 traceroute](#) on page 66.

Checking for noise on an optical line

Excessive noise on an optical line can result in dropped packets because of excessive CRC errors, NIC interoperability errors, or other conditions.

1. Enter the **show interface** command and check the output for CRC errors or TX discards; examine the fields shown underlined in the following example.

```
switch# show interface tengigabitethernet 66/0/55
TenGigabitEthernet 66/0/55 is up, line protocol is up (connected)
Hardware is Ethernet, address is 0005.3367.26d3
  Current address is 0005.3367.26d3
Pluggable media present
Interface index (ifindex) is 283874100484
MTU 2500 bytes
LineSpeed Actual      : 10000 Mbit
LineSpeed Configured : Auto, Duplex: Full
Flowcontrol rx: off, tx: off
Last clearing of show interface counters: 21:51:35
Queueing strategy: fifo
Receive Statistics:
  15433457505 packets, 32164575799774 bytes
  Unicasts: 15433454934, Multicasts: 2571, Broadcasts: 0
  64-byte pkts: 11357, Over 64-byte pkts: 242664576, Over 127-byte pkts: 0
  Over 255-byte pkts: 0, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
  Over 1518-byte pkts(Jumbo): 15190781568
  Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
  Errors: 0, Discards: 0
Transmit Statistics:
  21456965161 packets, 32549136821934 bytes
  Unicasts: 15313174675, Multicasts: 6143790486, Broadcasts: 0
  Underruns: 0
  Errors: 0, Discards: 0
Rate info:
  Input 3345.136864 Mbits/sec, 200572 packets/sec, 33.45% of line-rate
  Output 3386.493904 Mbits/sec, 281345 packets/sec, 33.86% of line-rate
Time since last interface status change: 1d00h24m
```

2. If errors are reported in the previous step, check the SFP transceiver and cable on the local switch and on the peer switch at the other end of the cable.
 - a) Enter the **show media interface** command on each switch and check the Vendor Name field to check that the optics are Brocade-certified. If the Vendor Name field shows anything other than BROCADE, replace the optics with Brocade-certified optics.

Replace any non-Brocade SFP transceiver.
 - b) Try replacing the SFP transceiver.
 - c) Try replacing the cable.

Ping fails

If pings do not successfully traverse the switch, try the following operations.

1. Trace the packet flow and check whether ARP or ICMP packets are getting dropped.
2. Trace which direction is failing by using interface statistics.
3. Locate the device that is dropping the packets.
4. Look for any error counters incrementing on that device.
5. Check the MAC address table to determine whether the MAC addresses are learned on the correct port or port-channel.

QoS configuration causes tail drops

Tail-drop queueing is the most basic form of congestion control. Normal operation is first-in, first-out (FIFO) until all buffers are exhausted. After that, new frames are dropped.

For the Brocade VDX 6740 series and the Brocade VDX 6940 series, you can check for ingress tail-drop at run time by issuing the **show qos rcv-queue interface <port>** command. You can check for egress tail-drops by running the **show qos tx-queue interface <port>** command. You can configure rcv-queue and tx-queue globally per rbridge level by using the **qos rcv-queue limit** and **qos tx-queue limit** commands. For more information, refer to the *Network OS Command Reference*. You can adjust these two global qos buffer values to the same egress port and source of congestion.

QoS is not marking or treating packets correctly

Use the Switched Port Analyzer (SPAN) feature to mirror the ingress and egress ports to check that QoS is marking and treating packets correctly. Refer to the “Configuring Switched Port Analyzer” section of the *Network OS Layer 2 Switching Configuration Guide* for details.

If the traffic is not marked properly, check whether the ingress/egress traffic on the port is applied with port-based QoS configuration or flow-based QoS configuration. In flow-based QoS, check if the traffic is applied with RBridge-level flow-based QoS or port-level flow-based QoS. Port-level flow-based QoS takes precedence over system-level or RBridge-level flow-based QoS.

RBridge ID is duplicated

Switches with the same RBridge ID cannot coexist in the same VCS Fabric cluster. Any attempt to add a switch with the same RBridge ID as an existing cluster switch will fail. The ISL between the two switches will not be formed; it will be segmented.

1. On the new switch, enter the **show vcs** command to determine its RBridge ID.

```
switch1# show vcs
Config Mode   : Local-Only
VCS ID       : 1
Total Number of Nodes      : 1
Rbridge-Id   WWN                               Management IP Status  HostName
-----
22           >10:00:00:05:33:13:B3:5A*  10.24.84.41   Online  switch1
```

- On any switch in the functioning VCS Fabric cluster, enter the **show vcs** command to display the RBridge IDs for all the switches in the cluster. In this example, we execute this from switch2, and find that switch2 has the same RBridge ID as switch1.

```
switch2# show vcs
Config Mode   : Local-Only
VCS ID       : 1
Total Number of Nodes      : 2
Rbridge-Id  WWN                               Management IP Status  HostName
-----
22          10:00:00:05:33:5F:EA:A4  10.24.81.65   Online  switch1
22          >10:00:00:05:33:67:26:78* 10.24.81.66   Online  switch2
```

- If the new switch has the same RBridge ID as any switch in the existing cluster, on the new switch, in privileged EXEC mode, enter the **vcs rbridge-id** command to change its RBridge ID to a unique value.

```
switch1# vcs rbridge-id 23
```

SNMP MIBs report incorrect values

If SNMP MIBs report incorrect values, complete the following steps.

- Ensure you are using a supported MIB browser.
- Ensure that the issue is seen consistently.
- Ensure that the SNMP configuration is correct.
- If the MIB browser is supported, the SNMP configuration is correct, and you still see the issue consistently, contact your switch support provider.

SNMP traps are missing

If SNMP traps are missing, complete the following procedure.

- Ensure that the correct SNMP configuration is enabled. Refer to the "Configuring SNMP" section of the *Network OS Administration Guide* for details.
- Ensure that the SNMP host is reachable.
- If the problem still persists, contact your switch support provider.

As a workaround, set a trap configuration for syslog messages.

Telnet operation into the switch fails

Assuming a correct IP address and correct login credentials, failure to access the switch using Telnet could be for one of the following reasons:

- The management port is down. Refer to [Verifying the status of the management port](#) on page 59 for details.
- Access to the management interface is denied by an ACL. Refer to [Checking for a deny ACL](#) on page 59 for details.
- The switch CPU is overloaded. Refer to [Checking for overloaded CPU](#) on page 59 for details.

Verifying the status of the management port

1. On the system console, enter the **show system** command to check the status of the management port, shown underlined in the following example.

```
switch# show system
Stack MAC                : 00:05:33:67:26:78
-- UNIT 0 --
Unit Name                : switch
Switch Status           : Online
Hardware Rev            : 107.4
TengigabitEthernet Port(s) : 60
Up Time                 : up 1 day, 2:52
Current Time            : 23:40:50 GMT
NOS Version             :
Jumbo Capable           : yes
Burned In MAC           : 00:05:33:67:26:78
Management IP           : 10.24.81.66
Management Port Status  : UP
-- Power Supplies --
PS1 is faulty
PS2 is OK
-- Fan Status --
Fan 1 is Ok
Fan 2 is Ok
Fan 3 is Ok
```

2. If the status of the management port is DOWN, enter the **interface management** command to configure the management port correctly. Refer to the “Configuring Ethernet management interfaces” section of the *Network OS Administration Guide*.
3. If the problem persists, contact your switch support provider.

Checking for a deny ACL

On the system console, enter the **show running-config ip access-list** command and check the output to determine whether an ACL is denying access to the management port.

Checking for overloaded CPU

An overloaded switch CPU can prevent Telnet access. Refer to [CPU use is unexpectedly high](#) on page 35.

Traffic is not being forwarded

If the traffic is not being forwarded, perform the following steps:

1. Check for db packet capture. Below are the commands to enable and view a capture.

```
db 8/0/1 rte enable capture all
db 8/0/1 rte start capture
db 8/0/1 rte show capture
```

After the **start capture** command, the system sends a stream and performs **show capture**. This displays most of the capture information:

- a) It shows all the fields resolved — whether it is trap, drop, or fwd.
- b) It shows the packet itself.
- c) It shows the Routing Engine (RTE) Layer 2 history, as in the result of the Layer 2 table hit or miss.
- d) It shows the RTE Layer 3 history, as in the result of the Layer 3 table hit or miss. If Layer 2 table has a success and Layer 3 table failed, then check for routing issues.

For example, in the entry for trap (Ping to box), the Routed fields should display `Ipv4Rtd` and the entries hit. For a TRAP hit, it should display `trape:1` (Bit set to 1 to indicate packet is trapped).

- e) Packet Capture displays the last four packets, Make sure those are the fwd packets (for example, check for SA DA MAC, pkttyp:0806).

```
db 8/0/1 rte enable capture all
db 8/0/1 rte start capture
db 8/0/1 rte show capture
```

- f) If the FWD and packets are not being forwarded, then it is an ASIC problem. If it is a DROP, proceed to the next step.
2. If result is DROP:
 - a) Execute the **show ip route** command. If the route is not present, then it is an RTM issue.
 - b) Execute the **show arp command**. If ARP is not resolved for the corresponding next hop, then it is an ARP issue. If it is VLAN along with ARP, MAC should be resolved. If MAC is not resolved then it is an L2SYS issue.
 - c) If step **2b** passes, enter the **debug show ip lpm** command to display the routes in hardware, and verify that the corresponding destination ARP address is present. If it is not present, then it is an L3FWD issue. Collect the information from **debug show ip lpm**, attach the file `/tmp/fib_wlv_ioct1`, along with supportSave data and contact support. Specify that it failed in this step.

Refer also to [Gathering troubleshooting information](#) on page 19, which provides information about Network OS supportSave files.

- d) If step **2c** passes and traffic is still dropping, then it is an ASIC issue.

NOTE

For additional information about packet capture, refer to [Using the packet capture utility](#) on page 74.

Trunk member not used

If you suspect that one or more members of a trunk are not being used, complete the following steps.

NOTE

The E1MG-SX-OM and E1MG-LX-OM modules are not supported by Network OS. Despite being Brocade products, these modules will return the error 'Optic is not Brocade qualified, optical monitoring is not supported' and must be replaced with a supported module.

1. Enter the **show running-config interface** command to determine which interfaces have trunking enabled.

```
switch# show running-config interface
interface Management 66/0
  no ip address dhcp
  ip address 10.24.81.66/20
  ip route 0.0.0.0/0 10.24.80.1
  ipv6 address ""
  no ipv6 address autoconfig
!
interface TenGigabitEthernet 66/0/1
  fabric isl enable
  fabric trunk enable
  no shutdown
!
interface TenGigabitEthernet 66/0/2
  fabric isl enable
  fabric trunk enable
  no shutdown
!
interface TenGigabitEthernet 66/0/3
  fabric isl enable
  fabric trunk enable
  no shutdown
!
      (output truncated)
```

2. Verify the status of the ISL port and link.
 - a) Enter the **show fabric isl** command to verify whether the ISL is up.
 - b) Enter the **show fabric islports** command to examine the status of each port.

