

Readme for Network OS v6.0.2

Appendix for Administrator's Guide

FASTFIND LINKS

[Product Version](#)

[Document Organization](#)

[Getting Help](#)

[Contents](#)

© 2016 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd.

Hitachi, Ltd., reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. This document contains the most current information available at the time of publication. When new or revised information becomes available, this entire document will be updated and distributed to all registered users.

Some of the features described in this document might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Data Systems Corporation at <https://portal.hds.com>.

Notice: Hitachi, Ltd., products and services can be ordered only under the terms and conditions of the applicable Hitachi Data Systems Corporation agreements. The use of Hitachi, Ltd., products is governed by the terms of your agreements with Hitachi Data Systems Corporation.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

Archivas, Essential NAS Platform, HiCommand, Hi-Track, ShadowImage, Tagmaserve, Tagmasoft, Tagmasolve, Tagmastore, TrueCopy, Universal Star Network, and Universal Storage Platform are registered trademarks of Hitachi Data Systems Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

Open source software components provided with Programs are licensed to you under the terms of the applicable license agreements included with such open source software components. HITACHI provides you with open source licensing information, documentations or corresponding source files under the terms of the license, such as GNU General Public License (GPL), which says that the distributors must distribute the source code. To get such information, contact your reseller with the software version.



Contents

Preface	i
Intended Audience	ii
Product Version	ii
Release Notes.....	ii
Document Organization	ii
Document Conventions.....	iii
Getting Help	iv
Comments.....	iv
What is NEW in Network OS v6.0.2	1-1
Software features	1-2
Improvement and Modified software features for Network OS v6.0.2	1-2
Deprecated software features	1-2
CLI changes	1-2
New Commands for Network OS v6.0.2	1-2
Modified Commands for Network OS v6.0.2	1-2
Deprecated commands.....	1-3
API changes	1-3
Newly supported standards and RFCs	1-3
Introduction of Logical Chassis Cluster mode.....	1-6
Fabric cluster mode.....	1-6
Logical chassis cluster mode.....	1-7
Transitioning between modes	1-7
Prerequisites of Network OS v6.0.2	2-1
Supported Optics and Option Cables.....	2-2
Optional Licensed Software.....	2-2
Software Upgrade and Downgrade	2-3
Migration Path	2-3
Firmware Installation	2-4

Upgrading to this Release	2-5
Configuring and managing switches.....	2-5
Configuring Ethernet management interfaces	2-5
Configuring a switch in logical chassis cluster mode	2-7
Scalability	2-13
Compatibility.....	2-16
Documentation Updates.....	2-16
IMPORTANT NOTES	2-17
Command Line Interface.....	2-17
Platform	2-19
Licensing	2-20
VCS	2-20
Logical Chassis.....	2-21
Brocade Trunks.....	2-22
Breakout Interfaces.....	2-22
Restrictions for Ports in 1G Mode.....	2-22
vLAG.....	2-22
Virtual IP Address Support	2-23
Security, Management ACLs, Authentication, Authorization	2-23
SPAN & RSPAN	2-24
MAC Learning Considerations in VCS.....	2-24
PVLAN.....	2-25
UDLD	2-25
STP/DiST	2-25
Edge Loop Detection (ELD)	2-26
Long Distance ISL Ports.....	2-26
AMPP and Port-Profiles	2-27
vCenter	2-28
QoS	2-29
FCoE.....	2-30
FlexPorts	2-30
Fibre Channel	2-31
Access Gateway	2-31
ND/RA.....	2-32
BFD	2-32
VRRP	2-32
Fabric Virtual Gateway (FVG)	2-33
OSPF v3	2-33
BGP	2-33
ACL.....	2-33
L2/L3 Multicast	2-34
VRF	2-34
BGP-VRF	2-35
Policy-based Routing (PBR).....	2-35

Inter-VRF Leaking (Static)	2-35
DHCP IP Helper	2-36
Dynamic ARP Inspection (DAI).....	2-36
DHCP-based Firmware download (DAD-DHCP Automatic Deployment) .	2-36
Link State Tracking	2-37
OpenFlow.....	2-37
Auto QoS for NAS	2-38
REST API	2-39
NetConf	2-39
VXLAN Gateway for VMware NSX	2-39
TCAM Profiles	2-40
Management VRF	2-40
Conversational MAC Learning.....	2-41
System level Flowbased QoS.....	2-41
Port level Flowbased QoS	2-41
Non-trivial Merge	2-41
Interoperability	2-42
MAPS.....	2-43
Miscellaneous	2-43

Bug Fixes and Known Issues	3-1
Bug Fixes	3-1
Known Issues	3-1



Preface

This document describes how to use the Network OS for Brocade 10Gbps DCB switch module for ComputeBlade 2500.

This preface includes the following information:

- [Intended Audience](#)
- [Product Version](#)
- [Release Notes](#)
- [Document Organization](#)
- [Document Conventions](#)
- [Getting Help](#)
- [Comments](#)

Notice: The use of the Network OS for Brocade 10Gbps DCB switch module and all other Hitachi Data Systems products are governed by the terms of your agreement(s) with Hitachi Data Systems.

Intended Audience

This document is intended for the personnel who are involved in planning, managing, and performing the tasks to prepare your site for Compute Blade installation and to install the same.

This document assumes the following:

- The reader has a background in hardware installation of compute systems.
- The reader is familiar with the location where the Compute Blade will be installed, including knowledge of physical characteristics, power systems and specifications, and environmental specifications.

Product Version

This document revision applies to Network OS for Brocade 10Gbps DCB switch module version v6.0.2

Release Notes

Release notes contain requirements and more recent product information that may not be fully described in this manual. Be sure to review the release notes before installation.

Document Organization

The following table provides an overview of the contents and organization of this document. Click the [chapter title](#) in the left column to go to that chapter. The first page of each chapter provides links to the sections in that chapter.

Chapter	Description
Chapter 1. What is NEW in Network OS v6.0.2	Describes the overview of Network OS 6.0.2.
Chapter 2. Prerequisites of Network OS v6.0.2	Describes the prerequisites of Network OS 6.0.2.
Chapter 3. Bug Fixes and Known Issues	Provides the information of Bug fixes and Known issues for Network OS 6.0.2.

Document Conventions

This document uses the following typographic conventions:

Convention	Description
Bold	Indicates text on a window, other than the window title, including menus, menu options, fields, and labels. Example: Click OK .
<i>Italic</i>	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: <i>copy source-file target-file</i> Note: Angled brackets (< >) are also used to indicate variables.
screen/code	Indicates text that is displayed on screen or entered by the user. Example: # <code>pairdisplay -g oradb</code>
< > angled brackets	Indicates a variable, which is a placeholder for actual text provided by the user or system. Example: # <code>pairdisplay -g <group></code> Note: Italic font is also used to indicate variables.
[] square brackets	Indicates optional values. Example: [a b] indicates that you can choose a, b, or nothing.
{ } braces	Indicates required or expected values. Example: { a b } indicates that you must choose either a or b.
vertical bar	Indicates that you have a choice between two or more options or arguments. Examples: [a b] indicates that you can choose a, b, or nothing. { a b } indicates that you must choose either a or b.
<u>underline</u>	Indicates the default value. Example: [<u>a</u> b]

This document uses the following icons to draw attention to information:

Icon	Meaning	Description
	WARNING	This indicates the presence of a potential risk that might cause death or severe injury.
	CAUTION	This indicates the presence of a potential risk that might cause relatively mild or moderate injury.
	NOTICE	This indicates the presence of a potential risk that might cause severe damage to the equipment and/or damage to surrounding properties.
	Note	This indicates notes not directly related to injury or severe damage to equipment.
	Tip	This indicates advice on how to make the best use of the equipment.

Getting Help

If you purchased this product from an authorized HDS reseller, contact that reseller for support. For the name of your nearest HDS authorized reseller, refer to the HDS support web site for locations and contact information. To contact the Hitachi Data Systems Support Center, please visit the HDS website for current telephone numbers and other contact information:
<http://support.hds.com>.

Before calling the Hitachi Data Systems Support Center, please provide as much information about the problem as possible, including:

- The circumstances surrounding the error or failure.
- The exact content of any error message(s) displayed on the host system(s).

Comments

Please send us your comments on this document: doc.comments@hds.com. Include the document title, number, and revision, and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Data Systems Corporation. **Thank you!**

What is NEW in Network OS v6.0.2

Brocade Network Operating System (Network OS) v6.0.2 is an update of the release of v6.0.1a1 new to deliver the bug fixes on Brocade 10G DCB switch module for Hitachi CB2500 (GG-BE4LSW2X1).

In addition, this release also improves to start supporting some features.

This chapter summarizes the features supported on this release in terms from the difference from the previous NOS release (NOS 6.0.1a1).

- [Software Features](#)
- [CLI Changes](#)
- [API changes](#)
- [Newly supported standards and RFCs](#)
- [Introduction of Logical Chassis Cluster mode](#)

Software features

The following section lists new, modified, and deprecated software features for this release. For information about which platforms support these features, refer to the NOS Feature support Matrix.

Improvement and Modified software features for Network OS v6.0.2

- New SNMP MIB to reflect the VCS specific details
- Support for Telnet and SSH for non-default VRF's
- Support for SNMP in non-default VRF's
- Logical Chassis Cluster mode
- VXLAN Gateway for VMware NSX
- Flexport on QSFP
- Enhancements in SNMP MIB functionality
- Support for configuring source-interface for SNMP traps & additional traps
- FCoE source interface improvement.

Deprecated software features

None

CLI changes

The following section lists new, modified, and deprecated commands for this release:

New Commands for Network OS v6.0.2

The following commands are new in this patch release:

- None

Modified Commands for Network OS v6.0.2

The following commands have been modified for this patch release:

- fcoeport ns-ip-registration
- snmp-server community

- snmp-server host
- snmp-server v3host

Deprecated commands

The following commands have been deprecated beginning with this release:

- None

API changes

Network OS follows the YANG model for CLI and NetConf/REST API. Hence relevant changes in above CLI Changes will get mirrored in API Changes as well.

Newly supported standards and RFCs

The following section lists RFCs and other standards newly supported in this release:

- None

This software generally conforms to Ethernet standards in a manner consistent with accepted engineering practices and procedures. In certain cases, Brocade might add proprietary supplemental functions to those specified in the standards, or choose to implement modifications to the standards for performance or behavioral improvements.

- IEEE 802.1D Spanning Tree Protocol
- IEEE 802.1s Multiple Spanning Tree
- IEEE 802.1w Rapid reconfiguration of Spanning Tree Protocol
- IEEE 802.3ad Link Aggregation with LACP
- IEEE 802.3ae 10G Ethernet
- IEEE 802.1Q VLAN Tagging
- IEEE 802.1p Class of Service Prioritization and Tagging
- IEEE 802.1v VLAN Classification by Protocol and Port
- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
- IEEE 802.3x Flow Control (Pause Frames)

The following draft versions of the Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) Standards are also supported on VDX products:

- IEEE 802.1Qbb Priority-based Flow Control
- IEEE 802.1 DCB Capability Exchange Protocol
(Proposed under the DCB Task Group of IEEE 802.1 Working Group)
- FC-BB-5 FCoE (Rev 2.0)

The VDX products conform to the following Internet IETF RFCs:

- RFC 2865 Remote Authentication Dial In User Service (RADIUS)
- RFC 1112 IGMP
- RFC 2236 IGMPv2
- RFC4601 PIM-SM
- RFC2131 DHCP
- RFC 2571 Architecture for Describing SNMP Framework
- RFC 3176 sFlow
- RFC 1157 SNMPv1/v2c
- RFC4510 Lightweight Directory Access Protocol (LDAP)
- RFC 3768 Virtual Router Redundancy Protocol (VRRP)
- RFC 2328 OSPF Version 2
- RFC 1587 OSPF NSSA Option
- RFC 3101 OSPF Not-So-Stubby-Area (NSSA) Option
- RFC 1765 OSPF Database Overflow
- RFC 2154 OSPF with Digital Signatures (MD-5 Support)
- RFC 3137 OSPF Stub Router advertisement
- RFC 2460 IPv6
- RFC 5340 OSPF for IPv6
- RFC 3623 Graceful OSPF Restart
- RFC 5187 OSPFv3 Graceful Restart (Helper Only)

