

Network OS Administrator's Guide

Supporting Network OS v5.0.1

Contents

Preface	11
Document conventions.....	11
Text formatting conventions.....	11
Command syntax conventions.....	11
Notes, cautions, and warnings.....	12
Brocade resources.....	12
Contacting Brocade Technical Support.....	13
Brocade customers.....	13
Brocade OEM customers.....	13
Document feedback.....	13
About This Document	15
Supported hardware and software.....	15
Using the Network OS CLI	15
What's new in this document.....	15
Introduction to Network OS and Brocade VCS Fabric Technology	17
Introduction to Brocade Network OS.....	17
Brocade VCS Fabric terminology.....	18
Introduction to Brocade VCS Fabric technology.....	18
Automation.....	19
Distributed intelligence.....	20
Logical chassis.....	21
Ethernet fabric formation.....	22
Brocade VCS Fabric technology use cases.....	23
Classic Ethernet access and aggregation use case.....	23
Large-scale server virtualization use case.....	25
Brocade VCS Fabric connectivity with Fibre Channel SAN.....	26
Topology and scaling.....	26
Core-edge topology.....	27
Ring topology.....	27
Full mesh topology.....	28
Basic Switch Management	31
Switch management overview.....	31
Connecting to a switch.....	31
Telnet and SSH overview.....	31
SSH server key exchange and authentication.....	32
Feature support for Telnet.....	32
Feature support for SSH.....	33
Firmware upgrade and downgrade considerations with Telnet or SSH.....	33
Using DHCP Automatic Deployment.....	33
Telnet and SSH considerations and limitations.....	36
Ethernet management interfaces.....	36
Brocade VDX Ethernet interfaces.....	36
Lights-out management.....	37
Stateless IPv6 autoconfiguration.....	37
Switch attributes.....	37

Switch types.....	37
Operational modes.....	38
Logical chassis cluster mode.....	38
Fabric cluster mode.....	41
Modular platform basics.....	41
Management modules.....	42
Switch fabric modules.....	42
Line cards.....	43
Supported interface modes.....	43
Slot numbering and configuration.....	43
Slot numbering.....	43
Slot configuration.....	43
Connecting to a switch.....	44
Establishing a physical connection for a Telnet or SSH session.....	44
Telnet services.....	45
Connecting with SSH.....	46
Using the management VRF.....	48
Configuring and managing switches.....	48
Configuring Ethernet management interfaces.....	49
Configuring a switch in logical chassis cluster mode.....	55
Configuring a switch in fabric cluster mode.....	64
Displaying switch interfaces.....	65
Displaying slots and module status information.....	66
Replacing a line card	66
Configuring High Availability.....	67
Disabling and enabling a chassis.....	68
Rebooting a switch.....	69
Troubleshooting switches.....	69
Configuring policy-based resource management.....	71
Configuring hardware profiles.....	73
Guidelines for changing hardware profiles.....	73
Using hardware profile show commands.....	74
Brocade support for OpenStack.....	74
Configuring OpenStack to access Network OS.....	75
Mixed-version fabric cluster support.....	75
Using Network Time Protocol.....	77
Network Time Protocol overview.....	77
Date and time settings.....	77
Time zone settings.....	77
Configuring NTP.....	77
Configuration considerations for NTP.....	77
Setting the date and time.....	78
Setting the time zone.....	78
Displaying the current local clock and time zone.....	79
Removing the time zone setting.....	79
Synchronizing the local time with an external source.....	79
Displaying the active NTP server.....	79
Removing an NTP server IP address.....	80
Configuration Management.....	81

Configuration management overview.....	81
Configuration file types.....	81
Displaying configurations.....	82
Displaying the default configuration.....	82
Displaying the startup configuration.....	83
Displaying the running configuration.....	83
Saving configuration changes.....	83
Saving the running configuration.....	83
Saving the running configuration to a file.....	83
Applying previously saved configuration changes.....	84
Backing up configurations.....	84
Uploading the startup configuration to an external host.....	84
Backing up the startup configuration to a USB device.....	85
Configuration restoration.....	85
Restoring the default configuration.....	85
Managing configurations on a modular chassis.....	86
Managing configurations on line cards.....	86
Managing configurations across redundant management modules.....	86
Managing configurations in Brocade VCS Fabric mode.....	87
Automatic distribution of configuration parameters.....	87
Downloading a configuration to multiple switches.....	87
Rejoining an offline node to a logical chassis cluster.....	88
Managing flash files.....	89
Listing the contents of the flash memory.....	89
Deleting a file from the flash memory.....	89
Renaming a flash memory file.....	89
Viewing the contents of a file in the flash memory.....	89
Configuring SNMP.....	91
Simple Network Management Protocol overview.....	91
SNMP Manager.....	91
SNMP Agent.....	91
Management Information Base (MIB).....	91
Basic SNMP operation.....	91
SNMP configuration.....	92
Configuring SNMP community strings.....	92
Configuring SNMP server hosts.....	93
Configuring multiple SNMP server contexts.....	95
Configuring SNMP server views.....	96
Configuring SNMP server groups.....	96
Configuring SNMP server users.....	96
Configuring SNMP server v3hosts.....	98
Managing SNMP access rights using ACLs	98
Displaying SNMP configurations.....	100
Configuring Brocade VCS Fabrics	101
Fabric overview.....	101
Brocade VCS Fabric formation.....	101
How RBridges work.....	102
Neighbor discovery.....	102
Brocade trunks.....	102