Refer to [Verifying the status of ISLs](#) on page 41 for details and corrective action.

- Enter the **show interface** command for each trunk link and examine the rate information to check for an equal distribution of traffic on the interfaces in the trunk. The rate information is shown underlined in the following example.

```
switch# show interface tengigabitethernet 66/0/12
TenGigabitEthernet 66/0/12 is up, line protocol is down (link protocol down)
Hardware is Ethernet, address is 0005.3367.26a8
  Current address is 0005.3367.26a8
  Pluggable media not present
  Interface index (ifindex) is 283871281409
  MTU 2500 bytes
  LineSpeed Actual      : Nil
  LineSpeed Configured : Auto, Duplex: Full
  Flowcontrol rx: off, tx: off
  Last clearing of show interface counters: 1d00h42m
  Queueing strategy: fifo
  Receive Statistics:
    0 packets, 0 bytes
    Unicasts: 0, Multicasts: 0, Broadcasts: 0
    64-byte pkts: 0, Over 64-byte pkts: 0, Over 127-byte pkts: 0
    Over 255-byte pkts: 0, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
    Over 1518-byte pkts(Jumbo): 0
    Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
    Errors: 0, Discards: 0
  Transmit Statistics:
    0 packets, 0 bytes
    Unicasts: 0, Multicasts: 0, Broadcasts: 0
    Underruns: 0
    Errors: 0, Discards: 0
  Rate info:
    Input 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
    Output 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Time since last interface status change: 1d03h16m
```

- Having found a trunk member that carries no traffic while the other trunk members are busy, from the same **show interface** command output, check the interface status, configuration, and error statistics.
 - If the interface is disabled, enable it with the **no shutdown** command.
 - If misconfiguration is apparent, refer to [Trunk member not used](#) for information on how to configure fabric trunks.
 - If you notice significant errors in the error statistics counters, depending on the error, check the SFP transceiver and cable on the local switch and on the peer switch at the other end of the cable.
 - Enter the **show media interface** command on each switch and check the Vendor Name field to check that the optics are Brocade-certified. If the Vendor Name field shows anything other than BROCADE, replace the optics with Brocade-certified optics.
 - Replace any non-Brocade SFP transceiver.
 - Try replacing the SFP transceiver.
 - Try replacing the cable.

Upgrade fails

If a failure occurs during firmware upgrade, complete the following steps.

- Revert to the previous firmware version.
- Contact your switch support provider to evaluate whether retrying the upgrade is appropriate.

VCS Fabric cannot be formed

A VCS Fabric can fail to form for several reasons:

- The VCS Fabric configuration is incorrect. The following configuration issues will prevent the VCS Fabric from forming:
 - The VCS ID on the constituent switches is not the same.
 - Multiple switches have the same RBridge ID.
 - ISL ports that connect the switches are not up.

Verifying the VCS Fabric configuration

To verify the VCS Fabric configuration, complete the following steps.

1. Enter the **show vcs** command on each switch to verify that the VCS ID on all switches is the same and the RBridge ID on each switch is different.
2. Enter the **show fabric isl** command to verify whether the ISL is up.
If ISL is down, refer to [ISL does not come up on some ports](#) on page 40.
3. Enter the **show fabric islports** command to examine the status of each port.

vLAG cannot be formed

A vLAG trunk can fail to form for several reasons:

- The link between the VCS Fabric switches does not exist. Refer to [Verifying the link between the VCS Fabric switches](#) on page 63.
- A bad connection causes abnormal reception or transmission of LACPDUs. Refer to [Verifying LACPDUs](#) on page 64.
- Port-channel numbers are not the same on the VCS Fabric switches. Refer to [Verifying the vLAG configuration](#) on page 64.
- The peer switches are not configured in the same LACP mode (static or dynamic). Refer to [Verifying the LACP mode of each switch](#) on page 64.
- A 1-Gbps port-channel has been upgraded to Network OS 2.1.x or later. Refer to [Explicitly setting the speed for a 1-Gbps port-channel](#) on page 64.

Verifying the link between the VCS Fabric switches

The link between switches could be broken for various reasons:

- A port is not activated.
- The ISL is segmented.
- The VCS Fabric is not properly formed.
- The CPU is overload.

Refer to [ISL does not come up on some ports](#) on page 40 for details on detecting and correcting the problem.

Verifying LACPDUs

LACPDUs should be transmitted and received on both ends of the vLAG. This procedure verifies whether that is happening, and also checks for PDU errors.

1. On both switches, enter the **show lacp counter** command to verify that LACPDUs are transmitted and received, and there are no error PDUs.

```
switch# show lacp counter 10
% Traffic statistics
Port          LACPDU      Marker      Pckt err
              Sent   Recv      Sent   Recv      Sent   Recv
% Aggregator  Po 10 1000000
Te 0/1        65     0         0     0         0     0
Te 0/2        64     0         0     0         0     0
Te 0/3        64     0         0     0         0     0
Te 0/4         0     0         0     0         0     0
```

In this case, LACPDUs are being transmitted by the switch, but none are being received.

2. If the output shows that LACPDUs are not being transmitted and received correctly, or packet errors are showing, contact your switch support provider.

Verifying the vLAG configuration

The port-channel number must be the same across all vLAG member switches, or the vLAG will not form.

1. On each vLAG member switch, in privileged EXEC mode, enter the **show port-channel summary** command.

```
switch# show port-channel summary
Static Aggregator: Po 15
Aggregator type: Standard
Member ports:
Te 0/6
Te 0/7
Te 0/14
Te 0/15
...
switch2# show port-channel summary
switch2#
```

2. If the port-channel does not appear on both switches, on the switch where it does not appear, in global configuration mode, enter the **interface port-channel** command to create the port-channel.

```
switch2(config)# interface port-channel 15
```

Refer to the “Configuring Link Aggregation” chapter of the *Network OS Layer 2 Switching Configuration Guide* for details.

Verifying the LACP mode of each switch

A vLAG must be configured either statically on both ends of the vLAG, or dynamically on both ends of the vLAG. Refer to “Configuring Link Aggregation” chapter of the *Network OS Layer 2 Switching Configuration Guide* for details.

Explicitly setting the speed for a 1-Gbps port-channel

To set the port speed to 1 Gbps, complete the following steps.

1. In interface configuration mode, shut down the port-channel.

```
switch(config-Port-channel-2)# shutdown
```


2. Set the port-channel speed to 1 Gbps.

```
switch(config-Port-channel-2)# speed 1000
```

3. Re-enable all port members in the port-channel.

```
switch(config-Port-channel-2)# no shutdown
```

Zoning conflict needs resolution

In case a zoning conflict is encountered, you must make sure that the zoning configuration on the local switch matches that on the switch that it is joining. Zone conflicts can be resolved by saving a configuration file with the **configUpload** command, examining the zoning information in the file, and performing a cut-and-paste operation so that the configuration information matches in the fabrics being merged.

After examining the configuration file, you can choose to resolve zone conflicts by using the **cfgDisable** command, followed by the **cfgClear** command, on the incorrectly configured segmented fabric. Then enter the **cfgSave** command, followed by the **portDisable** and **portEnable** commands, on one of the ISL ports that connects the fabrics. This causes a merge, making the fabric consistent with the correct configuration.

ATTENTION

Be careful when using the **cfgClear** command, because it deletes the defined configuration.

When merging two fabrics, multiple zoning CLI sessions can be launched on the same switch, or on different switches. The following describes these situations and how they are automatically resolved.

Dual-CLI sessions from the same switch: If you start a zone transaction from CLI-Session1 and then try to perform a zone modification from CLI-Session2, the CLI-Session2 zone transaction is not allowed, as CLI-Session2 is not the owner of the open transaction. If CLI-Session1 logs out, this ends the open transaction and aborts any current zone modifications. CLI-Session2 is then able to perform zone modifications. Therefore, the zone transaction locking mechanism works on a single switch from the CLI perspective and there is no dangling transaction.

Dual-CLI sessions from different switches: If you start a CLI zone transaction on Switch1 and started another CLI zone transaction on Switch2, when the zone transaction from Switch1 is committed, the open zone transaction from Switch2 is aborted by Switch1. The following message is posted on Switch2 at the time of zone commit from Switch1:

```
2014/01/09-21:45:26, [ZONE-1027], 3285, FID 128, INFO, switch, Zoning transaction aborted Zone Config
update Received
```

NOTE

The above is applicable only in fabric cluster mode. This does not apply to logical chassis cluster mode, because all zoning operations must be performed from the principal switch.

Using troubleshooting and diagnostic tools

This section describes the various troubleshooting and diagnostic tools available with Network OS and provides some guidelines for their use .

Refer also to [Gathering troubleshooting information](#) on page 19, which provides information about Network OS supportSave files.

Using Layer 2 traceroute


TRILL OAM provides the **I2tracertool** command to verify the fabric path continuity. When the **I2tracertool** command is used with extended options, it provides granular control over the Layer 2 path that a Layer 2 traceroute packet takes.

Layer 2 traceroute packets

To use the Layer 2 traceroute tool, you need to understand the structure of the Layer 2 traceroute packet when observed on the wire, when it is a request frame, and when it is a response frame.

The figure below shows what a normal Layer 2 packet looks like when traversing through an Ethernet fabric, without Layer 2 traceroute applied.