- RFC 4271 A Border Gateway Protocol 4 (BGP-4)
- RFC 1745 BGP - OSPF Interactions
- RFC 1997 BGP Communities Attributes
- RFC 2385 Protection of BGP Sessions via the TCP MD5 Signature options
- RFC 2439 BGP Route Flap Dampening
- RFC 4456 BGP Route Reflection
- RFC 5492 Capabilities Advertisement with BGP-4
- RFC 3065 Autonomous System Confederations for BGP
- RFC 2858 Multiprotocol Extensions for BGP-4
- RFC 2918 Route Refresh Capability for BGP-4
- RFC 4724 Graceful Restart Mechanism for BGP
- RFC 5798 VRRP Version 3 for IPv4 and IPv6
- RFC 4541 MLDv1 Snooping
- RFC 6987 OSPFv3 Stub Router Advertisement (R-bit in Router LSA not supported)

The Brocade VDX 6740x, VDX 2740 and VDX 2746 products conform to the following Fibre Channel standards:

- FC-GS-5 ANSI INCITS 427:2007 (includes the following)
- FC-GS-4 ANSI INCITS 387: 2004
- FC-SP-2 INCITS 496-2012 (AUTH-A, AUTH-B1 only)
- FC-DA INCITS TR-36: 2004 (includes the following)
- FC-FLA INCITS TR-20: 1998
- FC-PLDA INCITS TR-19: 1998
- FC-MI-2 ANSI/INCITS TR-39-2005
- FC-PI INCITS 352: 2002
- FC-PI-2 INCITS 404: 2005
- FC-PI-4 INCITS 1647-D, revision 7.1 (under development)

- FC-FS-2 ANSI/INCITS 424:2006 (includes the following)
- FC-FS INCITS 373: 2003
- FC-LS INCITS 433: 2007
- MIB-FA INCITS TR-32: 2003

Introduction of Logical Chassis Cluster mode

This release supports two operation modes, one is the Fabric Cluster mode (FC mode), the other is the Logical Chassis mode (LC mode). The logical chassis cluster mode is the operation mode that is newly supported from this release.

- **Fabric Cluster mode** - The second of two types of "VCS" modes for a switch. In this mode, the data path for nodes is distributed, but the configuration path is not distributed. Each node keeps its configuration database independently. Refer to chapter of Fabric cluster mode for more information.
- **Logical Chassis Cluster mode** - One of two types of "VCS" modes for a switch. This mode requires Network OS 6.0.2 or later. In this mode, both the data and configuration paths are distributed. The entire cluster is configured from the principal node. Refer to chapter of Logical chassis cluster mode for more information.

When a new switch boots up, the switch enters fabric cluster mode by default.

ATTENTION: Each time you change the Brocade VCS Fabric configuration, the switch resets to the default configuration and reboots automatically. Make sure to save the configuration before you change it.

ATTENTION: The switch module has different Ethernet paths for the management interface depending on the switch mode.

- When in FC mode, the management interface is not configurable from the switch module console but only from the Hitachi Compute Blade management module.
- When in LC mode, you must connect an Ethernet cable to the RJ45 (eth1) port on the faceplate to a network for management.

Fabric cluster mode

By default, the switch will boot up in fabric cluster mode and will attempt to form Inter-Switch Links (ISLs):

If the chassis is not connected to another switch, it forms a "single node VCS fabric." This means that the chassis operates as a standalone system, but the operational mode is always VCS-enabled. You cannot disable the VCS mode on any of the models listed above.

Logical chassis cluster mode

Logical chassis cluster mode characteristics

The following are the main characteristics of logical chassis cluster mode:

- The maximum number of nodes supported in a logical chassis cluster is 48.
- This mode supports in-band management (through eth0 on management modules) over virtual Ethernet (VE) interfaces.
- In-Band Management is supported in Logical Chassis mode on VDX devices.
- Physical connectivity requirements for logical chassis cluster deployment are the same as those for fabric cluster deployment.
- A single global configuration exists across all nodes, while each node can contain its unique local configuration. However, each node contains the local configuration information for all other nodes in the cluster.
- When an RBridge is rejoining the logical chassis cluster, the interface-level configuration is reset to the default values.
- Global and local configurations for the entire logical chassis cluster are performed from one node — the principal node only.
- Startup configurations are not maintained by the cluster; each node preserves its running configuration.
- A logical chassis cluster can be transitioned into a fabric cluster while preserving configurations, if you follow the steps provided later in this section
- An existing fabric cluster can be transitioned into a logical chassis cluster while preserving configurations, if you follow the steps provided later in this section
- Cluster-wide firmware upgrades can be performed.
- Cluster-wide supportSave can be performed.

For detail about the Logical Chassis Cluster mode, refer to the Network OS Administrator's Guide.

Transitioning between modes

The `vcs set-rbridge-id` command allows you to change the RBridge ID or optionally the VCS ID of a node, and also to transition the node from logical chassis cluster mode to fabric cluster mode, or back to logical chassis cluster mode again.

ATTENTION: Each time you change the Brocade VCS Fabric configuration, the switch resets to the default configuration and reboots automatically. Make sure to save the configuration before you issue this command.

To change the RBridge ID of a node to 10:

```
# vcs set-rbridge-id 10
```

To change the RBridge ID of a node to 10 and its VCS ID to 35:

```
# vcs set-rbridge-id 10 vcsid 35
```

To convert a node that is in fabric cluster mode to logical chassis cluster mode, while simultaneously changing its RBridge ID to 10 and its VCS ID to 35:

```
# vcs set-rbridge-id 10 vcsid 35 logical-chassis enable
```

To convert a node that is in logical chassis cluster mode to fabric cluster mode, while optionally simultaneously changing its RBridge ID to 10 and its VCS ID to 35:

```
# no vcs set-rbridge-id 10 vcsid 35 logical-chassis enable
```

NOTE: You can also use the `no vcs logical-chassis-enable` command without the option of setting the RBridge ID and VCS ID.

Refer to the Network OS Command Reference for detail.

Prerequisites of Network OS v6.0.2

In this chapter, the prerequisites for Network OS v6.0.2 are described.

This chapter provides fundamental information about Network OS v6.0.2 as exemplified by the optional license, the standard compliance, scalability and compatibility.

Additionally the important notes described about the features of Hitachi Network OS v6.0.2 release.

- [Supported Optics and Option Cables](#)
- [Optional Licensed Software](#)
- [Software Upgrade and Downgrade](#)
- [Configuring and managing switches](#)
- [Scalability](#)
- [Compatibility](#)
- [Documentation Updates](#)
- [IMPORTANT NOTES](#)

Supported Optics and Option Cables

The embedded switch module for CB2500 supports following optics types:

I/F type (Speed)	Brocade P/N	Description
1GbE	XBR-000190 (1-pack)	RJ-45 connector
10GbE	10G-SFPP-SR (1-pack)	Optical transceiver (10G-SR)
	10G-SFPP-SR-8 (8-pack)	
	10G-SFPP-TWX-0301 (1-pack)	10Gb SFP+ twinax cable 3m
	10G-SFPP-TWX-0308 (8-pack)	
	10G-SFPP-TWX-0501 (1-pack) 10G-SFPP-TWX-0501 (8-pack)	10Gb SFP+ twinax cable 5m
40GbE	40G-QSFP-SR4 (1-pack)	Optical transceiver (40G-SR4)
	40G-QSFP-SR4-INT (1-pack)	Optical transceiver (40G-SR4) (4x10G SFPP break-out capable) Breakout Optical cable is not included with this optics.
	40G-QSFP-QSFP-C-0301 (1-pack)	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 3m
	40G-QSFP-QSFP-C-0501 (1-pack)	40GbE Direct Attached QSFP+ to QSFP+ Active Copper cable, 5m
	40G-QSFP-4SFP-C-0301 (1-pack) 40G-QSFP-4SFP-C-0501 (1-pack)	4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 3m 4x10GbE Direct Attached QSFP+ to 4 SFP+ Copper Breakout Cable, 5m
8G FC	XBR-000163 (1-pack)	Optical transceiver (8G FC SWL)
	XBR-000164 (8-pack)	
	XBR-000153 (1-pack)	Optical transceiver (8G FC LWL)
	XBR-000172 (8-pack)	
16G FC	XBR-000174 (1-pack)	Optical transceiver (8G FC ELWL)
	XBR-000192 (1-pack) XBR-000193 (8-pack)	Optical transceiver (16G FC SWL)
FC QSFP	XBR-000245 (1-pack)	4x8G or 4x16G FC QSFP breakout.

Optional Licensed Software

This release supports the following licensed features:

- **Port upgrade license 18p** - Allows customers to instantly scale the fabric by provisioning additional ports via license key upgrade. (Applies to select models of switches).
- **40Gb activate license** - Allows customers to instantly scale the fabric by provisioning additional 40G ports via license key upgrade.
- **FC/FCoE license** - Allows customers to enable FlexPort technology on this switch module via license key upgrade.

Software licenses are available in following formats.

Software License	SKU Description
GG-BE4LSL6X1-Y	Port upgrade license 18p ,DCB
GG-BE4LSL7X1-Y	40Gb activate license ,DCB
GG-BE3LSL2X1-Y	FC/FCoE license, DCB

Software Upgrade and Downgrade

Migration Path

Recommended upgrade/downgrade migration paths in both fabric cluster and logical chassis cluster modes are summarized in table below.

From \ To	NOS v4.0.1_hitx	NOS v6.0.1a1	NOS v6.0.2 (This release)
NOS v4.0.1_hitx	N/A	FWDL (Config loss)	FWDL (Config loss)
NOS v6.0.1a1	FWDL "default-config" (Config loss)	N/A	FWDL with "coldboot"
NOS v6.0.2 (This release)	FWDL with "default-config" (Config loss)	FWDL with "coldboot"	N/A

FWDL - means "Firmware download" command.

NOTES:

1. **Upgrade/Downgrade from/to v4.0.1 hitx to/from v6.x (v6.0.1a1 or v6.0.2) cannot preserve switch configuration**, it means the default configuration is loaded after upgrading/downgrading.
2. Upgrade/Downgrade from/to v6.0.1a1 to/from v6.0.2, **the "coldboot" option must be required** when you perform FWLD (firmware download) command.
3. Before downgrading to lower releases, it is recommended to disable all new features that are not supported on lower releases by using the "no" version of the CLIs. Stray configurations left out before downgrade can cause undesired behavior.
For example, **the previous release for the switch (NOS 4.0.1 hitx and NOS 6.0.1a1) does not support the Logical chassis cluster mode. Before downgrading to previous releases, you must configure the switch operation mode to the Fabric Cluster mode.**
4. Limitations:
 - (a) If port-security feature is enabled, it is necessary to limit the Max OUI configuration to 13 ports to avoid switch instability during firmware download operation.

- (b) If DNS is enabled on the switch, it is necessary to ensure the DNS servers are valid and reachable before executing firmware download command. If the DNS servers are not reachable, it is necessary to remove/correct the DNS configuration before executing firmware download command as the command is not blocked and may cause timeouts during firmware download operation.
- (c) When downgrade from NOS6.0.2 to NOS6.0.1 is performed, if IPv6 VRRP Link local Virtual IP is not in fe80::/64 format (but in fe80::/10 format), user is expected to remove this config before performing downgrade.

Firmware Installation

In Fabric Cluster Mode:

- The "firmware download" command is required to be executed by logging on to each individual node.
- Under certain stress conditions firmware download might time out on a node, (e.g. due to excessive processing load on the processor). The firmware download command will recover the system automatically. It is required to wait for completion of recovery before retrying the firmware download command.
- While upgrading firmware on the node, it is recommended not to make any configuration changes before firmware download has been completed successfully.

In Logical Chassis Cluster Mode:

- When a new switch boots up, the switch enters fabric cluster mode by default.
- So firmware installation to this release (NOS v6.0.2) must be executed when the switch is configured as the fabric cluster mode. After updating to this release, change the operation mode to logical chassis mode as necessary.

General information on installing Brocade Network OS can be found in the Brocade Network OS Administrator's Guide. This section includes special considerations and caveats to be aware of when upgrading to or from this version of Brocade Network OS, as well as recommended migration paths to use to reach this version of Brocade Network OS.