Fabric formation.....	103
Fabric routing protocol	104
Configuring a Brocade VCS Fabric.....	104
Adding a new switch into a fabric.....	105
Configuring fabric interfaces.....	106
Configuring broadcast, unknown unicast, and multicast forwarding.....	107
Configuring VCS virtual IP addresses.....	108
Configuring fabric ECMP load balancing.....	110
Configuring Metro VCS.....	113
Metro VCS overview.....	113
Metro VCS details and configuration.....	113
Metro VCS using long-distance ISLs.....	114
Metro VCS using standard-distance ISLs.....	116
Metro VCS combined with vLAGs.....	118
Configuring a long-distance ISL.....	120
Configuring interconnected Ethernet Fabrics.....	121
Administering Zones.....	125
Zoning overview.....	125
Example zoning topology.....	125
LSAN zones	127
Managing domain IDs.....	128
Approaches to zoning.....	129
Zone objects.....	129
Zoning enforcement.....	130
Considerations for zoning architecture.....	131
Operational considerations for zoning.....	131
Configuring and managing zones	132
Zone configuration management overview.....	132
Understanding and managing default zoning access modes.....	133
Managing zone aliases.....	134
Creating zones.....	137
Managing zones.....	140
Zone configuration scenario example.....	146
Merging zones.....	148
Configuring LSAN zones: Device-sharing example.....	153
Configuring Fibre Channel Ports.....	157
Fibre Channel ports overview.....	157
Connecting to an FC Fabric through an FC Router.....	158
Fibre Channel port configuration.....	158
Using Fibre Channel commands.....	158
Activating and deactivating Fibre Channel ports.....	159
Configuring and viewing Fibre Channel port attributes.....	160
Configuring a Fibre Channel port for trunking.....	162
Monitoring Fibre Channel ports.....	162
Using Access Gateway.....	165
Access Gateway basic concepts.....	165
Switches supported for Access Gateway.....	165
Network diagrams.....	166

Access Gateway and native VCS modes.....	168
Access Gateway in a logical chassis cluster.....	169
Access Gateway ports.....	169
Access Gateway features and requirements.....	172
Enabling Access Gateway mode.....	175
Disabling Access Gateway mode.....	175
Displaying Access Gateway configuration data.....	176
VF_Port to N_Port mapping.....	178
Displaying port mapping.....	178
Configuring port mapping.....	180
Port Grouping policy.....	182
Displaying port grouping information.....	183
Creating and removing port groups.....	183
Naming a port group.....	184
Adding and removing N_Ports in a port group.....	185
Port Grouping policy modes.....	186
Trunking in Access Gateway mode.....	188
Setting up trunking for Access Gateway.....	189
Access Gateway under FlexPort.....	189
Configuring Access Gateway under FlexPort	189
Restoring N_Port login balance	190
N_Port monitoring for unreliable links.....	190
Setting and displaying the reliability counter for N_Port monitoring.....	191
Displaying Access Gateway N_Port utilization data	191
Using System Monitor and Threshold Monitor.....	193
System Monitor overview.....	193
Monitored components.....	193
Monitored FRUs.....	193
Configuring System Monitor.....	194
Setting system thresholds.....	194
Setting state alerts and actions.....	195
Configuring e-mail alerts.....	195
Viewing system SFP optical monitoring defaults.....	196
Displaying the switch health status.....	196
Threshold Monitor overview.....	197
CPU and memory monitoring.....	197
SFP monitoring.....	198
Security monitoring.....	199
Interface monitoring.....	200
Configuring Threshold Monitor.....	200
Viewing threshold status.....	201
CPU and memory threshold monitoring.....	201
Configuring SFP monitoring thresholds and alerts.....	202
Security monitoring.....	203
Configuring Interface monitoring.....	203
Pausing and continuing threshold monitoring.....	204
Using VMware vCenter	205
vCenter and Network OS integration overview.....	205
vCenter properties.....	205

vCenter guidelines and restrictions.....	205
vCenter discovery.....	206
vCenter configuration.....	206
Step 1: Enabling QoS.....	206
Step 2: Enabling CDP/LLDP	207
Step 3: Adding and activating the vCenter.....	207
Discovery timer interval	208
User-triggered vCenter discovery.....	208
Viewing the discovered virtual assets.....	209
Configuring Remote Monitoring.....	211
RMON overview.....	211
Configuring and managing RMON.....	211
Configuring RMON events.....	211
Configuring RMON Ethernet group statistics collection.....	212
Configuring RMON alarm settings.....	212
Monitoring CRC errors.....	213

Copyright Statement

© 2014, Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, the B-wing symbol, Brocade Assurance, ADX, AnyIO, DCX, Fabric OS, FastIron, HyperEdge, ICX, MLX, MyBrocade, NetIron, OpenScript, VCS, VDX, and Vyatta are registered trademarks, and The Effortless Network and the On-Demand Data Center are trademarks of Brocade Communications Systems, Inc., in the United States and in other countries. Other brands and product names mentioned may be trademarks of others.