FIGURE 1 Normal Layer 2 packet traversing a VCS fabric



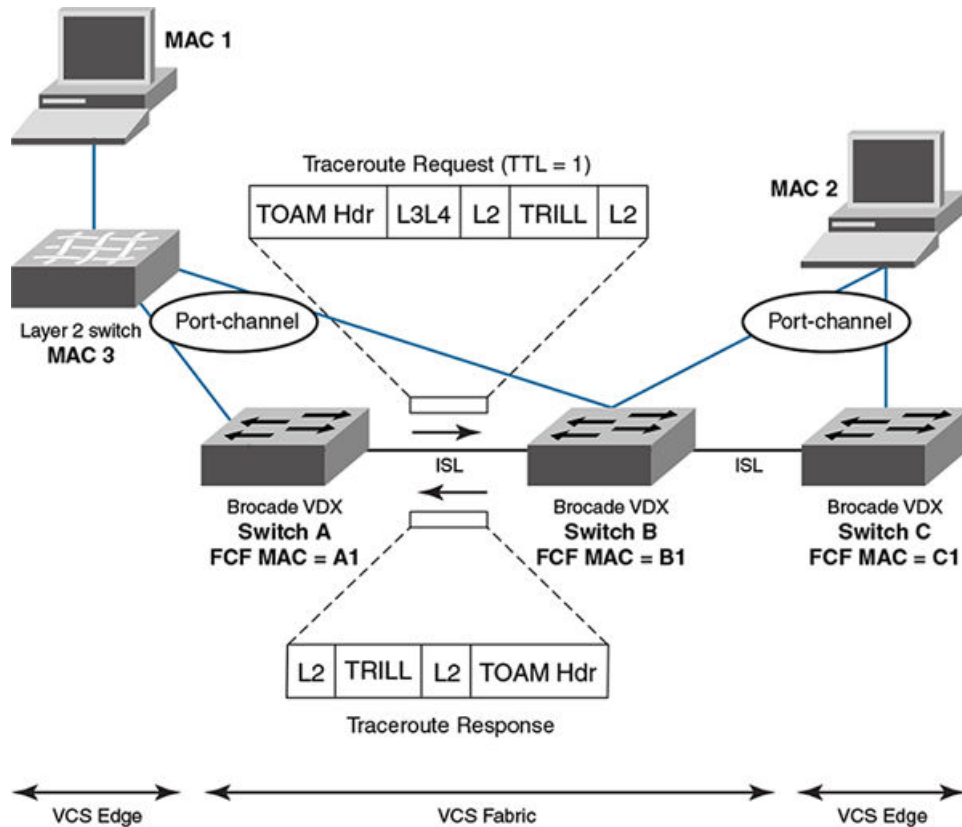
In the figure above, an Ethernet packet arrives from MAC 1 at the VCS fabric edge. TRILL header information is added while the packet passes through the VCS fabric. The TRILL information is removed on leaving the VCS fabric, and a regular Ethernet packet arrives at MAC 2. The table below shows the Layer 2 packet header details.

TABLE 4 Packet header details — Layer 2 packer traverses VCS fabric

Ethernet packet	TRILL packet — first hop	TRILL packet — second hop
L2 DA = MAC 2 L2 SA = MAC 1	Outer L2 DA = B1 Outer L2 SA = A1 Outer 802.1q tag Outer etype = TRILL TRILL destination RBridge ID = C TRILL source RBridge ID = A TRILL flags Inner L2 DA = MAC 2 Inner L2 SA = MAC 1 Inner 802.1q tag Inner etype = 0x800	Outer L2 DA = C1 Outer L2 SA = B1 Outer 802.1q tag Outer etype = TRILL TRILL destination RBridge ID = C TRILL source RBridge ID = B TRILL flags Inner L2 DA = MAC 2 Inner L2 SA = MAC 1 Inner 802.1q tag Inner etype = 0x800

When viewing packets while using the **I2tracertool** command, notice the TRILL OAM header information added to the packets as they traverse the VCS Fabric. Starting the trace on Switch A, TRILL OAM first verifies path continuity with its immediate neighbor, in this case Switch B. It does this as shown in the figure below, by sending a Layer 2 traceroute request packet with the time-to-live (TTL) TRILL attribute set to 1. Switch B replies with reachability information regarding the next hop.

FIGURE 2 Verifying path continuity with immediate neighbor



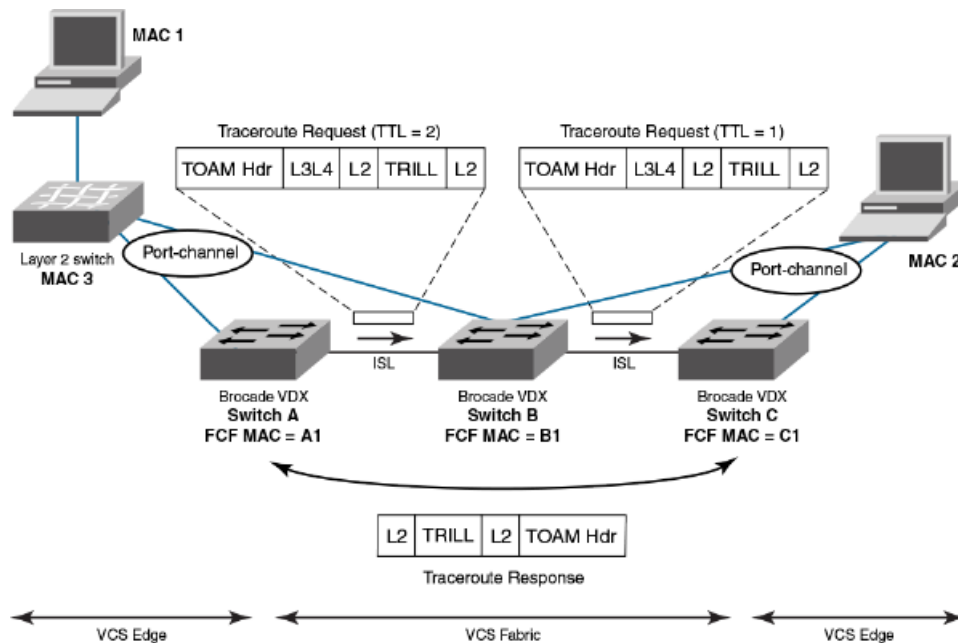
The table below shows the packet header information for the request and response. The added TRILL OAM information is shown in **bold**.

TABLE 5 Packet header details with Layer 2 traceroute – first hop

Traceroute request packet header	Traceroute reply packet header
Outer L2 DA = B1 Outer L2 SA = A1 Outer 802.1q tag Outer etype = TRILL TRILL destination RBridge ID = C TRILL source RBridge ID = A ATRILL flags: TTL = 1 Inner L2 DA = MAC 2 Inner L2 SA = MAC 1 Inner 802.1q tag Inner etype = 0x800 TOAM Opcode = 5 (request)	Outer L2 DA = B1 Outer L2 SA = A1 Outer 802.1q tag Outer etype = TRILL TRILL destination RBridge ID = A TRILL source RBridge ID = B BTRILL flags: TTL = MAX (63) Inner L2 DA = A1 Inner L2 SA = B1 Inner 802.1q tag Inner etype = TRILL OAM TOAM Opcode = 4 (reply) C reachable

Having successfully exchanged packets with the immediate neighbor (Switch B) and established the reachability of Switch C, the Layer 2 traceroute feature issues another request with TTL set to 2. Switch B decrements the TTL count and forwards the packet to Switch C, which returns a response to Switch A. Refer to the figure below.

FIGURE 3 Verifying path continuity— second hop TTL count



The table below shows the packet header information for the request and response packets. Information specific to the Layer 2 traceroute feature is show in **bold**.

TABLE 6 Packet header details with Layer 2 traceroute – second hop

Traceroute request – first hop (TTL = 2)	Traceroute request – second hop (TTL = 1)	Traceroute reply
Outer L2 DA = B1 Outer L2 SA = A1 Outer 802.1q tag Outer etype = TRILL TRILL destination RBridge ID = C TRILL source RBridge ID = A TRILL flags: TTL = 2 Inner L2 DA = MAC 2 Inner L2 SA = MAC 1 Inner 802.1q tag Inner etype = 0x800 TOAM Opcode = 5 (request)	Outer L2 DA = C1 Outer L2 SA = B1 Outer 802.1q tag Outer etype = TRILL TRILL destination RBridge ID = C TRILL source RBridge ID = B TRILL flags: TTL = 1 Inner L2 DA = MAC 2 Inner L2 SA = MAC 1 Inner 802.1q tag Inner etype = 0x800 TOAM Opcode = 5 (request)	Outer L2 DA = B1->A1 Outer L2 SA = C1->B1 Outer 802.1q tag Outer etype = TRILL TRILL destination RBridge ID = A TRILL source RBridge ID = C TRILL flags: TTL = MAX (63) Inner L2 DA = A1 Inner L2 SA = B1 Inner 802.1q tag Inner etype = TRILL OAM TOAM Opcode = 4 (reply)

Tracing a route with the `l2tracert` command

In the following example, the `l2tracert` command verifies the path between port 3/0/1 (source MAC address 0050.5685.0003) and port 2/0/9 (destination MAC address 0024.3878.3720).

1. enter the `show mac-address-table` command to display all known MAC addresses in the network.

```
device# show mac-address-table

VlanId  Mac-address  Type  State  Ports
-----  -
100     0024.3878.e720 Dynamic Active Po 11
100     0050.5685.0001 Dynamic Active Po 1
101     0000.0000.0003 Dynamic Active Po 1
101     0024.3878.e720 Dynamic Active Po 11
101     0050.5685.0003 Dynamic Active Po 1
Total MAC addresses : 5
```

From the output, choose the source and destination MAC address:

- Source MAC address: 0050.5685.0003
- Destination MAC address: 0024.3878.e720

2. Enter the `l2tracert` command.

```
device# l2tracert
Source mac address      : 0050.5685.0003
Destination mac address : 0024.3878.e720
Vlan [1-3962]          : 101
Edge rbridge-id [1-239] : 3
Extended commands [Y/N]? : y
Protocol Type [IP]     : IP
Source IP address      : 101.101.101.10
Destination IP address : 101.101.101.101
IP Protocol Type [TCP/UDP] : TCP
Source port number [0-65535] : 3000
Dest port number [0-65535] : 22
Rbridge  Ingress          Egress          Rtt (usec)
-----  -
3        Te 3/0/1(std-lag, Po 1) Te 3/0/20(isl)  0
2        Te 2/0/20(isl)          Te 2/0/9(std-lag, Po 11) 34041
```

Be advised of the following points:

- The MAC addresses used should be present in the MAC address-table (dynamic or static).
- Make use of IP parameters to influence path selection.

Using show commands

The table below lists some `show` commands that are often used for troubleshooting. Refer to the *Network OS Command Reference* for details of all `show` commands.

TABLE 7 show commands used for troubleshooting

Command group	Commands	Specific fields or purpose
System commands	show system show license show running-config show startup-config show logging raslog show version	

TABLE 7 show commands used for troubleshooting (continued)

Command group	Commands	Specific fields or purpose
	show chassis show environment show vlan brief show mac-address-table show process cpu show process memory show firmwaredownloadstatus	
Interface commands	show interface show media show ip int brief show qos flowcontrol interface show qos queue interface show qos rcv-queue interface show qos int	Check pause-frames Check the CoS statistics Check packet drops, buffer consumption, real-time queue statistics Check the QoS configuration on an interface
Diagnostic commands	show diags status show diags post results detailed show diag burninerrshow show diag burninstatus	
Feature commands	show port-channel detail show lacp counter show port-profile status show fcoe login show fcoe interface brief show fcoe internal fcf-mac-address show lldp neighbors detail show lldp statistics show qos interface all show udd statistics	
VCS Fabric commands	show vcs show fabric trunk all show fabric all show fabric isl show fabric islports show fabric route linkinfo show fabric route multicast show fabric route neighbor-state show fabric route pathinfo	

TABLE 7 show commands used for troubleshooting (continued)

Command group	Commands	Specific fields or purpose
	show fabric route topology show name-server detail all	

Using debug and system diagnostic commands

Diagnostic commands, such as "debug" and "show system internal" commands, are developed and intended for specialized troubleshooting. Brocade recommends that you work closely with Brocade technical support in executing such commands and interpreting their results.