Note: Installing Brocade Network OS may be service disruptive and any unsaved running configuration may be lost during the process. In Fabric cluster mode, running-config needs to be saved to startup-config in order to preserve the running-config across reboots. In Logical Chassis mode, running-config is always preserved across reboots. The firmware version migration path determines if the configuration across upgrade/downgrade shall be preserved.

Upgrading to this Release

Upgrade all nodes in the cluster at same time -- Service Disruptive Cluster Wide

- Download the firmware on all the switches running Network OS v4.0.1_hitx. Download the firmware on all the switches running Network OS v6.0.1a1 using the "coldboot" option. (Refer to the "Migration Path" for detail.)
- After all switches complete the firmware download, they will be automatically rebooted.
- Since all nodes reboot at the same time, this procedure is service disruptive.

Configuring and managing switches

The following sections describe how to configure and manage Brocade switches.

Configuring Ethernet management interfaces

In Fabric Cluster Mode:

When a new switch boots up, the switch enters fabric cluster mode by default. In this case, the management interface is not configurable from the switch module console but only from the Hitachi Compute Blade management module.

This section shows an example to configure the static IP address for the management interface for this switch when the switch module is set to the fabric cluster mode.

1. Login to the ComputeBlade 2500 Management Module, using a web browser with default administrator and password login settings.

Web address: `https://<Chassis IP Address>`
Login: `<Chassis Management Login ID>`
Password: `<Chassis Management Password>`

2. Navigate to **Resources > Systems > Network > Management LAN**.

3. Select IP Address (v4) tab if you need to configure IPv4 address or IP Address(v6) tab to configure IPv6 address. And then click the "Edit" button.
4. Scroll to the bottom of the pop-upped "Edit IP address" window page until you see Switch Module 1 or Switch Module 2.
5. If the switch is installed in slot 1, use section Switch Module 1. If the switch is installed in slot 2, use section Switch Module 2. Enter the details in the fields to configure the Switch Module.

The following figure shows an example when you configure the IPv4. Enter details for IP Address, Subnet Mask, and Default Gateway as shown in the following figure.

The screenshot shows a dialog box titled "Edit IP Address" with a close button in the top right corner. Below the title bar, it says "Edit the management LAN network IP address settings." The dialog contains several sections for configuration:

- Default Gateway:** 0.0.0.0
- Server Blade 12 (Not Installed):**
 - IP Address: 172.16.219.22
 - Subnet Mask: 255.255.0.0
 - Default Gateway: 0.0.0.0
- Server Blade 13:**
 - IP Address: 172.16.219.23
 - Subnet Mask: 255.255.0.0
 - Default Gateway: 0.0.0.0
- Server Blade 14:**
 - IP Address: 172.16.219.24
 - Subnet Mask: 255.255.0.0
 - Default Gateway: 0.0.0.0
- Server Blade 15 (Not Installed):**
 - IP Address: 172.16.219.25
 - Subnet Mask: 255.255.0.0
 - Default Gateway: 0.0.0.0
- Switch Modules:** (indicated by a blue arrow icon)
 - Switch Module 1:**
 - IP Address: 172.16.219.31
 - Subnet Mask: 255.255.0.0
 - Default Gateway: 172.16.219.222
 - Switch Module 2:**
 - IP Address: 172.16.219.32
 - Subnet Mask: 255.255.0.0
 - Default Gateway: 172.16.219.222

At the bottom of the dialog, there are two buttons: "Confirm" and "Cancel". Both buttons are highlighted with a red rectangular box.

6. Click Confirm and then, scroll to the bottom of the page to verify the information entered. Click OK.

In Logical Chassis Cluster Mode:

When the switch module operation mode is changed to the Logical chassis cluster mode (LC mode), you must connect an Ethernet cable to the RJ45 port on the faceplate to a network for management.

The IP address setting that was configured from Hitachi Compute Blade management module when the switch worked as FC mode can be used on the switch works as LC mode. **But the IP address configuration from the management module is no longer used on the LC mode.** Instead, you must configure the IP address configuration from the CLI of the principal switch module of the VCS fabric that the switch module belongs.

For detail about the way to configure the management IP address from the switch CLI, refer to the Network OS Administrator's Guide, Configuring static IP addresses.

Configuring a switch in logical chassis cluster mode

This section describes how to configure a switch in logical chassis cluster mode. If you need other use cases other than the scenarios described in this chapter, refer to the Network OS Administrator's Guide for detail.

Creating a logical chassis cluster

This section covers the basic steps to create a logical chassis cluster, with the assumption that all physical connectivity requirements have been met. The following figure is a representation of a five node logical chassis cluster.

To create a logical chassis cluster, follow the steps in the following example:

1. Log in to one switch that will be a member of the logical chassis cluster you are creating:
2. In privileged EXEC mode, enter the vcs command with options to set the VCD ID, the RBridge ID and enable logical chassis mode for the switch. The VCS ID and RBridge IDs shown are chosen for the purposes of this example.

```
switch# vcs vcsid 22 rbridge-id 15 logical-chassis enable
```

3. The switch reboots after you run the vcs command. You are asked if you want to apply the default configuration; answer yes.
4. Repeat the previous steps for each node in the cluster, **changing only the RBridge ID each time.** You must, however, set the VCS ID to the same value on each node that belongs to the cluster.

- When you have enabled the logical chassis mode on each node in the cluster, run the show vcs command to determine which node has been assigned as the cluster principal node. The arrow (>) denotes the principal node. The asterisk (*) denotes the current logged-in node.

```
switch# show vcs
Config Mode : Distributed
VCS Mode : Logical Chassis
VCS ID : 22
VCS GUID : bcab366e-6431-42fe-9af1-c69eb67eaa28
Total Number of Nodes : 3
Rbridge-Id WWN                               Management IP  VCS Status Fabric Status  HostName
-----
15          10:00:00:27:F8:1E:3C:8C   10.18.245.143  Offline   Unknown   sw0
16          >10:00:00:05:33:E5:D1:93*  10.18.245.152  Online    Online    cz41-h06-m-r2
17          10:00:00:27:F8:F9:63:41   10.18.245.158  Offline   Unknown   sw0
```

The RBridge ID with the arrow pointing to the WWN is the cluster principal. In this example, RBridge ID "16" is the principal.

- Set the clock and time zone for the principal node. Time should be consistent across all the nodes. Refer to Network Time Protocol overview on Network OS Administrator's Guide.
- Log in to the principal cluster and make any desired global and local configuration changes. These changes then are distributed automatically to all nodes in the logical chassis cluster.

NOTE: You can enter the RBridge ID configuration mode for any RBridge in the cluster from the cluster principal node. You can change the principal node by using the logical-chassis principal priority and logical chassis principal switchover commands. For more information about cluster principal nodes, refer to Selecting a principal node for the cluster for detail.

Taking precautions for mode transitions

Ensure that all nodes to be transitioned are running the same version of Network OS. Logical chassis cluster mode is supported starting with Network OS release 6.0.2.

If you are merging multiple global configuration files to create one new global configuration file, be sure that the same entity name does not exist in the merged file. For example, if mac access-list extended **test1** contains the entries shown in the following "Node 1 global configuration" and "Node 2 global configuration", when you merge the files you can rename mac access-list extended **test1** from Node 2 to mac access-list extended **test2**, as shown in the "Combined global configuration."

Node 1 global configuration

```
mac access-list extended test1
seq 10 permit any 1111.2222.333a ffff.ffff.ffff
seq 20 deny any 1111.2222.333b ffff.ffff.ffff
seq 30 deny any 1111.2222.333c ffff.ffff.ffff
seq 40 permit any any
```

Node 2 global configuration

```
mac access-list extended test1
seq 10 permit any 4444.5555.666d ffff.ffff.ffff
seq 20 deny any 4444.5555.666e ffff.ffff.ffff
seq 30 permit any any
```

Combined global configuration

```
mac access-list extended test1
seq 10 permit any 1111.2222.333a ffff.ffff.ffff
seq 20 deny any 1111.2222.333b ffff.ffff.ffff
seq 30 deny any 1111.2222.333c ffff.ffff.ffff
seq 40 permit any any
!
mac access-list extended test2
seq 10 permit any 4444.5555.666d ffff.ffff.ffff
seq 20 deny any 4444.5555.666e ffff.ffff.ffff
seq 30 permit any any
```

The local configuration for Node 2 also needs to be changed accordingly. In this example, one of the local configuration changes would be the interface TenGigabitEthernet. Instead of referencing **test1**, the local configuration file for Node 2 needs to reference **test2** because of the change that was made to the global configuration file. This is shown in the following "Node 2 local configuration..." sections.

Node 2 local configuration before matching the combined global configuration

```
interface TenGigabitEthernet 4/0/3
fabric isl enable
fabric trunk enable
switchport
switchport mode access
switchport access vlan 1
spanning-tree shutdown
mac access-group test1 in
no shutdown
```

Node 2 local configuration after matching the combined global configuration

```
interface TenGigabitEthernet 4/0/3
fabric isl enable
fabric trunk enable
switchport
switchport mode access
switchport access vlan 1
spanning-tree shutdown
mac access-group test2 in
no shutdown
```

ATTENTION:

- Note that the copy default-config to startup-config command in logical chassis cluster mode causes a cluster-wide reboot and returns the entire logical chassis cluster to the default configuration. Therefore, use this command only if you want to purge all existing configuration in the logical chassis cluster.

- Make sure that the backup files for global and local configurations are available in a proper SCP or FTP location that can be easily retrieved in logical chassis cluster mode during restore. Do not save the files in the local flash, because they may not be available on the principal node for replay of local configurations.

Converting a fabric cluster to a logical chassis cluster

You can convert an existing fabric cluster to a logical chassis cluster using the default configuration file.

1. Be sure all nodes are running the same firmware version. Logical chassis cluster functionality is supported in Network OS 4.0.0 and later.
2. Be sure all the nodes that you intend to transition from a fabric cluster to a logical chassis cluster are online. Run either the `show vcs` or `show vcs detail` command to check the status of the nodes.
3. Log in to one switch that you are converting from fabric cluster mode to logical chassis cluster mode.
4. In Privileged EXEC mode, enter the `vcs logical-chassis enable` command with desired options; for example you can convert all R Bridges with one command:

```
switch# vcs logical-chassis enable rbridge-id all default-config
```

NOTE: To convert a specific R Bridge from fabric cluster mode to logical chassis mode, use the R Bridge ID value in place of the "all" option. You can also specify a range, such as "1,3,4-6". Refer to the Network OS Command Reference for details.

The nodes automatically reboot in logical chassis cluster mode. Allow for some down time during the mode transition.

5. Run either the `show vcs` or the `show vcs detail` command to check that all nodes are online and now in logical chassis cluster (listed as "Distributed" in the command output) mode.
6. The `show vcs` command output can also be used to determine which node has been assigned as the cluster principal node.

```
switch# show vcs
```

R-Bridge	WWN	Switch-MAC	Status
1	> 11:22:33:44:55:66:77:81	AA:BB:CC::DD:EE:F1	Online
2	11:22:33:44:55:66:77:82	AA:BB:CC::DD:EE:F2	Online
3	11:22:33:44:55:66:77:83*	AA:BB:CC::DD:EE:F3	Online

The R Bridge ID with the arrow pointing to the WWN is the cluster principal. In this example, Rbridge-ID 1 is the principal.

7. Log in to the principal cluster and make any desired global and local configuration changes. These changes then are distributed automatically to all nodes in the logical chassis cluster.

NOTE: You can enter the RBridge ID configuration mode for any RBridge in the cluster from the cluster principal node.

NOTE: You can change the principal node by using the logical-chassis principal priority and logical chassis principal switchover commands. For more information about cluster principal nodes, refer to Selecting a principal node for the cluster.