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

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

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

Preface

- Document conventions..... 11
- Brocade resources..... 12
- Contacting Brocade Technical Support..... 13
- Document feedback..... 13

Document conventions

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

Text formatting conventions

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

Format	Description
bold text	Identifies command names Identifies keywords and operands Identifies the names of user-manipulated GUI elements
<i>italic text</i>	Identifies text to enter at the GUI Identifies emphasis Identifies variables
Courier font	Identifies document titles Identifies CLI output Identifies command syntax examples

Command syntax conventions

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

Convention	Description
bold text	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, --show WWN .
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options. In Fibre Channel products, square brackets may be used instead for this purpose.

Convention	Description
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member[member...]</i> .
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Notes, cautions, and warnings

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

NOTE

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

ATTENTION

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



CAUTION

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



DANGER

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

Brocade resources

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

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

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

Release notes are available on [MyBrocade](#) under Product Downloads.

White papers, online demonstrations, and data sheets are available through the [Brocade website](#).

Contacting Brocade Technical Support

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

Brocade customers

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

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

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

Brocade OEM customers

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

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

Document feedback

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

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

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

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

About This Document

- Supported hardware and software.....15
- Using the Network OS CLI15
- What's new in this document.....15

Supported hardware and software

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

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for Network OS 5.0.1, documenting all possible configurations and scenarios is beyond the scope of this document.

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

- Brocade VDX 2740 embedded switch
- Brocade VDX 6740
 - Brocade VDX 6740-48
 - Brocade VDX 6740-64
- Brocade VDX 6740T
 - Brocade VDX 6740T-48
 - Brocade VDX 6740T-64
 - Brocade VDX 6740T-1G
- Brocade VDX 8770
 - Brocade VDX 8770-4
 - Brocade VDX 8770-8

To obtain information about an OS version other than Network OS v5.0.1, refer to the documentation specific to that OS version.

Using the Network OS CLI

For complete instructions and support for using the Network OS v5.0.1 command line interface (CLI), refer to the *Network OS Command Reference*.

What's new in this document

This document supports the following features introduced in Network OSv5.0.0:

- Mixed-version fabric support
- Network Time Protocol (NTP) enhancements
- Access gateway enhancements

Beginning with Network OS v5.0.0, there are now five books that cover Network OS administration:

- Network OS Administration Guide

- Network OS Layer 2 Switching Configuration Guide
- Network OS Layer 3 Routing Configuration Guide
- Network OS Security Configuration Guide
- Network OS Troubleshooting Guide

This document supports the enhancements introduced in Network OSv5.0.1b and v5.0.2.

For complete information, refer to the Network OS Release Notes.

Additional documents for Network OS include the following:

- Network OS Command Reference
- Network OS Upgrade Guide
- Network OS Software Licensing Guide
- Network OS Message Reference
- Network OS Feature and Support RFC Matrix
- Network OS NETCONF Operation's Guide
- Network OS YANG Reference
- Network OS REST API Reference

Introduction to Network OS and Brocade VCS Fabric Technology

- [Introduction to Brocade Network OS](#)..... 17
- [Introduction to Brocade VCS Fabric technology](#)..... 18
- [Brocade VCS Fabric technology use cases](#)..... 23
- [Topology and scaling](#)..... 26

Introduction to Brocade Network OS

Brocade Network OS is a scalable network operating system available for the Brocade data center switching portfolio products, including the VDX product line.

Purpose-built for mission-critical, next-generation data centers, Network OS supports the following capabilities:

TABLE 1 Network OS capabilities

Simplified network management	<p>Brocade Virtual Cluster Switching (VCS) fabrics are self-forming and self-healing, providing an operationally scalable foundation for very large or dynamic cloud deployments. Multi-node fabrics can be managed as a single logical element, and fabrics can be deployed and easily re-deployed in a variety of configurations optimized to the needs of particular workloads.</p> <p>For more information on Brocade VCS Fabric technology, refer to Introduction to Brocade VCS Fabric technology on page 18 for an overview and Configuring Brocade VCS Fabrics on page 101 for configuration details.</p>
High resiliency	<p>Brocade VCS fabrics use hardware-based Inter-Switch Link (ISL) Trunking to provide automatic link failover without traffic interruption.</p>
Improved network utilization	<p>Transparent Interconnection of Lots of Links (TRILL)-based Layer 2 routing service provides equal-cost multipaths in the network, resulting in improved network utilization. Brocade VCS Fabric technology also delivers multiple active, fully load-balanced Layer 3 gateways to remove constraints on Layer 2 domain growth, eliminate traffic tromboning, and enable inter-VLAN routing within the fabric.</p> <p>Virtual Router Redundancy Protocol (VRRP) eliminates a single point of failure in a static, default-route environment by dynamically assigning virtual IP routers to participating hosts. The interfaces of all routers in a virtual router must belong to the same IP subnet. There is no restriction against reusing a virtual router ID (VRID) with a different address mapping on different LANs.</p> <p>Refer to Fabric overview on page 101 for additional information about TRILL.</p> <p>Refer to the "VRRP overview" section of the <i>Network OS Layer 3 Routing Configuration Guide</i> for additional information on VRRP/VRRP-E.</p>
Server virtualization	<p>Automatic Migration of Port Profile (AMPP) functionality provides fabric-wide configuration of network policies, achieves per-port profile forwarding, and enables network-level features to support Virtual Machine (VM) mobility.</p> <p>Refer to the "Configuring AMPP section" of the <i>Network OS Layer 2 Switching Configuration Guide</i> for more information about AMPP.</p>
Network convergence	<p>Data Center Bridging (DCB)-based lossless Ethernet service provides isolation between IP and storage traffic over a unified network infrastructure. Multi-hop Fibre Channel over Ethernet (FCoE) allows an FCoE initiator to communicate with an FCoE target that is a number of hops away.</p> <p>Refer to the "End-to-end FCoE" section of the <i>Network OS Layer 2 Switching Configuration Guide</i> for more information about multi-hop FCoE.</p>

In Network OS, all features are configured through a single, industry-standard command line interface (CLI). Refer to the *Network OS Command Reference* for an alphabetical listing and detailed description of all the Network OS commands.