General overview

You can perform the following operations related to debugging features:

- To enable debugging on a feature, use the **debug** command in privileged EXEC mode.

```
device# debug feature required-keywords
```

- To check whether debugging is enabled on a feature, use the **show debug** command in privileged EXEC mode.

```
device# show debug feature
```

- To disable debugging, use the **no debug** command.

```
device# no debug feature required-keywords
```

Use caution when debugging in real time on a production switch, because real-time debugging is CPU-intensive.

Brocade recommends checking the command output on a lab switch first, and then if the output looks acceptable, enable it on the production switch to get more data. In addition, to reduce CPU load, Brocade recommends using keywords such as **events** and **summary** that limit the extent of debugging rather than more comprehensive options such as **detail** and **all**.

Debugging operations are used mainly for debugging control plane protocols such as LACP and LLDP. For example, to view received LLDP packets on the console, use the following command:

```
device# debug lldp packets all rx
```

If the switch is accessed through Telnet, enable logging using a terminal monitor.

The following are the most frequently used debug commands:

- debug lldp packets interface [rx | tx | both]**
- debug lacp pdu [rx | tx]**
- debug spanning-tree bpdu [rx | tx]**
- debug dot1x packet**

Enhanced debugging

In addition, you can use a variety of commands to debug the status of various hardware components and processes, such as the following:

- Buffer-pool queue (BPQ) drops on ingress ports, with mapping of BPQs to particular protocol functions
- ASIC counters for a variety of internal functions (not limited to MAC addresses)

- Packet drops with granularity (Packet-drops as summations of multiple events do not indicate true causes.)

These commands are also supported by supportSave. The results are collected at least three times back-to-back for offline analysis.

The following table lists some of the commands that report counter status for ASICs and software components. For **show interface stats cpu** and additional **show system internal** commands, refer to the *Network OS Command Reference*.

Command	Description
show system internal ASIC counter blk	Displays packet-count for block-level packet transfer between ASIC blocks.
show system internal ASIC counter interface	Displays nonzero MAC counters for a specified Ethernet interface.
show system internal ASIC counter mem blk	Displays memory-error count for block-level packet transfer between ASIC blocks.
show system internal ASIC counter drop-reason	Displays the count of dropped packets, sorted by with drop reasons.

Using SPAN port and traffic mirroring

In certain instances, you may need to examine packets in transit across links to understand the traffic pattern on a specific port. In such situations, Switched Port Analyzer (SPAN) can be configured to copy the traffic (with the desired direction) on the specific Ethernet port to a mirror port where a sniffing device is connected. You can then analyze the packets captured by the sniffing device.

```
device(config)# monitor session 1
device(conf-mon-sess-1)# source tengigabitethernet 1/0/10 destination tengigabitethernet 2/3/15 direction
both

device# show monitor 1
Session :1
Description :Test SPAN Session
State :Enabled
Source interface : 1/0/10 (Up)
Destination interface : 1/0/15 (Up)
Direction :Both
```

Note the following guidelines:

- The source port cannot be an ISL port.
- The destination port cannot be an ISL, Layer 2, Layer 3, QoS, ACL, 802.1x, LAG member, LLDP, or port-profile port.
- Only edge ports are eligible for mirroring.

Using hardware diagnostics

The following diagnostic types currently exist:

- Power-on self-test (POST)
- Offline diagnostics

Online diagnostics are not currently supported on Brocade VDX switches.

Using POST diagnostics

POST is run on bootup and the results are stored. Use the **show diag post results** command to view the stored results.

To enable POST diagnostics, enter the **diag post rbridge-id rbridge-id enable** command. POST diagnostics are enabled by default.

Using offline diagnostics

Before proceeding, note the following Caution:



CAUTION

Offline diagnostics — otherwise known as system verification tests — are disruptive tests that check the individual hardware components thoroughly and report the findings. You must disable the chassis before running these tests. Do not run production traffic during this time.

Enter the **diag systemverification** command to run the entire set of offline diagnostics. This command can take up to two hours to finish, so Brocade recommends the less disruptive **diag systemverification short** command, which typically takes 10 to 15 minutes. Alternatively, you can run subsets of the offline commands that check various parts of the hardware. The table below shows the complete list of supported offline commands.

TABLE 8 Offline diagnostic commands

Offline diagnostic command	Purpose
diag burninerrclear	Clears the errors that are stored in the nonvolatile storage during the burn-in process.
diag clearerror	Clears the diagnostics failure status.
diag portledtest	Runs various action modes on the port LEDs and validates the functionality.
diag portloopbacktest	Sends frames between various ASICs on the switch and validates the ASIC functionality.
diag setcycle	Configures all the parameters required for the system verification test.
diag systemverification	Runs a combination of various hardware diagnostic tests.
diag turboramtest	Performs a turbo static RAM (SRAM) test of the ASIC chips.

The table below lists the show commands that provide output from offline diagnostics.

TABLE 9 Offline diagnostic show commands

Show offline diagnostic command	Purpose
show diag burninerrshow	Displays the errors that are stored in the nonvolatile storage during burn-in.
show diag burninstatus	Displays the diagnostics burn-in status.
show diag setcycle	Displays the current values used in system verification.
show diag status	Displays the currently running diagnostics tests.

For details of the commands listed in these tables, refer to the *Network OS Command Reference*.

Viewing routing information

The **show fabric route pathinfo** command displays routing and statistical information from a source port index on the local switch to a destination port index on another switch in the same VCS Fabric cluster, a different VCS Fabric cluster, a connected Fabric OS backbone fabric, or Fabric OS edge fabric. This routing information describes the full path that a data stream travels between these ports, including all intermediate switches.

The routing and statistics information are provided by every switch along the path, based on the current routing table information and statistics calculated continuously in real time. Each switch represents one hop.

Use the **show fabric route pathinfo** command to display routing information from a source port on the local switch to a destination port on another switch. The command output describes the exact data path between these ports, including all intermediate switches.

To use the **show fabric route pathinfo** command across remote fabrics, you must specify both the VCS ID (or Fabric ID) and the RBridge ID (or domain ID) of the remote switch. When obtaining path information across remote fabrics, the destination switch must be identified by its RBridge ID or domain ID. Identifying the switch by name or WWN is not accepted.

For details about the **show fabric route pathinfo** command, refer to the *Network OS Command Reference*.

Using the packet capture utility

When a packet is received at a switch's source port, it is routed through the switch to the destination port. If a problem occurs, SPAN or sFlow are commonly used. However, SPAN requires a network analyzer, and sFlow requires a collector.

The Network OS packet capture utility is a built-in sniffing mechanism that uses the pcap API to capture packets destined to the CPU. The results can then be displayed by means of a show command or exported as a .pcap file for offline analysis by means of a tool such as Wireshark. Packets from multiple interfaces can be captured simultaneously, and the capture is rate-limited to prevent overloading the CPU.

There are two ways to view the results of packet capture:

- By using the appropriate debug and show commands on the target switch.
- By viewing the results in an automatically generated file.

The following table lists packet capture commands for both IPv4 and IPv6.

TABLE 10 Packet capture commands

Command	Description
capture packet interface	<p>Captures IPv4 packets destined toward the CPU, as well as transit packets if a trap is enabled by means of ACL logging.</p> <p>NOTE This command can provide significant help in debugging, especially for Layer 2 TRILL and Layer 3 packets. Up to 100 packets per interface can be captured. Once the buffer is filled, the oldest packets are replaced with the most recent. Captured packets are stored in a circular buffer, and they are also written to an automatically generated <code>pktcapture.pcap</code> file, which can store up to 1500 KB of data in flash memory (the equivalent of approximately 10 KB packets, each having an average size of 100 bytes). Once this file is full, it is saved as <code>*_old.pcap</code> and data are written to a new <code>pktcapture.pcap</code> file.</p>
debug ipv6 packet	<p>Enables IPv6 packet capture on an interface or all interfaces.</p> <p>NOTE The packets are stored in either circular or linear buffers, with up to 20-56 packets stored per switch. Once the .pcap file is full, the file-rotation technique is used to manage incoming data by wrapping it in a circular buffer. Up to 1500 KB of data can be stored in flash memory. This represents 11180 of packets having an average size of 130 bytes. This feature is also supported on virtual Ethernet interfaces.</p>
debug ipv6 icmpv6	Enables the capture of IPv6 packets related to Internet Control Message Protocol (ICMP) version 6.
show capture packet interface	Displays information about captured IPv4 packets.
show debug ipv6 packet	Displays IPv6 packets captured through the packet capture utility on an interface or all interfaces, as well as the packet capture configuration on the switch.

For command details, refer to the *Network OS Command Reference*.

Note the following limitations:

- Support is provided only on physical interfaces (gigabit Ethernet), not on logical interfaces. To capture packets on logical interfaces, first enable the capture on the corresponding physical interfaces.
- Support for capturing transit traffic requires ACL logging.
- Packets that are dropped in the ASIC cannot be captured.



CAUTION

Capturing packets over multiple sessions and over long durations can affect system performance.

Debug Automation Suite for tracing data-path

The `ipfabric-ingressNode-debugv2.py` script enables you to trace the data path through the IP Fabric network and thoroughly check the integrity of the "route-programming" in VDX devices along the data path.

Traditionally, numerous set of commands are utilized to check the prefix health across various layers in software, ranging from BGP, RIB, FIB and Hardware-specific commands for forwarding path. However, this has been fully automated Network OS, with a suite of scripts that provide flow information and packet-path tracking across the IP Fabric. Once the prefix-health-check is completed on the device, it points to the next IP hop where a given string needs to be copy-pasted, in turn triggering another set of scripts that evaluate the health of the following node. This way, the script traces the data-path of the flow by thoroughly checking the health of all the intermediate routers. This tool complements the `traceroute` command, but tests the control-plane the dataplane programming in-depth on each node in the path.

Thus when you encounter a situation where certain flows or all flows are impacted, you can identify a specific impacted flow and enter its parameters (such as SA, DA, SIP, DIP, VLAN, and ingress interface) and trigger the automation suite on the VDX switch to trace the packet path from the ingress node VDX switch -> intermediate spine node VDX switches -> egress node VDX switch. This enables you to identify the specific switch impacting the traffic flow and pin-point the specific error.

If you are uncertain of the specific ingress interface for that flow (such as the port-channel members), you may provide a list of interfaces to check. The script runs through all interfaces to check for tagged or untagged traffic for that flow and identifies the interface and forwarding operational values.