Selecting a principal node for the cluster

Logical chassis cluster principal node behavior includes:

- All configuration for the logical chassis cluster must be performed on the principal node.
- By default, the node with the lowest WWN number becomes the principal node.
- You can run the show vcs command to determine which node is the principal node. An arrow in the display from this command points to the WWN of the principal node.
- You can select any node in the logical chassis cluster to become the principal by running the logical chassis principal priority command, followed by the logical-chassis principal switchover command, as shown in the following example (in this example, RBridge ID 5 is being assigned with the highest priority):

```
switch# configure
switch(config)# rbridge-id 5
switch(config-rbridge-id-5)# logical-chassis principal-priority 1
switch(config-rbridge-id-5)# end
switch# logical-chassis principal-switchover
```

A lower number means a higher priority. Values range from 1 to 128. Until you run the "logical-chassis principal switchover command", the election of the new principal node does not take effect.

Adding a node to a logical chassis cluster

Nodes can be dynamically added to an existing logical chassis cluster. If the proper physical connections exist between the existing logical chassis cluster and the new node, the process is automatic.

Log into the new node and run the "vcs logical-chassis enable" command with the desired options. You must assign the new node the VCS ID of the existing cluster.

You can run the show vcs command to verify that the status of the added node is "online."

Removing a node from a logical chassis cluster

If the "no vcs logical-chassis enable rbridge-id <rbridge-id | all> default-config command" is executed on a switch that is currently in logical chassis cluster mode, the switch boots in fabric cluster mode. The following is example:

```
# no vcs logical-chassis enable rbridge-id 239 default-config
```

Once the node is converted to fabric cluster mode, the Rbridge goes into offline state from the original cluster. To remove the configuration of the node, you must enter the "no vcs enable rbridge-id <rbridge-id>"command, as shown in the following example:

```
# no vcs enable rbridge-id 239
```

Once the node is removed, all configurations corresponding to that node are removed from the cluster configuration database. Similarly, the removed node does not retain any configurations corresponding to the other nodes in the cluster.

Rejoining a node to the cluster

Nodes that are temporarily isolated from a logical chassis cluster can re-join the cluster as long as no configuration or cluster membership changes have taken place on either the deleted node or the cluster. Run the "vcs logical-chassis enable" command with the desired options to rejoin the node to the cluster.

However, if configuration changes have occurred on either the node or cluster since the node was removed, you must reboot the node with its default configuration by issuing "copy default-config startup-config" on the segmented node.

Replacing a node in a logical chassis cluster

If a node in a logical chassis cluster becomes damaged and no longer be used, a similar node with identical capabilities can be used in its place.

The new node must use the same RBridge ID of the node that is being replaced. When the new node is detected, it joins the cluster as a previously known node instead of being considered a new node.

To replace a node that has an RBridge ID of 3 and then enter the WWN of the new node, follow the steps shown in the following example:

1. Add the new switch hardware to the network and connect all data cables.
2. Power on the replacement hardware and add the switch to the network as a standalone switch.
3. Run the following command on the principal switch:

```
switch# vcs replace rbridge-id 3  
Enter the WWN of the new replacement switch: 11:22:33:44:55:66:77:81
```

4. Assign the RBridge ID of 3 to the new node by running the following command on the new node:

```
switch# vcs rbridge-id 3
```

Converting a logical chassis cluster to a fabric cluster

To transition all nodes in a logical chassis cluster to a fabric cluster, using default configurations, perform these steps:

1. Make sure all the nodes that you intend to transition from a logical chassis cluster to a fabric cluster are online. Run either the show vcs or show vcs detail command to check the status of the nodes.
2. Log in to the principal node on the logical chassis cluster.
3. Run the following command to convert all RBridge IDs: "no vcs logical-chassis enable rbridge-id all default-config".

NOTE: To convert just one RBridge ID, specify the ID as shown in the following example: no vcs logicalchassis enable rbridge-id rbridge-id default-config.

The nodes automatically reboot in fabric cluster mode. Plan for some down time for this transition.

4. Run either the show vcs or show vcs detail command to check that all nodes are online and now in fabric cluster (listed as "Local-only" in the command output) mode.

Scalability

All scalability limits are subject to change. Limits may be increased after further testing has been completed, even after the release of a particular NOS version.

Network OS v6.0.1a1 Scalability Numbers	GG-BE4LSW2X1
Maximum # of dot1Q VLANs (Virtual-Fabric Disabled)	4,096
Maximum # of VLANs (dot1Q +Virtual-Fabric)	6,000
Maximum # of Service Virtual Fabric VLANs	2,000
Maximum # of Transport Virtual Fabric VLANs	1,000
Maximum # of MAC addresses per Switch	120,000
Maximum # of MAC addresses per Fabric (with CML)	256,000
Maximum # of MAC addresses across VxLAN tunnels per VCS cluster for VMware NSX	8,000
Maximum # of MAC-based Virtual-Fabric VLAN Classification per switch	256
Maximum # of Classified Virtual Fabric VLANs per Trunk Interface	2,000
Maximum # of port profiles(AMPP)	1,000
Maximum # of VLANs in port profiles	3,500

Network OS v6.0.1a1 Scalability Numbers	GG-BE4LSW2X1
Maximum # of dot1q VLANs that can be attached on VxLAN GW for VMware NSX	2,000
Maximum # of VLANs (dot1q VLANs attached to VxLAN GW for NSX + Virtual Fabric VLANs enabled on edge-interfaces)	(2,000+1,000)
Maximum # of VxLAN tunnels with VMware NSX	250
Maximum # of service-nodes with VMware NSX	5
Maximum # of MAC Associations for AMPP	8,000
Maximum # of per priority pause levels	3
Maximum # of VMware vCenters per Fabric	4
Maximum # of ELD instances in the fabric	2,000
Maximum # of IGMP Snooping Interfaces supported	512
Learning rate for IGMP snooping (groups/second)	512
Maximum # of L2 (IGMP Snooping) multicast groups	6,000
Maximum # of MLD Interfaces	256
Maximum # of MLD Groups	4,000
Learning rate for MLD snooping (groups/second)	512
# of L3 (S,G) forwarding Entries	2,000
# of L3 (*,G) forwarding Entries	256
PIM Interfaces Supported	32
IGMP interfaces supported	32
Learning Rate for PIM-SM (flows/second)	32
Maximum # of L2 ACL(ingress/egress) *	3,000/120
Maximum # of L3 ACL ipv4 (ingress/egress) *	1,500/120
Maximum # of class-maps	2,048
Maximum # of policy-maps	2,048
Maximum # of class-maps per policy map	50
Maximum Total # of L3 ACL ipv6 (ingress/egress) *	500/120
Maximum # of VF/FCoE interfaces/Logins (Per switch)	1,000
Maximum # of Enodes/FCoE Devices per Fabric	2,000
Maximum # of NPIV per Port	64
Maximum # of SAN Devices (FC + FCoE) per Fabric	3,000
Maximum # of MSTP instance	32
Maximum # of VLAN in PVST	128
Maximum # of LAGs (Port Channels)	60
Maximum # of members in a standard LAG	16
Maximum # of members in a Brocade Trunk (10G)	16
Maximum # of members in a Brocade Trunk (40G)	2
Maximum # of switches in a Fabric cluster mode **	48
Maximum # of switches in a Logical chassis mode **	48
Maximum # of ECMP Paths	16
Maximum # of VLAGs in a fabric	2,000
Maximum # of member ports in a VLAG	64
Maximum # of nodes in a VLAG	8
Maximum # of member ports per VLAG per Node	16
Maximum # of Management ACL	256

Network OS v6.0.1a1 Scalability Numbers	GG-BE4LSW2X1
Maximum # of ARP Entries *	16,000
Maximum # of OSPF areas	20
Maximum # of OSPF routers in a single area	64
Maximum # of OSPF adjacencies	100
Maximum # of OSPF routes *	8,000
# of OSPF Interfaces	100
# of OSPF enabled subnets	100
# of local subnets in a single area	100
Maximum # of OSPFv3 areas	9
Maximum # of OSPFv3 routers in a single area	64
Maximum # of OSPFv3 adjacencies	100
Maximum # of OSPFv3 routes *	1,500
# of OSPFv3 Interfaces	100
# of OSPFv3 enabled subnets	100
Maximum # of IPv4 routes in SW *	8,000
Maximum # of IPv6 routes in SW *	1,500
Maximum # of IPv4 static routes *	2,000
Maximum # of IPv6 static routes *	500
Maximum # of VRRP instances per system	255
Maximum # of VRRP v3 instances per system	255
Maximum # of VRRP instances per interface	32
Maximum # of routers participating in a VRRP-E session	8
Maximum # of virtual IP addresses per VRRP instance	16
Maximum # of FVG instances per system	256
Maximum # of FVG instances per interface	1
Maximum # of routers participating in a FVG session	32
Maximum # of Gateway IP addresses per FVG instance	1
Maximum # of IPv4 routes with ECMP supported *	8,000
Maximum # of IPv6 routes with ECMP supported *	1,500
Maximum # of L3 ECMP	16
Maximum # of IPv4 interfaces per system *	2,000
Maximum # of IPv6 interfaces per system *	512
Maximum # of VRF per node	32
Maximum # of VRFs support protocols per node	32
Maximum # of I-BGP peers	256
Maximum # of E-BGP peers	64
Maximum # of IPv4 BGP routes in HW *	8,000
Maximum # of IPv6 BGP routes in HW *	1,500
Maximum # of IPv4 RIB (IN + OUT) Routes *	110,000
Maximum # of IPv6 RIB (IN + OUT) Routes *	110,000
Maximum # BGP IPv4/IPv6 Peer Group	100
Maximum # of BFD sessions per node	100
Maximum # of UDLD enabled interfaces	64
Maximum # of PVLAN domain supported	1,000
Maximum # of Secondary vlans per PVLAN supported	24
Maximum # of primary vlans per PVLAN supported in promiscuous mode	24

Network OS v6.0.1a1 Scalability Numbers	GG-BE4LSW2X1
DHCP IP Helper Addresses per interface	16
DHCP IP Helper Ve interfaces	256
DHCP IP Helper physical ports	60
DHCP IP relay Address on the system	2,000
DHCP IPv6 Relay Address	2,000
Max Number of configurable PBR route maps	64
Max Number of configurable PBR stanzas	1,024
Max Number of TCAMs available for PBR	512
Max Number of configurable next hops within a single PBR stanza	128
Maximum # of OpenFlow L2 flows	1,000
Maximum # of OpenFlow L3 flows	1,000

* Parameters mentioned are applicable on specific HW profiles. Please check the Network OS Administrator's Guide for the specific HW profiles.

** Please consult your Brocade SE for best practices when designing a 48-node VCS Fabric. In Hybrid cluster environment (a cluster involving various VDX platforms), the scalability limit of the cluster is determined by the scalability limit of the lowest denominator. For instance, in such a fabric, if the MAC scalability limit on one VDX platform is lower than the other, then the fabric supports the lower scale value.

Compatibility

In VCS Fabric mode, 6.0.2 has the connectivity to only Brocade NOS 6.0.2 into same fabric. If 6.0.2 connects to other than Brocade NOS 6.0.2, the VCS Fabric will be segmented. Regarding to the connectivity to standard Ethernet switch in VCS Fabric mode, please refer to the Network OS Administrator's Guide.

Documentation Updates

When using the NOS 6.0.2 documentation, the embedded DCB switch module is equivalent to the VDX 6740x except where noted in the release note document. The most recent NOS documentation manuals are available on MyBrocade: <http://my.brocade.com/>

Following NOS 6.0.2 documents are recommended references.

1. Network OS Administrator's Guide : 53-1003768-02
2. Network OS Command Reference : 53-1004157-01
3. Network OS Layer 2 Switching Configuration Guide : 53-1003770-03

4. Network OS Layer 3 Routing Configuration Guide	: 53-1003771-01
5. Network OS Software Licensing Guide	: 53-1003773-01
6. Network OS Message Reference	: 53-1003777-02
7. Network OS MIB Reference	: 53-1003781-01
8. Network OS Defined Networking (SDN) Configuration Guide	: 53-1003782-02
9. Network OS Security Configuration Guide	: 53-1003769-03
10. Network OS Troubleshooting Guide	: 53-1003772-01
11. Network OS REST API Guide	: 53-1003775-02
12. Network OS NETCONF Operations	: 53-1003778-02
13. Network OS YANG Reference Manual	: 53-1003779-02
14. Network OS Puppet User's Guide	: 53-1003847-01
15. Network OS MAPS Administrator's Guide	: 53-1003783-02

IMPORTANT NOTES

This section contains information that you should consider before you use this NOS release.