Brocade VCS Fabric terminology

The following terms are used in this document.

TABLE 2 Network OS terms

Term	Definition
Edge ports	In an Ethernet fabric, all switch ports used to connect external equipment, including end stations, switches, and routers.
Ethernet fabric	A topologically flat network of Ethernet switches with shared intelligence, such as the Brocade VCS Fabric.
Fabric ports	The ports on either end of an Inter-Switch Link (ISL) in an Ethernet fabric.
Inter-Switch Link (ISL)	An interface connected between switches in a VCS fabric. The ports on either end of the interface are called ISL ports or Fabric ports. The ISL can be a single link or a bundle of links forming a Brocade trunk. This trunk can either be created as a proprietary Brocade trunk, or a standard IEEE 802.3ad based link aggregation.
RBridge	A physical switch in a VCS fabric.
RBridge ID	A unique identifier for an RBridge, each switch has a unique RBridge ID. In commands, the RBridge ID is used in referencing all interfaces in the VCS fabric. Refer to Configuring a Brocade VCS Fabric on page 104 for information about setting the RBridge ID.
VCS ID	A unique identifier for a VCS fabric. The factory default VCS ID is 1. All switches in a VCS fabric must have the same VCS ID.
WWN	World Wide Name. A globally unique ID that is burned into the switch at the factory.

Introduction to Brocade VCS Fabric technology

Brocade VCS Fabric technology is an Ethernet technology that allows you to create flatter, virtualized, and converged data center networks. Brocade VCS Fabric technology is elastic, permitting you to start small, typically at the access layer, and expand your network at your own pace.


Brocade VCS Fabric technology is built upon three core design principles:

- Automation
- Resilience
- Evolutionary design

When two or more Brocade VCS Fabric switches are connected together, they form an *Ethernet fabric* and exchange information among each other using *distributed intelligence*. To the rest of the network, the Ethernet fabric appears as a single *logical chassis*.

The following figure shows an example of a data center with a classic hierarchical Ethernet architecture and the same data center with a Brocade VCS Fabric architecture. The Brocade VCS Fabric architecture provides a simpler core-edge topology and is easily scalable as you add more server racks.

FIGURE 1 Comparison of classic Ethernet and Brocade VCS Fabric architectures




Automation

Resilience is a foundational attribute of Brocade Fibre Channel storage networks and resilience is also a requirement in modern data centers with clustered applications and demanding compute Service-Level Agreements (SLAs). In developing its VCS Fabric technology, Brocade naturally carried over this core characteristic to its Ethernet fabric design.

In traditional Ethernet networks running Spanning Tree Protocol (STP), only 50 percent of the links are active; the rest (shown as dotted lines in the following figure) act as backups in case the primary connection fails.

When you connect two or more Brocade VCS Fabric mode-enabled switches they form an Ethernet fabric (provided the two switches have a unique RBridge ID and same VCS ID), as shown in the following figure.

FIGURE 2 Ethernet fabric with multiple paths




The Ethernet fabric has the following characteristics:

- It is a switched network. The Ethernet fabric utilizes an emerging standard called Transparent Interconnection of Lots of Links (TRILL) as the underlying technology.
- All switches automatically know about each other and all connected physical and logical devices.
- All paths in the fabric are available. Traffic is always distributed across equal-cost paths. As illustrated in the figure, traffic from the source to the destination can travel across two paths.
- Traffic travels across the shortest path.
- If a single link fails, traffic is automatically rerouted to other available paths. In the topology shown in the figure, if one of the links in Active Path #1 goes down, traffic is seamlessly rerouted across Active Path #2.
- STP is not necessary because the Ethernet fabric appears as a single logical switch to connected servers, devices, and the rest of the network.
- Traffic can be switched from one Ethernet fabric path to the other Ethernet fabric path.

Distributed intelligence

With Brocade VCS Fabric technology, all relevant information is automatically distributed to each member switch to provide unified fabric functionality, as illustrated in the following figure.

FIGURE 3 Distributed intelligence in a Brocade VCS fabric



A Brocade VCS Fabric is designed to be managed as a single "logical chassis," so that each new switch inherits the configuration of the fabric, and the new ports become available immediately. The fabric then appears to the rest of the network as a single switch. This significantly reduces complexity for the management layer, which in turn improves reliability and reduces troubleshooting.

In addition, VCS Fabrics provide a NETCONF application programming interfaces (API), as well as extensions to OpenStack Quantum to orchestrate both physical and logical networking resources as part of virtual machine deployment to support multi-tiered application topologies.