The command syntax for the automated script suite is:

```
execute-script ipfabric-ingressNode-debugv2.py <DA-MAC address> <SA-MAC-Address> <DIP> SIP <Vlan>
<Interface number>
```

Refer to the *Network OS Command Reference* for complete details about the `execute-script` command.

Example of online help output:

```
switch# execute-script ipfabric-ingressNode-debugv2.py

Usage: execute-script ipfabric-ingressNode-debugv2.py [-verbose] <mac-da> <mac-sa> <dip> <sip> <vlan> <int-
list>
Usage: execute-script ipfabric-ingressNode-debugv2.py [-verbose] [-prefix-check] <dip> <vrf> <int-list>
Example: execute-script ipfabric-ingressNode-debugv2.py aaaa.bbbb.cccc dddd.eeee.ffff 10.1.1.81/24
10.1.1.82/24 10 0/21,0/31
Example: execute-script ipfabric-ingressNode-debugv2.py -verbose -prefix-check 10.1.1.81/24 blue 0/21,0/31

Points to note:
1. For routed flow, please use the correct DA based upon the default-gateway on the end-host
2. For non-IP switched flow, please use dip = sip = 0.0.0.0/0
3. For <int-list>, use comma-separated list of probable ingress int(s) for this flow (eg. PO members, ECMP
members)
```

4. For <int-list>, please ensure all interfaces listed are of same type ie. l2-switchport OR ISL link OR l3-routed

The ipfabric-ingressNode-debugv2.py script run on an ingress node runs through a thorough check of the integrity of the "route-programming" in that VDX device and provides detailed user-friendly logs and "Alerts" for any discrepancies encountered. At the end of the output, it provides a concise table summary (shown below) at the end and a string that can be copy-pasted over to be applied on the next-hop IP Address listed.

NOTE

The ipfabricingressNode-debugv2.py script supports a traditional leaf-spine IP-Fabric topology and traditional layer-3 networks for prefix-checks. The automation suite systematically runs through debugging steps and verifies vital data. However, it is not a replacement for the copy-support command.

The script involves the following debugging elements:

- Validating underlay Layer 2 and Layer 3 infrastructure for Leaf and Spine
- BGP peering between leaf and spine for EVPN advertisements of MAC and ARP
- BGP operational state and contents (EVPN and VRF)
- Overlay tunnels between the Leaf Nodes (VLAN to VNI mapping)
- Control plane packet tracing (Route, ARP, MAC)
- Programming of the routes in the VDX platform ASICs
- Tromboned traffic destined towards VCS LAG Partner
- Packet capture of that flow and operational value of forwarding decision

The above is applicable for following traffic flow types:

- Layer 2 Forwarding (local or overlay)
- Asymmetric Routing Forwarding (local or overlay)
- Local Peering
- Symmetric Routing Forwarding (local or overlay)

As simple prefix check in the routing table can be executed using CLI:

```
execute-script ipfabric-ingressNode-debugv2.py [prefix-check] <dip> <vrf> <int-list>
```

For Layer 3 routed traffic flow, the MAC DA in the input should be populated with the appropriate Router MAC (such as VE MAC, Anycast MAC, or VRRP vMAC). If the script encounters any suspicious discrepancies, it displays an "Alert" message, but does not take recovery actions. The ipfabric-ingressNode-debugv2.py script is built to be run on Leaf Nodes of the IP fabric that may be running on the Brocade VDX 6740 and Brocade VDX 6940.

The suite only supports IPv4 and accepts a layer-2 port as an ingress interface. It does not support the flow through the ingress physical routed port on the leaf. The ipfabric-ingressNode-debugv2.py script is executed using the execute-script command in privileged EXEC mode. Refer to the *Network OS Command Reference* for complete details about the **execute-script** command.

```
device# execute-script ipfabric-ingressNode-debugv2.py 50eb.1a81.0857 0067.6866.0010 72.1.77.2/24
62.1.66.2/24 6201
```

```
INPUT PARAMETERS:
=====
```

```
DA = 50eb.1a81.0857 , SA = 0067.6866.0010 , DIP = 72.1.77.2 , SIP = 62.1.66.2 , Ingress Vlan = 6201
```

```
DERIVING & VERIFYING FROM INPUT PARAMETERS:
=====
```

```
VDX Platform Model: 6740
```

Management Interface IP Addresses: ['10.20.234.67/21']

Verified the Vlan is valid & active..further checks..

VE MAC Address = 50eb.1a81.0857

VE Interface config: Anycast IP configured: 62.1.66.1/24 Anycast MAC configured: 0000.6768.0001

Extracting vrf, vni, VE information....<please wait>..

Appended list of VRF Import RT configured for this VRF: 6610:6210,7710:7200

DERIVED PARAMETERS:

=====

Interface VE IP Address = 62.1.66.1, VRF = vrf10, Int VE MAC addr (My DA list) = ['50eb.1a81.0857', '0000.6768.0001']

L3VNI = 7200 , L3VNI VE = 7200 VRF Import L3-RT configured for this VRF:

6610:6210,7710:7200

Comment: Flow classified as L3FWD as DA:50eb.1a81.0857 is in MYDA List:['50eb.1a81.0857', '0000.6768.0001']

VERIFYING MY-DA PROGRAMMED IN HW:

=====

VE MAC programmed in HW 50eb1a810856 matches that in Software 50eb1a810856 ..proceeding

DEBUG: Constructed Anycast MAC programmed for VDX6740 in HW = 000067680001

VE Anycast MAC 000067680001 programmed in HW matches that in Software 000067680001 ..proceeding

CHECK FOR DIP IN LOCAL NODE TO CLASSIFY FLOW:

=====

Route found for DIP: 72.1.77.2

SYMMETRIC ROUTING: It is an EVPN Route with Egress Port VE 7200 being the same as VRF L3VNI VE 7200

ECMP Route 1 : 72.1.77.0/24 EVPN Ve 7200 Bi

DEBUG: 2 Byte AS = 122

CHECKING VNI SYNC BETWEEN HW & SW

=====

Verified the VNI programmed in Software for VLAN 7200 is same as programmed in Hardware

IDENTIFYING THE EVPN BGP PEERS:

=====

Checking BGP table for Peering:

Total Number of BGP Peers (including non-eVPN peers): 3

Details about eVPN BGP peers:

eVPN Local BGP Peering IP: 67.67.67.67 / eVPN Remote BGP Peering IP: 122.122.122.122

eVPN Local BGP Peering IP: 67.67.67.67 / eVPN Remote BGP Peering IP: 125.125.125.125

SYMMETRIC FLOW VERIFICATION:

=====

Output derived from show ip route detail:

Path 1 : GW MAC Address: 0005.3365.377b Tunnel Interface: Tu 61441

Path 2 : GW MAC Address: 0005.3365.3633 Tunnel Interface: Tu 61441

SYMMETRIC RTG: CHECKING THE BGP VRF TABLES:

=====

Checking BGP VRF table for the Route existence:

Route FOUND in the BGP VRF Table:

BGP Best Route: *>i 72.1.77.0/24 77.77.77.77

Non-best Route: *i 72.1.77.0/24 77.77.77.77

Checking BGP EVPN table for the Route Properties:

Route FOUND in BGP EVPN table..checking further..

Checking BGP EVPN table for the RT values:

Ingress-Node RT string 6610:6210,7710:7200, MATCHES the route RT list ['7710:7200', '122:7200']
 ...so should be imported into the VRF on Ingress-Node

Egress VNI carried in the route = 7200

Verified the Route VNI 7200 matches the VNI for the tunnel VLAN 7200

(A) CHECK ROUTE-TABLE IN SOFTWARE:72.1.77.2 (DIP check in VRF Routing Table)

Step1: Checking Software Routing Table

```
Route found for DIP: 72.1.77.2
Number of ECMP Paths: 2
ECMP Route 1 : 72.1.77.0/24 EVPN Ve 7200 200/0
ECMP Route 2 : EVPN Ve 7200 Bi
Routes found in following Modules / Slots / Linecards: [0]
Next-hop Egress Interfaces for this prefix: ['Ve7200', 'Ve7200']
```

(B) CHECK ROUTE-TABLE IN SLOT: 72.1.77.2

Step2: Checking if route is programmed in specific Hardware Slots

```
Number of Slots = 0
Module 0: RIB Route programmed in Hardware : Yes
```

(C) CHECK ROUTE-TABLE IN HARDWARE: 72.1.77.2

Step3: Checking the hardware LPM & NHOP for required slots.....

```
CHECKING HARDWARE FOR MODULE: 0
(i) Module 0: LPM Entry FOUND in Hardware for Prefix 72.1.77.0/24
Module 0 LPM: Egress Interfaces registered in Hardware: ['vlan0.7200', 'vlan0.7200']
Less specific (Subnet) prefix programmed as TRAP in HW ..no further nhop checks
```

LOG-MSG: Reference to NH Entry NOT found in LPM table in Hardware / ASIC, ... possibly due to conversational behavior

CHECK FOR TUNNEL STATES TOWARDS EGRESS NODE:

```
Tunnels associated in SOFTWARE for the vlan 7200 on this Rbridge:
Tunnel 1 = Tu 61441
Tunnel 2 = Tu 61443
```

```
Tunnels associated in HARDWARE for the vlan 7200 on this Rbridge:
Found following tunnels associated with this vlan in HARDWARE
Tunnel 61441 Tunnel 61443
Number of Tunnels programmed in HW are same as in Software
```

```
Total list of Tunnels that are Operationally UP on this Rbridge:
Tunnel Number : 61441 Source IP: 66.66.66.66 Destination IP: 77.77.77.77 Oper State : up
Tunnel Number : 61442 Source IP: 66.66.66.66 Destination IP: 54.54.54.54 Oper State : up
Tunnel Number : 61443 Source IP: 66.66.66.66 Destination IP: 71.71.71.71 Oper State : up
```

***LOG-MSG: All Operationally UP tunnels are NOT associated with the vlan 7200

QUICK GLANCE SUMMARY OF INGRESS NODE FLOW STRUCTURE:

```
FLOW Information:
Flow Forward Type : 13fwd-symmetric
Flow DIP : 72.1.77.2
Flow Ingress Vlan : 6201
Flow Ingress VE VRF : vrf10
Flow Egress Vlan : 7200
Flow Egress VE VRF : vrf10
Flow Egress Vlan VNI : 7200
```