Command Line Interface

- Break command is not supported. Ctrl-c can be used as an alternative.
- Few commands may not display paginated output.
- For few clear and show commands "?" will not show all options for VRF. Tab completion will give all possible values.
- For certain commands (including no form with some commands), "?" will show unsupported additional options.
- Some CLI commands will generate an "Error: Access denied" message upon failure. This means the operation failed on the switch and may not be related to permissions.
- Tab completion and <ctrl>-c (cancel) does not work for some commands.

- Incorrect range might be displayed in the help text for some of the show commands.
- Range support is available for all the interfaces in Network OS v6.0.0. Following limitations are applicable:
 - Interface range command is supported on breakout ports of same connector. Range is not supported involving breakout ports of multiple connectors.
 - Interface range command does not support mix of regular ports and breakout ports.
 - Range command is not supported across multiple slots of the chassis
 - In some instances, there could be delay in starting of range operation after issued
 - When range issued for very large subset (e.g. 4k VLAN), timeout can happen or user may see temporary switch being unresponsive or high CPU. Brocade recommends using range in smaller chunks. Especially, while configuring VLANs/VEs, Brocade recommends range to be less than 500.
 - Range prompt doesn't get updated when few or all of interface in that range are deleted. Therefore, user should exit from Range sub-mode if few or all interfaces are deleted that are part of that range. New configuration performed on same range sub-mode may give unpredictable results.
- System does not warn user on deleting the ip config when vrf is configured.
- Redistributed connected/static routes may be shown twice as part of config.
- If "switchport trunk allowed vlan all" is already configured on any interface, then Vlan creation using range command will be slow as each vlan will get provisioned individually.
- Some unsupported debug commands may be seen. Brocade recommends not to run them on switches:
 - Show confd-state -, for debugging purpose only.
 - Show parser dump -, for debugging purpose only
 - Show notification stream -, for debugging purpose only
 - Show features - no use
 - Show ssm -, for debugging purpose only.
 - Autoupgrade command in config mode

- During "copy running-config startup-config" or "copy support" user might see occasional CPU spikes (to ~30-40%).
- show mac-address-table command on console with include option cannot be aborted with a break/ctl-C. Use a telnet session for the same.
- Short form of MAC-Address is not supported as filter in "show running-config".
- For ip access lists, display filtering based on sequence number alone does not work as expected.
- Certain oscmd commands may not work or give a different output under admin login
- If an alias exactly matches a partial keyword anywhere in the command line, pressing the TAB key for CLI command completion will claim that the input is invalid, and pressing the ENTER key will first replace the partial keyword with the alias expansion string. To avoid this, make sure that any partial keywords are not an exact match for an alias name.
- The authentication mode with primary & secondary sources of authentication cannot be updated to a configuration containing only the primary source. For example, the authentication mode cannot be changed from "radius local or radius local-auth-fallback" to 'radius'. The workaround is to remove the existing configuration and then configure it to the required configuration.
- NTP server with full length IPv6 address configuration can be used only with NTP key with less than 15 characters length.
- The "logging syslog server" command returns an error on the "secure" keyword. Use "secure port" to assign a nondefault port number.
- OSPFv3 on default VRF can be created without mentioning VRF name but while removing default VRF needs to be provided like "no ipv6 router ospf vrf default-vrf".

Platform

- 1G copper SFPs do not support exchanging flow-control settings during the auto-negotiation process. It is recommended to configure static mode of configuration of flow-control on both the ends of the desired link.
- System verification/ offline diagnostics tests need "chassis disable" before the test and "chassis enable" followed by immediate reboot.
- The GG-BE4LSW2X1 switch does not support IP fragmentation. MTU errors are reported in "show interface" as "Errors" under the "Transmit Statistics".
- Logical SAN is not supported in fabric cluster mode.

Licensing

- On the switch that has Flexport FC capable interfaces, enabling FibreChannel ports requires only the FCoE license to be installed and does not require any Port Upgrade license. The Port Upgrade license only controls Ethernet ports (number of ports or speed supported).
- An Integrated Routing license is NOT required on FOS-based SAN platforms running FOS 7.0.1 or above for FCR interoperability connectivity with VCS fabrics and the VDX 674x platforms including the GG-BE4LSW2X1 switch. Please refer to the FOS v7.0.1 Admin Guide documentation on configuring FOS platforms for connectivity to VDX 674x switches and VCS fabrics.

VCS

- Loopback connection is not supported in VCS mode. If a loopback connection is done (either using loopback plugs or port to port connections on the same switch), those interfaces become ISL interfaces.
- A node with default configuration will not join a cluster if the intermediate nodes between the node being defaulted and rest of the cluster are also undergoing reload. If the node boots up earlier than the intermediate nodes, it will form its own VCS and not join the parent cluster. In such situations, reload the node that is required to join the cluster.
- Fabric Cluster Mode:
 - When a new switch is added to an existing VCS fabric and if the new switch takes the role of principal node, the other switches in the fabric will receive the configuration of the distributed features such as Virtual IP and VM-Aware Network Automation from the newly added switch. This will cause the existing distributed configuration to be overwritten by the newly added switch in the principal role. This can be avoided by following the new switch addition procedures in the Admin Guide.
 - After a cluster reboot, Brocade recommends to do both "show fabric all" and "show vcs" to ensure that cluster is entirely formed without any issue. User might see that 'show vcs' takes an additional 2-3 minutes to show all participating switches. This is an existing behavior and doesn't affect data path functionality in most cases.
- "show fabric isl" & "show fabric trunk" may show the interfaces in random order without sorting
- The default-configuration behavior may be different depending on the default-configuration triggers.

Note the following results for the given actions.

Default config Trigger	Global Config (i.e. virtual-fabric)	Local Config (i.e. SFP breakout)
------------------------	--	-------------------------------------

copy default-config startup-config	Preserved	Preserved
VCS-ID and/or Rbridge-ID change	Preserved	Removed
firmware download default config	Removed	Removed
Write-erase	Removed	Removed

Logical Chassis

- Configurations are not auto preserved on mode transitions (between Fabric Cluster and Logical Chassis mode). Please follow the mode transition procedure as outlined in the Network OS Admin Guide. Non-default User Id/password will be lost when migrating from FC to LC.
- "show vcs" output displaying "Coordinator" indicates "Principal" node role.
- Principal priority value ranges from 1 to 128, 1 being the highest. Recommend to set higher principle priority to VDX 8770.
- User should not make configuration change during Logical Chassis firmware upgrade or while ISL toggling to prevent the switch segmenting from the cluster due to configuration mis-match.
- Upon Node segmentation from the cluster, user should run "copy default start" or exercise the default-config boot feature on the segmented switch to bring it back to the cluster.
- Number of config snapshots saved on switch is limited to 4 per rbridge ID. So on 24 node cluster, a max of $24 * 4 = 96$ snapshots are possible.
- For Netconf and SNMP, user has to poll using individual node Management IP.
- Creating a snapshot with "\" in snapshot-id creates the snapshot file with incorrect name.
- There will not be any raslog to the user when replacement of a node fails.
- With large configs, while a switch is rejoining a fabric with default config, "%Error:Could not find Interface" may be printed temporarily. The switch will recover and join the fabric.
- Config changes during principle switch-overs are not supported and may segment the cluster.
- An Rbridge in bare-metal state can join a VCS with or without the pre-provisioning mapping. For the scenario without the pre-provisioning mapping, the Rbridge must be in bare-metal state and in the "Offline" state of the VCS.
- Disabling virtual-fabric may take up to 10 minutes depending on the number of ISLs and VLAN interfaces configured in the VCS.

Brocade Trunks

- The GG-BE4LSW2X1 Brocade trunk (BTRUNK) can support up to 16 member links with a maximum throughput of 160G using 16x10G ports in the same trunk group. On these platforms traffic may not be distributed evenly across all member of a trunk at lower traffic rates.
- The GG-BE4LSW2X1 Brocade trunk (BTRUNK) can support up to 2x40G member links in the same trunk group for a maximum throughput of 80G.

Breakout Interfaces

- The LED state for a breakout interface is deterministic. For all other supported platforms, the LED state for a breakout interface is not deterministic.
- In breakout mode, there is only SFP and no per-breakout media information. The show media command will displays the same media information for all breakout interfaces. The TX Power Field in the show media command is not supported by the 40G optics.
- On 40G native mode - Breakout configuration is not blocked. If configured on one side, other side of link won't able be identify peer port config is breakout and link won't be stable.
- On the GG-BE4LSW2X1 switch, the breakout ports are FlexPort capable, and may be configured to connect to FC switches with 4x16G breakout supported cables and optics.

Restrictions for Ports in 1G Mode

- RMON stats are calculated incorrectly for packet sizes 64-127 bytes.
- 1G ports cannot form ISL links. Only 10G ports can be used to form ISL links.
- Brocade Trunks cannot be formed with 1G. Brocade Trunks are only supported on 10G.
- A LAG cannot be created between 1G and 10G ports.
- FCoE configuration is NOT supported on 1G ports.
- DCBX configuration for FCoE is not supported on 1G ports.

vLAG

- LAGs are created with default speed of 10G. Therefore Brocade recommends end user to set required speed manually based on member speed using "speed" command.

- When configuring LACP LAG between VDX & non-Brocade switches it is highly recommended to enable the VLAG ignore-split on the VDX. Ignore split option is enabled by default.
- The port-channel interface "load-balance" is not the same as "fabric port-channel <#> loadbalance"
 - The port-channel interface "load-balance" command configures load-balancing on the actual vLAG member links (effective on Rbridges directly participating in the vLAG).
 - The "fabric port-channel <#> load-balance" configures load-balancing on Rbridges NOT participating in the vLAG, but connecting to neighboring vLAG participating Rbridges.

Virtual IP Address Support

- A separate gateway cannot be configured for Virtual IP address. Default gateway will be the same as the gateway address for the management port of this switch.
- There is no Virtual MAC address associated with the Virtual IP address, physical MAC will be used.
- For VCS Virtual IP address to work correctly the management port's IPv4 address should be assigned, functional and both address should be in same subnet.

Security, Management ACLs, Authentication, Authorization

- Login authentication service (aaa authentication login cli):
 - With "local" option specified as secondary authentication service, local authentication will be tried only when the primary authentication service (TACACS+/Radius/LDAP) is either unreachable or not available.
 - Behavior of "local" option in pre-4.1.0 releases is changed to the "local-auth-fallback" option.
 - When login authentication configuration is modified, the user sessions are not logged out as in previous releases. All connected user sessions can be explicitly logged out using "clear sessions" CLI.
- ACLs are not supported for egress traffic flows on management interfaces.
- Configuring TACACS+ or RADIUS without a key is not supported. If no key is configured, the switch uses a default key of "sharedsecret".

- There is a possibility that locked user accounts will get unlocked after a reboot if the running-config (before reboot) is different from startup-config of user accounts.
- Encrypted text (taken from running-config of any user account password with encryption turned on) should not be used as input for clear-text password for the same user. This may result in login failure of the user subsequently.
- When the ACL is applied to a management interface, only the top 256 rules will be applied if the ACL contains more than 256 rules.
- When more than 250 rules ACL's are configured (over supported scale), they may be partially installed & effective.
- Access to ONLY the following Active Directory (AD) servers is supported by Brocade LDAP client:
 - Windows 2000
 - Windows 2003
 - Windows 2008 AD

SPAN & RSPAN

- CPU-originated packets cannot be output spanned.
- If SPAN has to be supported to multiple locations, please use RSPAN on vlan.
- Spanning of LAG port is not supported. To span a LAG, user should individually enable spanning on all the member ports of the LAG. However flow based SPAN is supported on LAG port.
- A profiled port cannot be a SPAN destination.
- SPAN destination port statistics will keep incrementing even when port is operational or admin down.

MAC Learning Considerations in VCS

- Under rare circumstances, end user might see mac address sync up issues on few nodes of a cluster (where 1 or more MAC addresses might be missing in some nodes). Brocade recommends to do "clear mac-address-table dynamic" in such cases.
- Static mac addresses will be displayed even when interfaces are down. This may cause black-holing of the traffic.
- Under certain conditions, MAC addresses may not be learnt even though ARP's may be learnt for those same MAC addresses.