Distributed intelligence has the following characteristics:

- The fabric is self-forming. When two Brocade VCS Fabric mode-enabled switches are connected, the fabric is automatically created and the switches discover the common fabric configuration.
- The fabric is masterless. No single switch stores configuration information or controls fabric operations. Any switch can fail or be removed without causing disruptive fabric downtime or delayed traffic.
- The fabric is aware of all members, devices, and virtual machines (VMs). If the VM moves from one Brocade VCS Fabric port to another Brocade VCS Fabric port in the same fabric, the port-profile is automatically moved to the new port, leveraging Brocade's Automatic Migration of Port Profiles (AMPP) feature.

Logical chassis

All switches in an Ethernet fabric are managed as if they were a single logical chassis. To the rest of the network, the fabric looks no different from any other Layer 2 switch. The following figure illustrates an Ethernet fabric with two switches. The rest of the network is aware of only the edge ports in the fabric, and is unaware of the connections within the fabric.

FIGURE 4 Logical chassis in Ethernet fabric



Each physical switch in the fabric is managed as if it were a blade in a chassis. When a Brocade VCS Fabric mode-enabled switch is connected to the fabric, it inherits the configuration of the fabric and the new ports become available immediately.

Ethernet fabric formation

Brocade VCS Fabric protocols are designed to aid the formation of an Ethernet fabric with minimal user configuration. Refer to [Brocade VCS Fabric formation](#) on page 101 for detailed information about the Ethernet fabric formation process.

All supported switches are shipped with Brocade VCS Fabric mode disabled. Refer to [Configuring Brocade VCS Fabrics](#) on page 101 for information about disabling and enabling Brocade VCS Fabric mode on your switches.

Automatic neighbor node discovery

When you connect a switch to a Brocade VCS Fabric mode-enabled switch, the Brocade VCS Fabric mode-enabled switch determines whether the neighbor also has Brocade VCS Fabric mode enabled. If the switch has Brocade VCS Fabric mode enabled and the VCS IDs match, the switch joins the Ethernet fabric.

Refer to [Configuring Brocade VCS Fabrics](#) on page 101 for information about changing the VCS ID.

Automatic ISL formation and hardware-based trunking

When a switch joins an Ethernet fabric, ISLs automatically form between directly connected switches within the fabric.

If more than one ISL exists between two switches, then Brocade ISL trunks can form automatically. All ISLs connected to the same neighboring Brocade switch attempt to form a trunk. The trunks are formed only when the ports belong to the same port group. No user intervention is necessary to form these trunks.

Refer to [Configuring fabric interfaces](#) on page 106 for information about enabling and disabling ISLs and trunks.

Principal RBridge election

The RBridge with the lowest World Wide Name (WWN) in the Ethernet fabric is elected as the principal RBridge.

The role of the principal RBridge is to decide whether a new RBridge joining the fabric conflicts with any of the RBridge IDs already present in the fabric. If a conflict arises, the principal RBridge keeps the joining RBridge segmented.

Refer to [Configuring a Brocade VCS Fabric](#) on page 104 for information about setting the RBridge ID.

Brocade VCS Fabric technology use cases

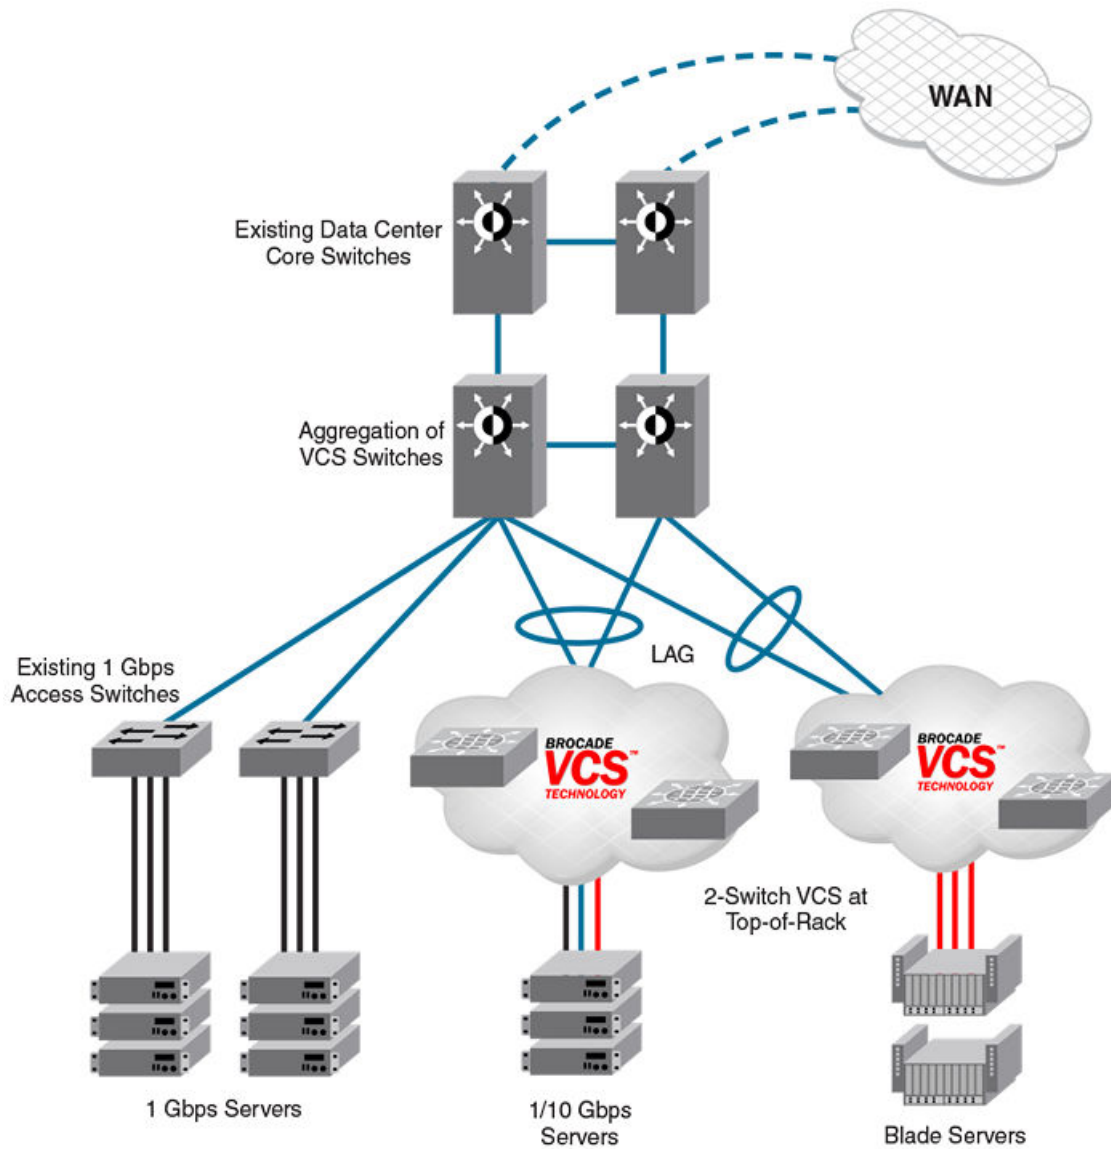
This section describes the following use cases for Brocade VCS Fabric technology:

- Classic Ethernet
- Large-scale server virtualization

Classic Ethernet access and aggregation use case

Brocade VCS Fabric can be deployed in the same fashion as existing top-of-rack switches, as shown in the following figure. In the right-most two server racks, a two-switch Ethernet fabric replaces the Ethernet switch at the top of each rack.

FIGURE 5 Pair of Brocade VDX switches at the top of each server rack



The servers perceive a single top-of-rack switch, allowing for active/active connections running end-to-end.


Brocade VCS Fabric technology in this use case provides the following advantages:

- Multiple active-active connections, with increased effective bandwidth
- Preserves existing architecture
- Works with existing core and aggregation networking products
- Co-exists with existing access switches
- Supports 1- and 10-Gbps server connectivity
- Works with server racks or blade servers

Large-scale server virtualization use case

The following figure shows an example of a logical two-tier architecture with Brocade VCS Fabrics at the edge. Each Brocade VCS Fabric appears as a single virtual switch to the switches outside the fabric, which results in flattening the network.

FIGURE 6 Collapsed, flat Layer 3 networks enabling virtual machine mobility



Brocade VCS Fabric technology in this use case provides the following advantages:


- Optimizes the multipath network (all paths and Layer 3 gateways are active, no single point of failure, and STP is not necessary)
- Increases sphere of virtual machine (VM) mobility

Brocade VCS Fabric connectivity with Fibre Channel SAN

Using the FlexPort feature on the Brocade VDX 6740, Fibre Channel ports provide support for connecting a Brocade VCS Fabric to a Fibre Channel SAN. Fibre Channel routers provide the connectivity, which provides access to Fibre Channel devices while preserving isolation between the fabrics. Brocade zoning allows you to determine which FCoE devices can access which storage devices on the Fibre Channel SAN.

Brocade VDX 6740 switches can be deployed into your Brocade VCS Fabric as access-level switches, aggregation-level switches, or as a means of attachment to Brocade VCS Fabric aggregation-level switches. Brocade recommends deployment as access-level switches to minimize congestion issues for storage traffic and isolating FCoE traffic from non-FCoE traffic. The following figure shows such a deployment.

FIGURE 7 Brocade VDX 6740 switches deployed as access-level switches



Topology and scaling


Up to 24 switches can exist in a Brocade VCS Fabric. Although you can use any network topology to build a Brocade VCS Fabric, the following topics discuss the scaling, performance, and availability considerations of topologies more commonly found in data centers:

- [Core-edge topology](#) on page 27
- [Ring topology](#) on page 27
- [Full mesh topology](#) on page 28

Core-edge topology

Core-edge topology devices connect to edge switches, which are connected to each other through core switches. The example shown in the following figure uses three core switches. You could use more or fewer switches in the core, depending on whether you need higher availability and greater throughput, or a more efficient use of links and ports.

FIGURE 8 Core-edge topology




This topology is reliable, fast, and scales well. It is reliable because it has multiple core switches. If a core switch or a link to a core switch fails, an alternate path is available. As you increase the number of core switches, you also increase the number of link or core switch failures your cluster can tolerate.

High performance and low latency are ensured because throughput is high and the hop count is low. Throughput is high because multiple core switches share the load. Two hops get you from any edge switch to any other edge switch. If you need greater throughput, simply add another core switch.

Scaling the topology also requires additional core switches and links. However, the number of additional links you need is typically not as great as with, for example, a full mesh topology.

Ring topology

Ring topology connects each node to exactly two other nodes, forming a single continuous pathway. Data travels from node to node, with each node along the path handling every packet of the data. The following figure shows a ring topology.