BGP Derived Information:

```

Flow VRF Import L3 RT                : 6610:6210,7710:7200
Flow VRF Import L2 RT                : None
Flow IP Route found in BGP VRF Table  : 72.1.77.0/24
Flow IP Route found in BGP EVPN Table : 72.1.77.0/24
Flow MAC Route found in BGP MAC-VRF Table : N/A
Flow ARP Route found in BGP MAC-VRF Table : N/A
Flow Egress LeafNodes VTEP IP        : 77.77.77.77
Flow SpineNodes IP's                 : ['122.122.122.122', '125.125.125.125']
-----
RIB Derived Information:
Flow IP Route found in RIB            : 72.1.77.0/24
Flow Egress Interface                 : Tu 61441
Flow Egress Node Gateway MAC Address  : ['0005.3365.377b', '0005.3365.3633']
-----
TUNNEL Information: ( Tu 61441 )
Flow Egress Int Tunnel SrcIP          : 66.66.66.66
Flow Egress Int Tunnel DestIP         : 77.77.77.77
-----
HARDWARE Programming:
Flow IP ROUTE L3 LPM Fwding Decision  : ['Trap']
Flow IP ROUTE L3 LPM NH HW Fwding Decision : ['None']
Flow ARP L3 EXM Fwding Decision       : N/A
Flow ARP L3 EXM NH HW Fwding Decision  : N/A
Flow L2 HW Fwding Entry                : []
-----
CHECK ALERTS:
Programming Consistency in SW & HW    : True
Number of Alerts found                 : 0
=====

```

Paste the below string on following Spine neighbor nodes for further tracing the packet path:

Reference Format: `execute-script ipfabric-spineNode-debugv<x>.py <DIP> <DIP-prefix> <ingressNode-bgpPeerIP> <egress-vlan> <fwd-type> <destn-Leaf-Node-IP's> <egressVNI> <MAC-Address> <tunnel-sip> <tunnel-dip> <RT>`

```

Spine Node (122.122.122.122): execute-script ipfabric-spineNode-debugv2.py 72.1.77.2 72.1.77.0/24
67.67.67.67 7200 13fwd-symmetric 77.77.77.77 7200 0005.3365.377b 66.66.66.66 77.77.77.77
6610:6210,7710:7200

```

```

Spine Node (122.122.122.122): execute-script ipfabric-spineNode-debugv2.py 72.1.77.2 72.1.77.0/24
67.67.67.67 7200 13fwd-symmetric 77.77.77.77 7200 0005.3365.3633 66.66.66.66 77.77.77.77
6610:6210,7710:7200

```

```

Spine Node (125.125.125.125): execute-script ipfabric-spineNode-debugv2.py 72.1.77.2 72.1.77.0/24
67.67.67.67 7200 13fwd-symmetric 77.77.77.77 7200 0005.3365.377b 66.66.66.66 77.77.77.77
6610:6210,7710:7200

```

```

Spine Node (125.125.125.125): execute-script ipfabric-spineNode-debugv2.py 72.1.77.2 72.1.77.0/24
67.67.67.67 7200 13fwd-symmetric 77.77.77.77 7200 0005.3365.3633 66.66.66.66 77.77.77.77
6610:6210,7710:7200
=====

```

Beginning with Network OS 7.0.1b, the automation suite has been enhanced to include dataplane packet captures. The suite spawns out a new thread to capture the packets for that specific flow in parallel to checking the health of the node. The packet capture enables to catch conditions that causes blackholing or misrouting of traffic due to an unexpected behavior. Apart from the above, the automation suite flags incrementing CRC errors, IP MTU errors, L2 MTU errors and enables the user to locate and recover them promptly.

In general, the data-path debug automation suite enables the user to pin-point the problem area for that flow in their IP Fabric Network and enable quick resolution of that specific node. The single command with flow information utilizes more than 80 commands and stitches the data to present a summary in as user-consumable format.

Example of Layer 2 Flow switched over VXLAN tunnel with packet capture:

QUICK GLANCE SUMMARY OF INGRESS NODE FLOW STRUCTURE:

```

=====
FLOW Information:
Flow Forward Type                : 12fwd
Flow DIP                          : 10.0.0.81

```

```

Flow Ingress Vlan                : 2001
Flow Ingress VE VRF              : vrf101
Flow Egress Vlan                 : 2001
Flow Egress VE VRF              : vrf101
Flow Egress Vlan VNI            : 2001
-----
BGP Derived Information:
Flow VRF Import L3 RT           : 101:101
Flow VRF Import L2 RT           : 2001
Flow IP Route found in BGP VRF Table : N/A
Flow IP Route found in BGP EVPN Table : N/A
Flow MAC Route found in BGP MAC-VRF Table : N/A
Flow ARP Route found in BGP MAC-VRF Table : N/A
Flow Egress LeafNodes VTEP IP   : 200.2.2.2
Flow SpineNodes IP's            : ['10.11.29.2']
-----
RIB Derived Information:
Flow IP Route found in RIB       : N/A
Flow Egress Interface           : Tu 61441
Flow Resolved MAC Address        : ['0004.0000.0000']
-----
TUNNEL Information: ( Tu 61441 )
Flow Egress Int Tunnel SrcIP     : 200.1.1.1
Flow Egress Int Tunnel DestIP    : 200.2.2.2
Flow Physical Egress Int(s) towards Tunnel DIP : ['Fo 29/0/51 (105/0/13)']
Flow Tunnel DIP Fwding Decn in HW : [['FWD']]
-----
HARDWARE Programming:
Flow IP ROUTE L3 LPM Fwding Decision : [N/A]
Flow IP ROUTE L3 LPM NH HW Fwding Decision : [N/A]
Flow ARP L3 EXM Fwding Decision      : N/A
Flow ARP L3 EXM NH HW Fwding Decision : N/A
Flow L2 HW Fwding Entry              : Present
-----
DATAPLANE information:
Traffic Detected on Interface       : 29/0/37
Traffic Type (tagged / untagged)    : tagged
Traffic Classification              : IPv4 Bridged
Traffic forwarding flags: RTE# FWD, DROP, TRAP : 1 1 0 0
Traffic forwarding decn to multicast : 0
Traffic DROP Reason                 : 0x0
Traffic TRAP Reason                 : 0x0
-----
CHECK ALERTS:
Programming Consistency in SW & HW : True
Number of Alerts found              : 0

```

Example of Symmetric Routed Flow with packet captures:

```

QUICK GLANCE SUMMARY OF INGRESS NODE FLOW STRUCTURE:
=====
FLOW Information:
Flow Forward Type                 : 13fwd-symmetric
Flow DIP                          : 200.107.1.11
Flow Ingress Vlan                 : 3957
Flow Ingress VE VRF              : vrf101
Flow Egress Vlan                 : 7101
Flow Egress VE VRF              : vrf101
Flow Egress Vlan VNI            : 7101
-----
BGP Derived Information:
Flow VRF Import L3 RT           : 101:101
Flow VRF Import L2 RT           : None
Flow IP Route found in BGP VRF Table : 200.107.1.11/32
Flow IP Route found in BGP EVPN Table : 200.107.1.0/24
Flow MAC Route found in BGP MAC-VRF Table : N/A
Flow ARP Route found in BGP MAC-VRF Table : N/A
Flow Egress LeafNodes VTEP IP   : 200.1.1.1
Flow SpineNodes IP's            : ['10.101.223.2']
-----

```



```

RIB Derived Information:
  Flow IP Route found in RIB          : 200.107.1.11/32
  Flow Egress Interface                : Tu 61443
  Flow Egress Node Gateway MAC Address : ['0027.f8c3.dd7d']
-----
TUNNEL Information: ( Tu 61443 )
  Flow Egress Int Tunnel SrcIP        : 200.223.223.223
  Flow Egress Int Tunnel DestIP       : 200.1.1.1
  Flow Physical Egress Int(s) towards Tunnel DIP : ['Fo 223/0/2 (101/1/9)']
  Flow Tunnel DIP Fwding Decn in HW   : [['FWD']]
-----
HARDWARE Programming:
  Flow IP ROUTE L3 LPM Fwding Decision : ['Fwd']
  Flow IP ROUTE L3 LPM NH HW Fwding Decision : [['FWD']]
  Flow ARP L3 EXM Fwding Decision      : N/A
  Flow ARP L3 EXM NH HW Fwding Decision : N/A
  Flow L2 HW Fwding Entry              : Present
-----
DATAPLANE information:
  Traffic Detected on Interface        : 223/0/5:1
  Traffic Type (tagged / untagged)    : tagged
  Traffic Classification               : IPv4 routed
  Traffic forwarding flags: RTE# FWD, DROP, TRAP : 2 1 0 0
  Traffic forwarding decn to multicast : 0
  Traffic DROP Reason                 : 0x0
  Traffic TRAP Reason                 : 0x0
-----
CHECK ALERTS:
  Programming Consistency in SW & HW   : True
  Number of Alerts found                : 0

```

Example of Symmetric Flow flagged for IP MTU violation:

QUICK GLANCE SUMMARY OF INGRESS NODE FLOW STRUCTURE:

```

=====
FLOW Information:
  Flow Forward Type                   : 13fwd-symmetric
  Flow DIP                             : 220.107.1.32
  Flow Ingress Vlan                   : 2023
  Flow Ingress VE VRF                 : vrf103
  Flow Egress Vlan                    : 2022
  Flow Egress VE VRF                 : vrf103
  Flow Egress Vlan VNI                : 2022
-----
BGP Derived Information:
  Flow VRF Import L3 RT               : 103:103
  Flow VRF Import L2 RT               : None
  Flow IP Route found in BGP VRF Table : 220.107.1.32
  Flow IP Route found in BGP EVPN Table : 220.107.1.0/24
  Flow MAC Route found in BGP MAC-VRF Table : N/A
  Flow ARP Route found in BGP MAC-VRF Table : N/A
  Flow Egress LeafNodes VTEP IP       : 200.2.2.2
  Flow SpineNodes IP's                : ['10.11.29.2']
-----
RIB Derived Information:
  Flow IP Route found in RIB          : 220.107.1.32/32
  Flow Egress Interface                : Tu 61441
  Flow Egress Node Gateway MAC Address : ['0027.f8b9.8919']
-----
TUNNEL Information: ( Tu 61441 )
  Flow Egress Int Tunnel SrcIP        : 200.1.1.1
  Flow Egress Int Tunnel DestIP       : 200.2.2.2
  Flow Physical Egress Int(s) towards Tunnel DIP : ['Fo 29/0/51 (105/0/13)']
  Flow Tunnel DIP Fwding Decn in HW   : [['FWD']]
-----
HARDWARE Programming:
  Flow IP ROUTE L3 LPM Fwding Decision : ['Fwd']
  Flow IP ROUTE L3 LPM NH HW Fwding Decision : [['FWD']]
  Flow ARP L3 EXM Fwding Decision      : N/A
  Flow ARP L3 EXM NH HW Fwding Decision : N/A

```

```

Flow L2 HW Fwding Entry : Present
-----
DATAPLANE information:
Traffic Detected on Interface : 29/0/37
Traffic Type (tagged / untagged) : tagged
Traffic Classification : IPv4 routed
Traffic forwarding flags: RTE# FWD, DROP, TRAP : 1 1 0 0
Traffic forwarding decn to multicast : 0
Traffic DROP Reason : 0x0
Traffic TRAP Reason : 0x0
Traffic Errors from Interface Statistics : L2 MTU Violation on TenGigabitEthernet 29/0/37
Traffic Errors from Interface Statistics : Egress IP MTU violation on Fo 29/0/51
-----
CHECK ALERTS:
Programming Consistency in SW & HW : True
Number of Alerts found : 3