PVLAN

- Following PVLAN features are not supported:
 - IGMP on PVLANS but there is no error message displayed if operator configures IGMP snooping on PVLAN
 - ARP & Routing in PVLAN domain
 - Enabling Routing in Primary and Secondary Vlans.
 - CLI to enable Local Proxy ARP on primary VLAN.
 - IP Configuration on PVLANS
 - Vx Configuration on both Primary and Secondary Vlans
 - AMPP on PVLANS
 - In case of MSTP if a primary VLAN is added to the instance automatically secondary VLAN also added to the instance.
 - When the operator wants to delete the host association on a host port recommended to use "no switchport" rather than "no switchport private-VLAN host-association". This is applicable only when the host port is untagged. When the host port is tagged both the commands can be used.
 - Primary VLAN ID needs to be lower than the secondary VLAN IDs. If primary VLAN ID is greater than secondary there is an issue with config replay.

UDLD

- The UDLD protocol is not supported on the members of a Brocade trunk.
- The UDLD protocol is not compatible with Cisco's proprietary UDLD protocol.
- UDLD needs to use the higher timer in Scale and Stress environment. UDLD may flap during HA failover.

STP/DiST

- VDX does not support tunneling non-standard BPDUs and thus IEEE BPDUs (0180:C200:0000) generated as tagged packets in STP/RSTP/MSTP modes may not be tunneled successfully across VCS fabric. However, VDX supports tunneling standards' based BPDUs such as untagged IEEE BPDUs and tagged or untagged PVST BPDUs (0100:0CCC:CCCD). Post 3.0.1, the tagged IEEE BPDU can be tunneled across VCS fabric using command: "tunnel tagged-ieee-bpdu" under interface configuration.

- In Fabric Cluster mode, global spanning-tree configurations (STP enable, STP Vlan configurations, STP over vLAG configurations) have to be performed in all the switches in VCS at the same time. For example, to run spanning-tree, it has to be enabled on all the switches including switches that don't have any edge ports.
- By default global spanning-tree and interface level spanning-tree will be disabled, user has to explicitly enable on the desired ports. vlan spanning-tree state is default enabled
- BPDU tunnel configurations are permitted only when spanning-tree is disabled in VCS.
- For cisco proprietary Per Vlan Spanning Tree protocols (PVST and RPVST) user needs to configure Brocade switch to send BPDU on Cisco multicast destination mac address "0100.0ccc.cccd" for non-native vlans. By default, NOS 4.1.0 software use's brocade "0304.0800.0700" multicast mac to send BPDU's on non-native vlans. Since NI/FI/Cisco boxes use Cisco multicast mac address to send spanning tree BPDU on non-native vlans, this configuration is needed in VDX switches to interoperate. This is an interface specific configuration.

Below is the example to configure Cisco BPDU mac for PVST and RPVST under interface mode,

```

VDX 6740-VCS1# conf t
VDX 6740-VCS1(config)# protocol spanning-tree rpvst
VDX 6740-VCS1(config-rpvst)# exit
VDX 6740-VCS1(config)# interface Port-channel 100
VDX 6740-VCS1(config-Port-channel-100)# spanning-tree bpdu-mac ?
Possible completions:
  0100.0ccc.cccd Cisco Control Mac
  0304.0800.0700 Brocade Control Mac
VDX 6740-VCS1(config-Port-channel-100)# spanning-tree bpdu-mac 0100.0ccc.cccd
VDX 6740-VCS1(config-Port-channel-100)# exit
VDX 6740-VCS1(config)#

```

Edge Loop Detection (ELD)

- ELD is supported on the edge ports that are connected either by end-hosts OR another switch OR another VCS.
- ELD is also supported for edge interfaces connected to hosts too.
- The edge-loop-detection port-priority with the higher number takes priority for shutting down the loop interface. If the port-priority is the same, the highest interface ID followed by the highest Rbridge-ID are used as the deciding metric.

Long Distance ISL Ports

- Long distance ISL configuration ("long-distance isl" command) is not allowed if CEE Map/fcoeport is configured on any edge ports in the same port group.

- CEE Map modification is not allowed when long distance ISL is configured.
- A maximum of three PFCs can be supported on a long distance ISL configured platform.
- When long distance ISL is configured on the switches, all ports in the port group will be bounced.
- Both side of long distance link should have long-distance-isl configuration. Otherwise end to end PFC might not work appropriately.
- For 10Km/Extended Range long distance configuration all other ISLs in the port group will be disabled.
- For 2Km/5 Km long distance configuration, one other ISL will be allowed to come online in the port group.
- For 2 km, 5 km and 10 km long-distance, use Brocade supported Long Range (LR) optics for direct connectivity.
- For 30 km long-distance, use Brocade-supported Extended Range (ER) optics for direct connectivity.
- "long-distance isl" command based extended fabrics are supported only on 10G interfaces. 40G and 100G interfaces do not support "long-distance isl" command, however can extend distances for non-lossless traffic up to 2Km using standard ISLs. On standard ISLs, the 10G, 40G and 100G interfaces support lossless traffic up to 1Km.
- The "long-distance-isl" command will not be supported on the SO-10GE-ZR-CX and 10G-SFPP-ZR 80km optics.

AMPP and Port-Profiles

- Port-profile status does not reflect the remote interface info in VCS fabric mode.
- Native VLAN support inside AMPP does not honor the global enable/disable flag.
- SPAN destination port cannot be a profiled port.
- Brocade recommends deleting all manually created port-profiles when migrating from a legacy AMPP environment to VM Aware Network Automation.
- Vmkernel related port-profiles may unapply/reapply during HA resulting in vmotion failures.

- Default port-profile configuration is changed from Network OS v4.0.1_hit1. The "switch port trunk allow VLAN all" that was present in prior version is removed. Other configuration stays the same.
- From Network OS v4.0.1_hit1, user defined port-profile-domain is introduced to control the VM mobility. Port-profile created must be explicitly associated with a profile domain.
- From Network OS v4.0.1_hit1, after upgrade, a new port-profile named UpgradedVlanProfile is auto-created. This profile has the single VLAN profile that contains the "switch port trunk allow VLAN all". This is the configuration that is present in the default port-profile prior to Network OS v4.1.0.
- Mac-based classification allowed only on access port-profile and C-tag classification allowed only on trunk port-profile.
- When a port becomes a profiled-port, all SERVICE VFs in that domain are provisioned on this port.
- "Switch trunk allow VLAN all" can only be present in one domain, it cannot co-exist with other c-tag based classifications in that domain.
- User is not allowed to edit/delete the default-profile-domain when Service VF is disabled.
- New port-profile is not auto added to the default domain when Service VF is enabled. It can only be explicitly added to or removed from the default profile-domain.
- On disabling Service VF UpgradedVlanProfile should be re-configured with "switchport trunk allowed VLAN all" in Default-profile-domain if it is removed /modified.
- Newly created port-profiles which is not part of any domain should be added to the defaultprofile-domain explicitly while disabling the Service VF.
- SERVICE VF classification cannot conflict across port-profiles in the same port-profile domain, but it can conflict across PP in different domains. i.e. a port-profile-domain cannot contain conflicting SERVICE VF classifications.

vCenter

- Receiving more than five vCenter events within a span of 30 seconds, results in asset discovery getting initiated. Post discovery cluster configuration will be in sync with vCenter.
- vCenter auto-profile is automatically added/deleted to the default port-profile-domain in Service VF enabled/disabled mode.

- Modifying/editing the auto port-profiles in the default-domain is not recommended, which may cause auto-pp application failure during vCenter operation and end up in traffic failure.
- Adding/removing the auto-port-profile to the user-created domain when Service VF is enabled is not recommended which may cause auto-pp application failure during vCenter operation and end up in traffic failure.
- vCenter auto-profile does not support SERVICE VF classification.

QoS

- It is recommended to use the same CoS Tail drop threshold on all members of a port-channel to avoid unpredictable behavior.
- In a hybrid logical-chassis, if a user configures a platform specific feature, it will be configured only on the rbridges which support that feature.
- Asymmetric pause is supported on 1G port interfaces
- It is recommended to enable maximum 2 PFC s on edge interfaces. Flow control is disabled by default on all interfaces.
- Priority 7 is reserved for control traffic on VDX switches. User data traffic should use priorities 0 through 6. Priority 3 is used for the FCoE lossless traffic by default.
- Brocade VDX architecture prioritizes Unicast traffic over Broadcast or Multicast traffic under port congestion.
- The interface queues operate in Strict Priority mode when there are no ISLs online on the switch. This could result in potential starvation of the CPU queue if line-rate traffic flows through an interface.
- Byte count is not supported for match ACL statistics.
- Byte count is not supported for RED statistics.
- The "count log" option in ACL is not supported for Flow based QoS and SysFBQ.
- The CLI "qos trust cos" is not applicable in VCS mode. However, "show qos int" will show as cos is trusted on ports on which "cos-mutation" or "cee default" config is applied.
- Configuring an interface with a nondefault DSCP-to-traffic class-map is allowed. However, configuring an interface with a nondefault CoS-to-traffic class-map is not supported.

FCoE

- VLAN's which are reserved for FCoE may not be used for any other purpose. Brocade recommends not having FCoE ports and Long Distance ISL ports in the same portgroups. This configuration will NOT be prevented by the CLI; however it can result in unpredictable behavior for FCoE traffic.
- Brocade recommends that for all LAGs with FSB, the fcoeport config must be applied on the LAG itself and for all LAGs with directly attached CNAs, the fcoeport config must be applied on the member ports.
- If FCoE priority is changed from default to non-default, user might see that FCoE login may not happen. Please toggle the interface using "shutdown" followed by "no shutdown" to work this around.
- Binding an enode mac to FCoE interface is not allowed in range context, as only one enode mac can be bound to one FCoE interface.
- While providing range for FCoE interfaces, it's recommended to provide the range only in ascending order. For ex: interface fcoe 1/48/11-38 is recommended, interface fcoe 1/48/38-11 is not recommended.
- FCoE traffic may not be mirrored using RSPAN. Workaround is to use SPAN.
- In use cases with FSB, it is noticed that after converting dynamic port-channel to static, hosts and targets don't see each other.
- When an RBridge is removed from VCS cluster, it has to be manually removed from fcoe fabric-map.
- In NETWORK OS V6.0.1a1, up to four FCoE Vlan's are supported in VDX . But on a single VDX, All member ports in a LAG have to be configured with the same FCoE Vlan. Different LAG can be configured with different FCoE Vlan.
- In NETWORK OS V6.0.1a1, it is recommended user define different fabric-map for Remote Logical SAN and Local Logical SAN configuration. If user configures a fabric-map to work on Remote Logical SAN first and then later change the same fabric-map to become Local Logical SAN, it may cause FCoE port continuously flapping.

FlexPorts

- The port transceiver optic type must match the configured Flexport type. If a port is configured as Fibre Channel then an appropriate FC SFP+ transceiver must be used; likewise when the port is configured as an Ethernet port then an appropriate Ethernet SFP+ transceiver must be used. The same applies to QSFP+ transceivers – the transceiver type must match the configured Flexport type for the QSFP links.
- Only Brocade-branded SFPs are supported.

- Setting the connector-group speed to HighMixed allows only one FC port speed (16G) but the port speed configuration can still be set to auto.
- Changing the connector-group speed always disrupts any other active ports in the connector-group, independent of configured Flexport type.
- The FCoE Base license is required to enable any FibreChannel interface.

Fibre Channel

- F-Port can support only 63 NPIV devices.
- Loop devices are not supported.
- Long distance is not supported on Fibre Channel ports.
- Proprietary features such as QoS, D-Port, FAPWWN are not supported on Fibre Channel ports.
- Credit Recovery is supported on Fibre Channel ports.
- FEC is supported on Fibre Channel E/Ex ports only (no support on F/N ports).
- Trunking is not supported on Fibre Channel ports running at 2G or 4G speeds.
- Fibre Channel trunks are limited to 2 FC trunks per port group.
- Fibre Channel trunks only form with ports running at the same speed.