FIGURE 9 Ring topology

This topology is highly scalable, yet susceptible to failures and traffic congestion. It is highly scalable because of its efficient use of interswitch links and ports; an additional node requires only two ports to connect to the ring. It is susceptible to failures because it provides only one path between any two nodes. Throughput of the fabric is limited by the slowest link or node. Latency can be high because of the potentially high number of hops it takes to communicate between two given switches. This topology is useful where economy of port use is critical, but availability and throughput are less critical.

Full mesh topology

Full mesh topology connects each node to all other cluster nodes, as shown in the following figure.

FIGURE 10 Full mesh topology



This topology is highly reliable and fast, but it does not scale well. It is reliable because it provides many paths through the fabric in case of cable or node failure. It is fast with low latency because you can get to any node in the fabric in just one hop. It does not scale well because each additional node increases the number of fabric links and switch ports exponentially. This topology is suitable for smaller fabrics only.

Basic Switch Management

• Switch management overview.....	31
• Ethernet management interfaces.....	36
• Stateless IPv6 autoconfiguration.....	37
• Switch attributes.....	37
• Switch types.....	37
• Operational modes.....	38
• Modular platform basics.....	41
• Supported interface modes.....	43
• Slot numbering and configuration.....	43
• Connecting to a switch.....	44
• Using the management VRF.....	48
• Configuring and managing switches.....	48
• Configuring policy-based resource management.....	71
• Brocade support for OpenStack.....	74
• Mixed-version fabric cluster support.....	75

Switch management overview

In addition to connecting to Brocade switches, an understanding of switch types, attributes, and operational modes is essential to the successful installation and management of networks. This chapter introduces the operational modes, command modes and submodes, and other switch-related activities (such as troubleshooting and managing high-availability scenarios), providing an essential reference for everyday management operations.

Connecting to a switch

You can connect to your switch through a console session on the serial port, or through a Telnet or Secure Shell (SSH) connection to the management port or the inband mgmt-vrf port. You can use any account login present in the local switch database or on a configured authentication, authorization, and accounting (AAA) server for authentication. For initial setup procedures, use the pre-configured administrative account that is part of the default switch configuration.

The switch must be physically connected to the network. If the switch network interface is not configured or the switch has been disconnected from the network, use a console session on the serial port.

- Refer to the *Brocade VDX Hardware Reference* manuals for information on connecting through the serial port.
- Refer to [Configuring Ethernet management interfaces](#) on page 49 for details on configuring the network interface.

Telnet and SSH overview

Telnet and Secure Shell (SSH) are mechanisms for allowing secure access to management functions on a remote networking device. SSH provides a function similar to Telnet, but unlike Telnet, which offers no security, SSH provides a secure, encrypted connection to the device.

SSH and Telnet support is available in privileged EXEC mode on all Brocade VDX platforms. Both IPv4 and IPv6 addresses are supported.

Telnet and SSH services are enabled by default on the switch. When the Telnet server or SSH server is disabled, access to the switch is not allowed for inbound Telnet or SSH connections, thereby restricting remote access to the switch.

Network OS supports up to 32 Telnet or SSH sessions on a switch.

In configuration mode, the CLI can be used to disable Telnet or SSH service on the switch. Doing so will terminate existing inbound Telnet or SSH connections and block any new inbound Telnet or SSH connections to the switch. Additional inbound Telnet or SSH connections will not be allowed until the Telnet server or SSH server is re-enabled. If you have admin privileges, you can re-enable inbound Telnet or SSH connections from configuration mode.

If you are in logical chassis cluster mode (refer to [Operational modes](#) on page 38), the command for enabling or disabling Telnet or SSH services is not distributed across the cluster. The RBridge ID of the node should be used to configure the service on individual nodes.

In operational mode, you can use the **show** command to display whether Telnet or SSH is enabled or disabled on the switch.

SSH server key exchange and authentication

The Secure Sockets Handling (SSH) protocol allows users to authenticate using public and private key pairs instead of passwords. In password-based authentication, the user must enter a password for authentication purposes. In public-key authentication, the user should have a private key in the local machine and a public key in the remote machine. The user should be logged in to the local machine to be authenticated. If a passphrase is provided while generating the public and private key pair, it must be entered to decrypt the private key while getting authenticated.

SSH key-exchange specifies the method used for generating the one-time session keys for encryption and authentication with the SSH server. A user is allowed to configure the SSH server key-exchange method to DH Group 14. When the SSH server key-exchange method is configured to DH Group 14, the SSH connection from a remote SSH client is allowed only if the key-exchange method at the client is also configured to DH Group 14.

The following steps briefly describe public-key authentication:

1. The user generates a pair of encryption keys in a local machine using the **ssh-keygen** command, along with the public and private key, as shown below. Messages encrypted with the private key can only be decrypted by the public key, and vice-versa.

```
switch# ssh-keygen -t rsa
generates RSA public and private keypair
switch# ssh-keygen -t dsa
generates DSA public and private keypair
```

2. The user keeps the private key on the local machine, and uploads the public key to the switch.
3. When attempting to log in to the remote host, the user receives an encrypted message from the remote host containing the public key. After the message is decrypted in the local host by means of the private key, the user is authenticated and granted access.

The **ssh-keygen** command is not distributed across the cluster. The RBridge ID of the node should be used to configure service on individual nodes.