```

Example of L2Fwd Flow flagged for ingress tail-drops:

```

QUICK GLANCE SUMMARY OF INGRESS NODE FLOW STRUCTURE:
=====
FLOW Information:
Flow Forward Type : l2fwd
Flow DIP : 10.0.0.82
Flow Ingress Vlan : 2001
Flow Ingress VE VRF : vrf101
Flow Egress Vlan : 2001
Flow Egress VE VRF : vrf101
Flow Egress Vlan VNI : 2001
-----
BGP Derived Information:
Flow VRF Import L3 RT : 101:101
Flow VRF Import L2 RT : :2001
Flow IP Route found in BGP VRF Table : N/A
Flow IP Route found in BGP EVPN Table : N/A
Flow MAC Route found in BGP MAC-VRF Table : N/A
Flow ARP Route found in BGP MAC-VRF Table : N/A
Flow Egress LeafNodes VTEP IP : 200.1.1.1
Flow SpineNodes IP's : ['20.31.129.2']
-----
RIB Derived Information:
Flow IP Route found in RIB : N/A
Flow Egress Interface : Tu 61441
Flow Resolved MAC Address : ['0003.0000.0000']
-----
TUNNEL Information: ( Tu 61441 )
Flow Egress Int Tunnel SrcIP : 222.3.3.3
Flow Egress Int Tunnel DestIP : 200.1.1.1
Flow Physical Egress Int(s) towards Tunnel DIP : ['Te 31/0/42 (129/2/33)']
Flow Tunnel DIP Fwding Decn in HW : [['FWD']]
-----
HARDWARE Programming:
Flow IP ROUTE L3 LPM Fwding Decision : [N/A]
Flow IP ROUTE L3 LPM NH HW Fwding Decision : [N/A]
Flow ARP L3 EXM Fwding Decision : N/A
Flow ARP L3 EXM NH HW Fwding Decision : N/A
Flow L2 HW Fwding Entry : Present
-----
DATAPLANE information:
Traffic Detected on Interface : 31/0/29
Traffic Type (tagged / untagged) : tagged
Traffic Classification : IPv4 Bridged
Traffic forwarding flags: RTE# FWD, DROP, TRAP : 0 1 0 0
Traffic forwarding decn to multicast : 0
Traffic DROP Reason : 0x0
Traffic TRAP Reason : 0x0
Traffic Errors from Interface Statistics : Ingress Tail-drops on TenGigabitEthernet 31/0/29
-----
CHECK ALERTS:

```

```

Programming Consistency in SW & HW           : True
Number of Alerts found                       : 2

```

Example of Asymmetric Routed flow with packet captures and unexpected forwarding error:

NOTE

This example highlights the fact that prefix is properly programmed in hardware, however, while forwarding of traffic, it encounters unexpected operational values and thus in turn drops traffic.

QUICK GLANCE SUMMARY OF INGRESS NODE FLOW STRUCTURE:

```

=====
FLOW Information:
  Flow Forward Type           : 13fwd-asymmetric
  Flow DIP                   : 201.107.1.101
  Flow Ingress Vlan          : 2001
  Flow Ingress VE VRF        : vrf101
  Flow Egress Vlan           : 2002
  Flow Egress VE VRF         : vrf101
  Flow Egress Vlan VNI       : 2002
-----
BGP Derived Information:
  Flow VRF Import L3 RT      : 101:101
  Flow VRF Import L2 RT      : 2002
  Flow IP Route found in BGP VRF Table : N/A
  Flow IP Route found in BGP EVPN Table : N/A
  Flow MAC Route found in BGP MAC-VRF Table : N/A
  Flow ARP Route found in BGP MAC-VRF Table : False
  Flow Egress LeafNodes VTEP IP : 200.2.2.2
  Flow SpineNodes IP's       : ['10.11.29.2']
-----
RIB Derived Information:
  Flow IP Route found in RIB   : 201.107.1.0/24
  Flow Egress Interface        : Tu 61441
  Flow Resolved MAC Address    : ['None']
-----
TUNNEL Information: ( Tu 61441 )
  Flow Egress Int Tunnel SrcIP : 200.1.1.1
  Flow Egress Int Tunnel DestIP : 200.2.2.2
  Flow Physical Egress Int(s) towards Tunnel DIP : ['Fo 29/0/51 (105/0/13)']
  Flow Tunnel DIP Fwding Decn in HW : [['FWD']]
-----
HARDWARE Programming:
  Flow IP ROUTE L3 LPM Fwding Decision : ['Trap']
  Flow IP ROUTE L3 LPM NH HW Fwding Decision : ['None']
  Flow ARP L3 EXM Fwding Decision : N/A
  Flow ARP L3 EXM NH HW Fwding Decision : N/A
  Flow L2 HW Fwding Entry : N/A
-----
DATAPLANE information:
  Traffic Detected on Interface : 29/0/37
  Traffic Type (tagged / untagged) : tagged
  Traffic Classification : IPv4 routed
  Traffic forwarding flags: RTE# FWD, DROP, TRAP : 1 0 1 1
  Traffic forwarding decn to multicast : 0
  Traffic DROP Reason : 0x7c
  Traffic TRAP Reason : 0x53
  Traffic Errors from Interface Statistics : Unknown RTE Drops on interface 29/0/37
  Traffic Errors from Interface Statistics : Unknown RTE Traps on interface 29/0/37
-----
CHECK ALERTS:
  Programming Consistency in SW & HW           : True
  Number of Alerts found                       : 6

```


TACACS+ Accounting Exceptions

- [TACACS+ command-accounting limitations](#)..... 85
- [Unsupported Network OS command line interface commands](#)..... 85

TACACS+ command-accounting limitations

TACACS+ command accounting is subject to the following limitations:

- A number of Network OS CLI commands are not supported by TACACS+ accounting; refer to [Unsupported Network OS command line interface commands](#) on page 85 for a listing of unsupported operational and configuration commands.

Unsupported Network OS command line interface commands

The following table lists the Network OS CLI commands that are not supported in privileged EXEC mode.

TABLE 11 Unsupported Network OS CLI commands in privileged EXEC mode

Command name	Command Description
<code>cipherset</code>	Configures FIPS-compliant secure ciphers for LDAP and SSH.
<code>clear</code>	Clears the specified parameter.
<code>clear arp</code>	Clears Address Resolution Protocol (ARP) configuration data.
<code>clear counters</code>	Clears statistics from the switch.
<code>clear dot1x</code>	Clears IEEE 802.1X Port-Based Access Control configuration data.
<code>clear fcoe</code>	Clears FCoE configuration data.
<code>clear ip</code>	Clears Internet Protocol (IP) configuration data.
<code>clear lacp</code>	Clears Link Aggregation Control Protocol (LACP) configuration data.
<code>clear lldp</code>	Clears Link Layer Discovery Protocol (LLDP) configuration data.
<code>clear mac-address-table</code>	Clears the MAC address table.
<code>clear mcagt</code>	Clears the MCAGT agent.
<code>clear policy-map-counters</code>	Clears the policy map counters.
<code>clear sflow</code>	Clears sFlow configuration data.
<code>clear spanning-tree</code>	Clears Spanning Tree Protocol (STP) configuration data.
<code>clear vrrp</code>	Clears Virtual Router Redundancy Protocol (VRRP) configuration data.
<code>configure</code>	Configures access mode.
<code>copy</code>	Copies data.
<code>debug</code>	Sets debugging options.
<code>delete</code>	Delete a specified file.
<code>dir</code>	Displays a directory listing.
<code>dot1x</code>	Executes IEEE 802.1X Port-Based Access Control options.
<code>exit</code>	Exits to the top level and optionally runs a command.
<code>fips</code>	Executes FIPS-related operations.

TABLE 11 Unsupported Network OS CLI commands in privileged EXEC mode (continued)

Command name	Command Description
help	Provides help information.
history	Configures the size of the history log.
logout	Terminates the current login session.
mac-rebalance	Rebalances MAC on a port channel
ping	Executes the ping command.
quit	Terminates the current session.
rename	Renames a file.
reload	Reboots the system.
resequence	Re-orders a list.
send	Sends a message to terminal of one or all users.
terminal	Configures terminal properties.
show arp	Displays the Address Resolution Protocol (ARP) configuration.
show bpdu-drop	Displays the Bridge Protocol Data Unit (BPDU) drop configuration.
show cee maps	Displays CEE maps.
show cipherset	Displays ciphers for LDAP and SSH.
show cli	Displays CLI session parameters.
show clock	Displays the date and time settings.
show diag	Displays diagnostic information.
show dot1x	Displays IEEE 802.1X Port-Based Access Control configuration data.
show edge-loop-detection globals	Displays system-wide Edge-Loop-Detection status information.
show fcoe login	Displays the FCoE CNA Login information.
show file	Displays the contents of a file.
show history	Displays command history.
show interface	Displays interface status and configuration.
show ip	Displays Internet Protocol (IP) information.
show lacp counter	Displays Link Aggregation Control Protocol (LACP) counters.
show lldp	Displays Link Layer Discovery Protocol (LLDP) configuration data
show monitor	Displays interface status and configuration.
show netconf-state	Displays NETCONF statistics.
show ntp	Displays the active NTP server.
show parser dump	Displays a parser dump.
show policy-map	Displays the configured rate-limiting policy maps.
show port	Displays port parameters.
show port-channel	Displays the port-channel configuration.
show port-profile	Displays the port profile configuration
show qos	Display the Quality of Service (QoS) configuration.
show running-config	Displays the running configuration.
show sflow	Displays the sFlow configuration.
show spanning-tree	Displays the Spanning Tree Protocol configuration.
show ssm	Displays the switch services subsystem.
show startup-db	Displays the startup configuration.

TABLE 11 Unsupported Network OS CLI commands in privileged EXEC mode (continued)

Command name	Command Description
show storm-control	Displays storm control configuration.
show statistics	Displays accounting information.
show system	Displays runtime system information.
show rmon	Displays the Remote Monitoring Protocol (RMON) configuration.
show vcs	Displays VCS information.
show vlan	Displays the VLAN configuration
show mac-address-table	Displays the MAC address table.
show startup-config	Displays the contents of the startup-configuration file.
show zoning	Displays zoning information.
traceroute	Executes the traceroute command.

The following table lists the Network OS CLI commands that are not supported in global configuration mode.