Access Gateway

- AG enable/disable command is moved to configuration mode in 6.0.1a1. From principal node AG mode can be changed on individual nodes under rbridge-id-ag configuration mode.
- All AG configurations have to be done under rbridge-id-ag sub mode. Prefix "ag" is not allowed any more.
- The switch can be operated as Fibre Channel Forwarder (FCF) by disabling Access Gateway mode.
- AG does not bridge the VCS and SAN fabrics because hosts connected to the AG switch are registered in the SAN name-server only. Therefore, all zoning operations for AG are done on the SAN fabric.
- At least one N-port must be online in order for FCoE devices to log in.
- After enabling Remote Logical SAN on AG switch, FCoE devices connected to AG switch will not login with "fcoeport default" provisioning and needs to be configured as "fcoeport <logical-san>".

ND/RA

- Proxy ND is not supported.

BFD

- Static Route BFD, BGP BFD and OSPFv2/v3 BFD
 - For Single HOP BFD sessions configured with source IP as secondary IP is not supported, since significance of Source IP in BFD configuration is only to determine on which interface BFD session should be started and hence interfaces' Secondary IP is not used as source in BFD PDU.
 - BFD is not supported on leaked routes.
 - BFD for multi-HOP BFD neighbor reachable via multiple paths with equal cost(ECMP) will not be supported since BFD requires BFD session to be created for the neighbor for each ECMP path.
 - BFD is not supported for OSPFv2 & OSPFv3 virtual links.
 - For single hop BFD sessions, BFD will consider the interval values that are configured on interface, and not the non-default values that are configured within the global command.
- BFD for VxLAN tunnels
 - BFD session may not come online or may flap if VCS cluster is in transient state during reload, vLAG failover, fabric split, chassis disable/enable and such scenarios. It is required to have a stable VCS cluster in order for BFD sessions on VxLAN tunnels to work as expected.
 - BFD parameters are not configurable on VCS VxLAN Gateway for Type NSX. The parameters are derived from NSX controller.

VRRP

- VRRP and VRRP-E cannot be enabled together. Command "protocol vrrp-extended" is added to specifically enable VRRPE.
- IPv6 and IPv4 VRRP sessions cannot be configured with the same VRRP group-ID on the same Layer3 interface.
- If an IPv6 VRRP session is configured with only global VIP address without Link-Local VIP, VIP configuration will fail for that session during download of configuration from file.
- "show vrrp summary" and "show ipv6 vrrp summary" will display all sessions in default vrf. In earlier NOS versions, these commands displayed sessions across all vrf.

Fabric Virtual Gateway (FVG)

- FVG co-existence with VRRP/VRRP-E
 - FVG ipv4 or FVG ipv6 with default global mac cannot be enabled with VRRP but can be enabled with VRRPE-E.
 - FVG ipv4 or FVG ipv6 with non-default global mac cannot be enabled either with VRRP or VRRPE-E.

OSPF v3

- OSPFv3 HA with Graceful restart is not supported but GR-helper mode functionality is supported. VRF-Lite-Capability CLI and support for Down bit handling is not available in OSPFv3 as in OSPFv2. When the BGP4+ route from the MPLS cloud is redistributed into OSPFv3 domain the redistributed route is always installed in the OSPFv3 routing table.

BGP

- Conditional advertisement of default-route using route-map match prefix not supported.
- Over a link-local eBGP session, updates are not carrying the new nexthop that is set using a route-map.
- EBGP TTL Security Hack Protection is not supported.

ACL

- L2 User ACL deny rule can prevent trapping of L3 control frames.
- IPv6 ACLs at ingress are not applicable for packets with Link local source address.
- ACL Logging at egress can impact forwarding traffic at high rates.
- Counters for hard-drop ACLs may not count accurately.
- Statistics are not supported for hard-drops at Egress.
- For Private VLANs, Egress ACLs on Primary VLAN is applied only for all traffic which ingresses primary VLAN i.e.
 - If the traffic ingresses from Primary VLAN but gets translated to Secondary VLAN at egress, ACL on primary VLAN at egress is still applicable to it.

- If the traffic ingresses from Secondary VLAN but gets translated to Primary VLAN at egress, ACL on primary VLAN at egress is still not applicable to it.

L2/L3 Multicast

- The following PIM features are not supported in this release:
 - IP version 6
 - VRF
 - Configuring the switch as the BSR (Bootstrap Router) candidate.
 - Configuring the switch as the Rendezvous Point or Rendezvous Point candidate. The RP must be configured outside the VCS cluster.
- In fabric cluster mode, IGMP Snooping must be enabled in all the switches in the VCS cluster.
- Statistics for MLDv1 is done on a VLAN basis across VCS.
- Multiple IP subnetting support: PIM FHR and LHR operation are not supported on secondary subnets.

VRF

- Under VRF submode there is a syntax change for the address-family ipv4 command.

Old format	: address-family ipv4 [max-route <value>]
New format	: address-family ipv4 unicast max-route <value>

Note: "max-route" command is now moved to address - family submode.
- There is no provision to configure "max-routes" for default-vrf.
- There is no use case for "rd" configuration in VRF and this command will be deprecated in next release.
- On configure VRF on an interfaces, all previous IP config on that interface will be deleted.
- IP Services like telnet are supported only on mgmt-vrf.
- User will be able to access VDX switches only through interfaces belonging to mgmt-vrf.
- Removing VRF address family on a non-default VRF will delete all relevant address-family configurations including the interface and protocol configuration for that VRF.

- Support added for SNMP Infrastructure in non-default VRFs (the supported number of VRFs is 6).
- Support added for SSH in non-default VRFs (the supported number of VRFs is 6).

BGP-VRF

- Local - as <num> can be configured for particular VRF under "address-family ipv4 unicast vrf <vrfname>" and is not supported under "address-family ipv6 unicast vrf <vrf-name>".
- "maxas-limit in" can be configured for particular VRF under "address-family ipv4 unicast vrf <vrfname>" and is not supported under "address-family ipv6 unicast vrf <vrf-name>".
- When route-map is applied to BGP, and route-map has multiple 'set ip next-hop' statements in a single instance, BGP will consider the last 'set ip next-hop' in the route-map.

Policy-based Routing (PBR)

- If a PBR route-map is applied to an interface that is actively participating in a control protocol and the ACL specified in the route-map also matches the control protocol traffic the control protocol traffic will be trapped to the local processor and not be forwarded according to the route-map.

Inter-VRF Leaking (Static)

- S+ symbol routes indicates leaked routes.
- VRF route leak cascading is not supported– only one level of indirection.
- User should avoid making Static, dynamic and connected route conflict with routes in target VRF when configuring route leak.
- For bidirectional traffic with router leak, user needs to configure route leak in both direction separately.
- Route leak configuration to next hop IP on the same box on different VRF is not a valid configuration, but CLI will be accepted.
- Precaution needs to be taken when leaking default routes - this can result in routing loops.
- Switch management from non-management VRF by leaking route from non-management to management VRF is not supported.

DHCP IP Helper

- There is no HA support for dhcp relay statistics. When a switchover happens, the statistics will not be replicated to the new active MM.

Dynamic ARP Inspection (DAI)

- The arps learnt on trusted ports would be deleted when DAI is enabled or DAI filter changed.
- Static arps not permitted by DAI filter would be promoted to active state. Administrator is responsible for configuring static ARPs in sync with DAI acls.
- ARP packets more than 190 bytes on a DAI enabled vlan will be dropped.
- ARP access-list with longer names is not effective (greater than 20 characters)

DHCP-based Firmware download (DAD-DHCP Automatic Deployment)

- In order for successful version upgrade using DAD method, switch should undergo 2 reloads. For switch in factory default, there is additional reboot to cancel bare metal mode.
- If firmware download is skipped only config download is allowed.
- For dual MM chassis, dual MM must be in sync for DAD to function.
- DAD is a disruptive.
- In FIPS mode, DAD is not supported.
- Cluster principal node failover is not supported.
- DAD over in-band is not supported. Virtual Fabrics is not supported with DAD. You must disable Virtual Fabrics before starting the DAD process in the global configuration file or in the script.
- DAD must complete and succeed on Principle node alone before turn on power for all secondary nodes.
- When the switch is in Factory default, DAD is enabled upon power up the switch
- DAD executes only if the switch configuration is the default configuration. If the configuration on the switch is not the default configuration, DAD exits.

- If the switch is in the default configuration before DAD is triggered, DHCP will remain enabled after the deployment completes. However, this setting can be overwritten by the switch-specific configuration file or the dad.py script.
- You must enable DHCP in the DCMD default configuration to ensure that the switch receives its IP address from the preconfigured DHCP server.
- The factory default DAD runs only once in a DHCP-enabled environment. Irrespective of whether this process is a success or failure, DAD will not be triggered again after a reboot or power off. You can run DAD manually using the dhcp auto-deployment enable command if required.
- Must set ztp=0 in dad configuration file since ZTP (Zero Touch Provisioning) is enabled by default.
- The "vcsmode" value in dad.conf MUST be set to "LC" regardless of whether the existing cluster is in LC or FC mode. If "vcsmode" set to "FC" value in dad.conf, the DAD request can fail.
- DAD is enabled automatically upon switch reboot when you use NOSCLI "write erase" command.

Zero Touch Provisioning (ZTP) consideration

All nodes can either be powered up at the same time or enabled from the CLI. This is the key difference vs regular DAD process.

Link State Tracking

- The "track enable/disable" command can only be used to enable or disable the tracking. In order to remove tracking configuration from internal database for a given interface "track remove all" command must be used.
- [UPDATED for 6.0.1a] When there is no uplink interface configured, the track disable command will remove tracking configuration from internal database and this behavior is applicable only in 6.0.1a patch and not in prior releases. If the "track min-link" number is greater than number of uplinks, then the downlink will be shut down with a warning message

OpenFlow

- Once an interface becomes OpenFlow enabled, very limited set of conventional commands are allowed which includes some of the QoS related configuration commands. For complete list of allowed commands please refer to "NETWORK OS V6.0.1a SDN Configuration Guide"
- Priority-tagged frames are not supported.
- L3 Generic flows (incoming port as "wildcard") are not supported.

- PUSH/POP operations can only be associated with action type OFPAT_OUTPUT inside a flow-mod.
- Type of an existing GROUP cannot be changed.
- Existing "clear counter all" command applies to OpenFlow ports as well.
- Pre-provisioned flow-mods will not be replayed to a new slot coming online. GROUP and METER configurations will be always replayed.
- On Mercury, queue statistics should be interpreted as wire-vlan (COS) priority statistics.
- Actual number of supported flow-mods (L2/L3) may be less since MAX scale values include per port default miss entries, and single LLDP entry is needed for topology discovery. This applies to all supported platforms.
- [UPDATED for 6.0.1a] For layer 3 rules, switch can't differentiate between tagged and untagged flows when matching against rules. This applies to all supported platforms.
- [UPDATED for 6.0.1a] Filtering options are not supported for show openflow CLIs. Show openflow commands with filter option show the complete output.
- [UPDATED for 6.0.1a] For the port based flow mod, if the ivid reference is active, egress tagging is not cleared. The new flow mod will not be installed If the previous flow mod has created the egress tagging behavior. This case has to be handled by work-around flow mods or take the port off from openflow and bring it back.
- [UPDATED for 6.0.1a] With default rcv-queue and after coldboot, group select traffic may not be correct, need to do shut/no shut on the interface. This issue is not there with non-default rcv-queue.
- [UPDATED for 6.0.1a] With large number of flows, "show openflow flow <>" may take 20 seconds to display packet counts.

Auto QoS for NAS

- From Network OS v5.0.1 onwards, 'nas auto-qos' configuration appears below 'cee-map' configurations in running-config. In earlier versions, it was the other way round.
- As a result of this, if file replay is done using the Network OS v6.0.0 config (with auto-nas configuration) on any previous version (say, Network OS v4.1.0), 'nas auto-qos' configuration will be lost. User will have to reconfigure 'nas auto-qos' configuration manually.

REST API

- REST configuration for startup-config datastore is not supported.
- Only one command can be configured with one REST request. Configuring multiple commands in a single request is not supported.
- Versioning in API is not supported
- Pagination and Range is not supported.
- Higher level of resource can be obtained with the header "-H "Resource-Depth: x".
- Action related operational commands are not supported.
- Maximum 30 sessions are supported.
- [UPDATED for 6.0.1a] An FCoE Base license is required for the FCoE device to log in. Each FCoE device must have a VF port to log in.

NetConf

- Netconf configuration for startup-config datastore is not supported
- Configuring multiple commands in a single request is supported for configuration/deletion of vlan, switch port, trunk port, VE and rules under IP ACL only.
- Range is not supported.
- Maximum 16 sessions supported.

VXLAN Gateway for VMware NSX

- VCS VXLAN Gateway for NSX is supported only in the VCS Logical Chassis mode.
- A maximum of 4 RBridges are supported in a VXLAN enabled VCS Cluster. VXLAN Gateway should be enabled on all the RBbridges of the VCS Cluster.
- Only 1 VTEP Gateway is supported in a VXLAN enabled VCS Cluster.
- VxLAN GW for VMware NSX and VF Extension cannot be enabled in the same VCS fabric.
- VMware NSX vSwitch with vSphere version 5.5 (ESXi 5.5), XenServer 6.2, and KVM on Ubuntu 12.04 are supported as hypervisors.
- Only one-to-one VLAN to VNI mapping is supported.

- Service and Transport VF cannot be attached to VxLAN GW.
- Tunnel interfaces cannot be used as SPAN (Switch port Analyzer) destination.
- Only Ingress ACL can be applied on tunnels.
- Ingress/Egress QoS policies cannot be applied to tunnels.
- Unicast/Multicast routing between VXLAN and VLAN/VXLAN is not supported.
- When using the command "show overlay-gateway name <name> VLAN statistics" for debugging overlay-gateway VLANs, it should be noted that the statistics information is limited to 256 VLANs (Rx) and 116 VLANs (Tx).
- When multiple VMware NSX Service Nodes are setup, only one of the node would be used for handling BUM traffic. During service node failover scenarios another Service node would be selected for handling BUM traffic, if BFD is enabled for all the Service nodes.
- ALL the VE interfaces should run VRRP-E with the same VRID and same virtual-mac to terminate the incoming packets on other VLANs.
- Tunnels egressing/ingressing through an ISL port is not supported.
- Fabric-Virtual-Gateway (FVG) based VTEP is not supported. CLIs for configuring FVG as VTEP are available under overlay-gateway, however these CLIs should not be used as the functionality is not available in this release.
- In-band management for VCS VxLAN GW with default-vrf is not supported.
- Load balancing between multiple Service node tunnels is not supported.

TCAM Profiles

- The number of routes the user can create may not match the max scale numbers due to reserved routes/entries which are created for internal use.

Management VRF

OSPF/BGP/PIM/VRRP/VRRPe is not supported on Management VRF. The following are not supported on in-band ports when they are part of Management VRF:

- DHCP Client functionality
- Auto-config address
- Out-of-band management ports can only be part of Management VRF.
- In-band management ports can be part of Management VRF or default VRF.

- Switch cannot be managed from leaked routes pointing to Management-VRF.
- Address family on Management VRF cannot be removed.
- Firmware download/supportsave is not supported on in-band ports. This limitation is applicable for Management vrf/default-vrf/ non-default vrf.

Conversational MAC Learning

- Conversational MAC Learning and 'Disable Source MAC Learning' cannot be enabled simultaneously.

System level Flowbased QoS

- System Flow based QoS is not supported on the Egress direction.
- QoS can operate on either of three modes - MLS, CEE and MQC. Hence once service-policy is configured, the interface will be assumed to be in MQC mode and existing MLS and CEE commands will not be supported on the same interface. Un-configuring the policy will put the interface back to default mode which is MLS.
- For Policer, aggregation is possible only within a single chip. Hence when policer is applied on port-channel interface, multi-chip aggregation is not expected.
- SFLOW as action is not supported on Port-Channel interface.
- Any ACL that is used in Flowbased QoS class-map as a match criterion is considered as "QoS ACL" and is special in nature. Hence behavior in some aspects may differ from that of regular "User ACL".
- System based QoS is not supported in egress direction.

Port level Flowbased QoS

- Policer action or SPAN action or both can be applied in egress direction for Port Level Flowbased QoS.
- No other QoS actions are supported in egress direction for port level flowbased QoS.

Non-trivial Merge

- Non-trivial merge is not supported for global configuration. There are a few exceptions in Local configuration as well which are not supported for non-trivial merge. This is because these configurations modify global configuration indirectly.

- Modifying the local configurations listed below will result in both a local and global configuration change thereby causing configuration mismatch when ISLs are brought up during fabric formation resulting in node segmentation.

Command (Local Configuration)	Description
/hardware/flexport <interface tuple>/type fibre-channel	Converting an Ethernet interface to Fibre-Channel causes global configuration changes because the Ethernet interface can have configurations in these global configs L2Sys, SPAN, IGMPs, MLDs.
/rbridge-id <#>/vrf <name>	The creation of a VRF on an RBridge will internally create a global partition object which is not visible to the user and used to track the same VRFs created across rbridges in the cluster.

Interoperability

- In a VPC environment where the Brocade VDX side has the active LACP settings and the Cisco side has the passive settings on the vLAG, the port-channel takes over 30 seconds to come up.
Workaround: Reverse the settings and have the Brocade VDX LACP settings passive and the Cisco side set as active. The port channel will then restore after about 10 seconds.
- VDX interop with Cisco Nexus switch with 'peer-switch' enabled on VPC is not supported.
- When interoperating with Brocade 8000, it is recommended to set the mac-aging time to 0 on the VDX switch to prevent any adverse impact caused by certain errors generated by the Brocade 8000.
- ADX HA Sync packets use UDLD PDU's which may be dropped by VDX . To enable forwarding, we recommend configuring dot1q tagging to treat UDLD packets as datapackets to be forwarded across VCS.Virtual Fabric.
- PIM-SM is not supported on Virtual Fabric.
- For frames forward on a transport fabric, ingress CTAG tagging is preserved at the egress port regardless of the egress tagging classification.
- Default-VLAN can only be configured using TRANSPORT VF IDs.
- The "no vcs virtual-fabric enable" command execution time is dependent on the number of ISLs and VLANs in the VCS.
- The virtual-fabric resource allocation are platform dependent as follows:
 - uses TCAM table

MAPS

- Only BNA v12.4.2 (to be made available separately) supports NOS v6.0.1. It is required to first upgrade to BNA v12.4.2 and then upgrade the switches to NOS v6.0.1.
- Only one MAPS policy can be active at any time.
- All MAPS thresholds, policies, rules and groups are pre-defined in NETWORK OS v6.0.1a and may not be modified.
- MAPS port level alerting in NETWORK OS v6.0.1a1 is not available for FlexPorts configured in Fibre Channel mode.
- MAPS configuration and monitoring is applicable per switch, so users need to apply configuration on each switch being monitored.
- Rules for RX_SYM_ERR are triggered when breakout cable is connected on a 40G interface not configured for breakout.

Miscellaneous

- Brocade VDX switches load balance internal and external traffic based on hash functions using standard network headers as keys. Due to this implementation, users may experience traffic imbalance depending upon application flow definition.
- Packet drops will be seen for a short duration due to routing changes with link flaps and/or node failovers.
- On both ISL and Edge ports, sFlow sampling is supported only in inbound direction.
- SFlow collectors are not queried in snmp v1, v2 & v3 versions
- If multiple VLANs are configured on a switch, then in order to enable certain features such as IGMP or PVST it is recommended that specific features be enabled on a per-VLAN basis instead of enabling them globally.
- The VLANs 4093, 4095 and 1002 are reserved and used for internal cluster operations.
- "clear ip route all" need to be issued once the maximum number of routes supported by a router is exceeded.
- SNMPset operation is supported for certain MIB objects
- SNMP supports 2k OCTET-STRING size for MIB objects.

- Snmpwalk operation on TCP MIB (RFC 4022) may become very slow and timeouts may happen on all VDX platforms . The snmpwalk timeout should be set to at least 3 seconds while walking the TCP MIB.
- Under rare conditions, the switch may bootup with default configuration on power-cycling the switch.
- Firmware downgrade is not blocked if the scale configured would not be supported in the downgraded release
- Please make sure to not have large no of unreachable TACACS+ accounting server configured, else it might cause unit to reboot. This issue is hit only with large config (4K vlan etc. and 20K lines or config)
- Configuration of more than one In-band management port on a single switch is not recommended.
- Under certain stress conditions the 'copy support' command might time out for some modules. In such cases, it is recommended to retry 'copy support' with a higher timeout multiplier value.
- It is highly recommended to copy the configuration file to running-config and then save the running-config to startup-config, instead of directly copying the external configuration file to startup-config, especially when using fabric distributed features such as Zoning, VM Aware Network Automation and Virtual IP.
- To replace the failure parts with new one during OS operating if a switch becomes damaged and no longer be used, the management module firmware for Hitachi ComputeBlade 2500 version A0145 or higher version must be required. If the version of the management module is lower than A0140, the failure parts replacement during OS operation is NOT supported. For detail, refer to the management module user's guide.
- The followings show unsupported features for the embedded DCB switch module in NOS v6.0.2 against Brocade Network OS release.
 - NOS firmware update via ISSU
 - Sending SYSLOG messages via in-bound management is not supported.

Bug Fixes and Known Issues

In this chapter, Bug Fixes and Known Issues for this release are described.

- [Bug Fixes](#)
- [Known Issues](#)

Bug Fixes

Regarding NOS Bug Fix information on NOS v6.0.2, the embedded DCB switch module (GG-BE4LSW2X1) is equivalent to the VDX 6740x except where noted in this section. For detail information of the bug fixes from NOS v6.0.1a1, refer to the following release notes (available on MyBrocade: <http://my.brocade.com/>)

- Release note for NOS v6.0.2 : nos6.0.2_releasenotes_vXX.pdf

Known Issues

Regarding NOS Known Issue information on NOS v6.0.2, the embedded DCB switch module (GG-BE4LSW2X1) is equivalent to the VDX 6740x except where noted in this section. For detail information of the Known Issue of the NOS v6.0.2 release, refer to the following release notes (available on MyBrocade: <http://my.brocade.com/>)

- Release note for NOS v6.0.2 : nos6.0.2_releasenotes_vXX.pdf

Additionally, this section lists the known issues of embedded DCB switch module (GG-BE4LSW2X1) with NOS v6.0.2.

Item	Descriptions	
#1	Summary	On the Link State Tracking feature, internal ports become online even though one (or some) track port(s) is (are) offline after the switch reload.
	Symptom	The Link State Tracking feature allows for the monitoring and detection of the link state of external port(s) (tracked port), and the control of a/some internal ports, which is brought to a 'down' state (offline) when the tracked external port(s) goes to a 'down' state (offline). Even though the internal port goes to offline due to this feature, internal port becomes online after the switch reload.
	Technical Severity	Medium
	Workaround	Avoid reloading the switch when the link state of tracked external port(s) is (are) the 'down' state.
	Reported in release	NOS v6.0.1a1

Item	Descriptions	
#2	Summary	Trap packets from the switch has IPv6 address of management port for the out-bound management even though the Trap packets issued from the management interface via in-bound management.
	Symptom	Trap packets from the switch has IPv6 address of management port for the out-bound management even though the Trap packets issued from the management interface via in-bound management.
	Technical Severity	Medium
	Workaround	Source IPv6 address on the Trap packets should be replaced by IPv6 address for the in-bound management port as necessary.
	Reported in release	NOS v6.0.2

Item	Descriptions	
#3	Summary	Applying ACL on Ethernet management interface fails with "% Error: Internal Error" error message
	Symptom	The issue is observed when an attempt is made to apply ACL policy on Ethernet management interface with the following conditions: <ul style="list-style-type: none"> - An ACL policy is already enforced on the Ethernet management interface - The new ACL policy and enforced policy names differ only in letter cases (e.g. "TEST001" vs "Test001").
	Technical Severity	Medium
	Workaround	Before a new access-list is added to a management port, the access lists must be removed first. Or, create ACLs with distinctive policy names (e.g. with different letters, numbers etc.).
	Reported in release	NOS v6.0.2

Hitachi Data Systems

Corporate Headquarters

2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.
www.hds.com

Regional Contact Information

Americas

+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa

+44 (0) 1753 618000
info.emea@hds.com

Asia Pacific

+852 3189 7900
hds.marketing.apac@hds.com