TABLE 12 Unsupported Network OS CLI commands in global configuration mode

Command name	Command Description
abort	Aborts the current configuration session.
diag	Manages diagnostic commands.
do	Executes an operational command while in global configuration mode.
end	Terminates the current configuration session.
exit	Exits from the current mode.
help	Provides help information.
pwd	Displays the current mode path.
service	Performs password encryption services.
top	Exits to the top level and optionally runs a command.
no vlan	Disables VLAN configuration.

Supported NTP Regions and Time Zones

- Africa..... 89
- America..... 89
- Antarctica..... 91
- Arctic..... 91
- Asia..... 91
- Atlantic..... 92
- Australia..... 92
- Europe..... 92
- Indian..... 93
- Pacific..... 93

Africa

The table below lists region and city time zones supported in the Africa region.

Africa/Luanda	Africa/Banjul	Africa/Mogadishu
Africa/Ouagadougou	Africa/Conakry	Africa/Sao_Tome
Africa/Bujumbura	Africa/Malabo	Africa/Mbabane
Africa/Porto-Novo	Africa/Bissau	Africa/Ndjamena
Africa/Gaborone	Africa/Nairobi	Africa/Lome
Africa/Kinshasa	Africa/Monrovia	Africa/Tunis
Africa/Lubumbashi	Africa/Maseru	Africa/Dar_es_Salaam
Africa/Bangui	Africa/Tripoli	Africa/Kampala
Africa/Brazzaville	Africa/Casablanca	Africa/Johannesburg
Africa/Abidjan	Africa/Bamako	Africa/Lusaka
Africa/Douala	Africa/Nouakchott	Africa/Harare
Africa/Djibouti	Africa/Blantyre	
Africa/Algiers	Africa/Maputo	
Africa/Cairo	Africa/Windhoek	
Africa/El_Aaiun	Africa/Niamey	
Africa/Asmara	Africa/Lagos	
Africa/Ceuta	Africa/Kigali	
Africa/Addis_Ababa	Africa/Khartoum	
Africa/Libreville	Africa/Freetown	
Africa/Accra	Africa/Dakar	

America

The table below lists region and city time zones supported in the America region.

America/Antigua	America/Guatemala	America/Edmonton
America/Anguilla	America/Guyana	America/Cambridge_Bay
America/Curacao	America/Tegucigalpa	America/Yellowknife
America/Argentina/Buenos_Aires	America/Port-au-Prince	America/Inuvik
America/Argentina/Cordoba	America/Guadeloupe	America/Dawson_Creek
America/Argentina/San_Luis	America/Jamaica	America/Vancouver
America/Argentina/Jujuy	America/St_Kitts	America/Whitehorse
America/Argentina/Tucuman	America/Cayman	America/Thunder_Bay
America/Argentina/Catamarca	America/St_Lucia	America/Iqaluit
America/Argentina/La_Rioja	America/Marigot	America/Pangnirtung
America/Argentina/San_Juan	America/Adak	America/Resolute
America/Argentina/Mendoza	America/Martinique	America/Rankin_Inlet
America/Argentina/Rio_Gallegos	America/Montserrat	America/Winnipeg
America/Argentina/Ushuaia	America/Mexico_City	America/Rainy_River
America/Aruba	America/Cancun	America/Regina
America/Barbados	America/Merida	America/Montevideo
America/St_Barthelemy	America/Monterrey	America/St_Vincent
America/La_Paz	America/Mazatlan	America/Caracas
America/Noronha	America/Chihuahua	America/Tortola
America/Belem	America/Hermosillo	America/St_Thomas
America/Fortaleza	America/Tijuana	America/New_York
America/Recife	America/Managua	America/Detroit
America/Araguaina	America/Panama	America/Kentucky/Monticello
America/Maceio	America/Lima	America/Indiana/Indianapolis
America/Bahia	America/Miquelon	America/Indiana/Vincennes
America/Sao_Paulo	America/Puerto_Rico	America/Indiana/Knox
America/Campo_Grande	America/Asuncion	America/Indiana/Winamac
America/Cuiaba	America/Paramaribo	America/Indiana/Marengo
America/Santarem	America/El_Salvador	America/Indiana/Vevay
America/Porto_Velho	America/Grand_Turk	America/Chicago
America/Boa_Vista	America/Swift_Current	America/Indiana/Tell_City
America/Manaus	America/Dawson	America/Indiana/Petersburg
America/Eirunepe	America/Santiago	America/Menominee
America/Rio_Branco	America/Bogota	America/North_Dakota/Center
America/Nassau	America/Costa_Rica	America/North_Dakota/New_Salem
America/Belize	America/Havana	America/Denver
America/St_Johns	America/Dominica	America/Boise
America/Halifax	America/Santo_Domingo	America/Shiprock
America/Glace_Bay	America/Guayaquil	America/Phoenix

America/Moncton	America/Grenada	America/Los_Angeles
America/Goose_Bay	America/Cayenne	America/Anchorage
America/Blanc-Sablon	America/Godthab	America/Juneau
America/Montreal	America/Danmarkshavn	America/Yakutat
America/Toronto	America/Scoresbysund	America/Nome
America/Nipigon	America/Thule	America/Port_of_Spain

Antarctica

The table below lists region and city time zones supported in the Antarctica region.

Antarctica/McMurdo	Antarctica/Mawson	Antarctica/Vostok
Antarctica/South_Pole	Antarctica/Davis	Antarctica/DumontDURville
Antarctica/Rothera	Antarctica/Casey	Antarctica/Syowa

Arctic

The table below lists region and city time zones supported in the Arctic region.

Arctic/Longyearbyen		
---------------------	--	--

Asia

The table below lists region and city time zones supported in the Asia region.

Asia/Dubai	Asia/Tokyo	Asia/Gaza
Asia/Kabul	Asia/Bishkek	Asia/Qatar
Asia/Yerevan	Asia/Phnom_Penh	Asia/Yekaterinburg
Asia/Baku	Asia/Pyongyang	Asia/Omsk
Asia/Dhaka	Asia/Seoul	Asia/Novosibirsk
Asia/Bahrain	Asia/Kuwait	Asia/Krasnoyarsk
Asia/Brunei	Asia/Almaty	Asia/Irkutsk
Asia/Thimphu	Asia/Qyzylorda	Asia/Yakutsk
Asia/Shanghai	Asia/Aqtobe	Asia/Vladivostok
Asia/Harbin	Asia/Aqtau	Asia/Sakhalin
Asia/Chongqing	Asia/Oral	Asia/Magadan
Asia/Urumqi	Asia/Vientiane	Asia/Kamchatka
Asia/Kashgar	Asia/Beirut	Asia/Anadyr
Asia/Nicosia	Asia/Colombo	Asia/Riyadh
Asia/Tbilisi	Asia/Rangoon	Asia/Singapore

Asia/Hong_Kong	Asia/Ulaanbaatar	Asia/Damascus
Asia/Jakarta	Asia/Hovd	Asia/Bangkok
Asia/Pontianak	Asia/Choibalsan	Asia/Dushanbe
Asia/Makassar	Asia/Macau	Asia/Dili
Asia/Jayapura	Asia/Kuala_Lumpur	Asia/Ashgabat
Asia/Jerusalem	Asia/Kuching	Asia/Taipei
Asia/Kolkata	Asia/Katmandu	Asia/Samarkand
Asia/Baghdad	Asia/Muscat	Asia/Tashkent
Asia/Tehran	Asia/Manila	Asia/Ho_Chi_Minh
Asia/Amman	Asia/Karachi	Asia/Aden
		Asia/Srednekolymsk

Atlantic

The table below lists region and city time zones supported in the Atlantic region.

Atlantic/Bermuda	Atlantic/Faroe	Atlantic/Azores
Atlantic/Cape_Verde	Atlantic/South_Georgia	Atlantic/St_Helena
Atlantic/Canary	Atlantic/Reykjavik	
Atlantic/Stanley	Atlantic/Madeira	

Australia

The table below lists region and city time zones supported in the Australia region.

Australia/Lord_Howe	Australia/Sydney	Australia/Darwin
Australia/Hobart	Australia/Brisbane	Australia/Perth
Australia/Currie	Australia/Lindeman	Australia/Eucla
Australia/Melbourne	Australia/Adelaide	

Europe

The table below lists region and city time zones supported in the Europe region.

Europe/Andorra	Europe/Gibraltar	Europe/Warsaw
Europe/Tirane	Europe/Athens	Europe/Lisbon
Europe/Vienna	Europe/Zagreb	Europe/Bucharest
Europe/Mariehamn	Europe/Budapest	Europe/Belgrade
Europe/Sarajevo	Europe/Dublin	Europe/Kaliningrad
Europe/Brussels	Europe/Isle_of_Man	Europe/Moscow

Europe/Sofia	Europe/Rome	Europe/Volgograd
Europe/Minsk	Europe/Jersey	Europe/Samara
Europe/Zurich	Europe/Vaduz	Europe/Stockholm
Europe/Prague	Europe/Vilnius	Europe/Ljubljana
Europe/Berlin	Europe/Luxembourg	Europe/Bratislava
Europe/Copenhagen	Europe/Riga	Europe/San_Marino
Europe/Tallinn	Europe/Monaco	Europe/Istanbul
Europe/Madrid	Europe/Chisinau	Europe/Kiev
Europe/Helsinki	Europe/Podgorica	Europe/Uzhgorod
Europe/Paris	Europe/Skopje	Europe/Zaporozhye
Europe/London	Europe/Malta	Europe/Simferopol
Europe/Guernsey	Europe/Amsterdam	Europe/Vatican
Europe/Oslo		

Indian

The table below lists region and city time zones supported in the Indian region.

Indian/Cocos	Indian/Antananarivo	Indian/Mahe
Indian/Christmas	Indian/Mauritius	Indian/Kerguelen
Indian/Chagos	Indian/Maldives	Indian/Mayotte
Indian/Comoro	Indian/Reunion	

Pacific

The table below lists region and city time zones supported in the Pacific region.

Pacific/Pago_Pago	Pacific/Kwajalein	Pacific/Palau
Pacific/Rarotonga	Pacific/Saipan	Pacific/Guadalcanal
Pacific/Easter	Pacific/Noumea	Pacific/Fakaofu
Pacific/Galapagos	Pacific/Norfolk	Pacific/Tongatapu
Pacific/Fiji	Pacific/Nauru	Pacific/Funafuti
Pacific/Truk	Pacific/Niue	Pacific/Johnston
Pacific/Ponape	Pacific/Auckland	Pacific/Midway
Pacific/Kosrae	Pacific/Chatham	Pacific/Wake
Pacific/Guam	Pacific/Tahiti	Pacific/Honolulu
Pacific/Tarawa	Pacific/Marquesas	Pacific/Efate
Pacific/Enderbury	Pacific/Gambier	Pacific/Wallis
Pacific/Kiritimati	Pacific/Port_Moresby	Pacific/Apia
Pacific/Majuro	Pacific/Pitcairn	