Feature support for Telnet

The following features are not supported with Telnet:

- Displaying Telnet sessions
- Terminating hung Telnet sessions

Feature support for SSH

SSHv2 is the supported version of SSH, but not all features typically available with SSHv2 are supported on the Brocade VDX family of switches.

The following encryption algorithms are supported:

- **3des** Triple-DES (default)
- **aes256-cbc** : AES in CBC mode with 256-bit key
- **aes192-cbc** : AES in CBC mode with 192-bit key
- **aes128-cbc** : AES in CBC mode with 128-bit key

The following Hash-based Message Authentication Code (HMAC) message authentication algorithms are supported:

- **hmac-md5** : MD5 encryption algorithm with 128-bit key (default).
- **hmac-md5-96** : MD5 encryption algorithm with 96-bit key.
- **hmac-sha1** : SHA1 encryption algorithm with 160-bit key.
- **hmac-sha1-96**: SHA1 encryption algorithm with 96-bit key.

SSH user authentication is performed with passwords stored on the device or on an external authentication, authorization, and accounting (AAA) server.

The following features are not supported with SSH:

- Displaying SSH sessions
- Deleting stale SSH keys

Firmware upgrade and downgrade considerations with Telnet or SSH

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

Upgrading to Network OS v4.0 or later is automatically allowed because the Telnet server and SSH server status are enabled by default.

Using DHCP Automatic Deployment

DHCP Automatic Deployment (DAD) is a method used to bring up the switch with new firmware or a preset configuration automatically. In Network OS 4.1.0 and later, you can automatically bring up a switch with new firmware or a preset configuration, omitting the need for logging in to the switch console to configure the switch. If you are in fabric cluster mode, you can apply either the default configuration or a preset configuration. For node replacement in logical chassis cluster mode, the switch is set to the default configuration.

NOTE

The DAD process is disruptive to traffic.

You must be using DHCP to use DAD. You utilize the DHCP process to retrieve certain parameters (for example, the firmware path, VCS ID, VCS mode, RBridge ID, and preset configuration file) needed by the DAD process to perform the firmware and configuration downloads. Currently, only DHCPv4 is supported.

You must enable DAD from the CLI, after which the switch is rebooted automatically. After the DAD process is triggered and completed, DAD is automatically disabled. If you attempt to download new firmware that is already installed on the switch, the DAD process is aborted.

NOTE

If DAD is enabled, you are warned when initiating a firmware download from the CLI that the firmware download will be unsuccessful.

After a firmware download begins, DAD will report firmware download success or failure status.

DAD depends on DHCP automatic firmware download to load the firmware and configuration onto the switch. For this to occur, you must first adhere to the following dependencies:

- The management interface of the switch must be set up as DHCP. After setting up the management interface on a switch in fabric cluster mode, you must use the **copy running-config startup-config** for the configuration to take effect.
- The DHCP server must have the FTP server IP address and configuration file path.
- The configuration file is on the FTP server and it contains the firmware path, new configuration file path, VCS ID, VCS mode, and RBridge ID.

DAD supports the following typical use cases:

- Invoking a firmware upgrade (and optional configuration download) on many switches at the same time in fabric cluster mode.
- Replacing a switch in a cluster by upgrading the firmware and setting up the switch to a preset configuration. (In this instance, DAD must be completed on the new switch hardware (in order to update the firmware) before the new switch can be incorporated into the cluster.)

Note the following considerations when using DAD:

- The DAD process is disruptive.
- Configuration download is not supported during a firmware downgrade.
- Configurations are not downloaded if the DAD process is aborted due to a sanity check failure, or if you are downloading the same firmware version before the firmware download started.
- If an existing firmware download session is either occurring or paused, such as during a firmware commit, DAD is not triggered. Instead, the last firmware download session continues. If a firmware download is in progress and you attempt to enable DAD, you are prompted to try again later.
- In-Service Software Upgrade (ISSU) is not supported at this time.
- For dual management module (MM) chassis, the dual MM must be in sync from the chassis bootup (not from HA failover). In a chassis system, both MMs are rebooted at the same time.

Configuring the DAD process for replacing logical chassis cluster switches

Provides procedures for configuring DHCP Automatic Deployment (DAD) when replacing switches in logical chassis cluster mode.

The following procedure configures DHCP Automatic Deployment (DAD) when replacing switches in logical chassis cluster mode.

For logical chassis cluster switches, the DAD process applies to a new switch and the principal switch.

1. Disconnect the existing switch from the cluster.
2. Connect the new switch to the cluster. The new switch should be the same model and use the same cable connection as the old switch.
The new switch should successfully load Network OS. Note, however, that the new switch cannot join the cluster just yet.
3. From the principal switch, manually run the node replacement with the WWN and RBridge ID.

4. Establish a DAD environment for the new switch. (Make sure DHCP is enabled on the management interface.)
 - a) The management interface of the switch must be set up as DHCP. After setting up the management interface on a switch in fabric cluster mode, you must use the **copy running-config startup-config** command for the configuration to take effect.
 - b) The DHCP server must have the FTP server IP address and configuration file path.
 - c) The configuration file is on FTP server and it contains the firmware path, new configuration file path, VCS ID, VCS mode, and RBridge ID.
 - d) The DHCP server and FTP server must be up-and-running.
 - e) DAD must be enabled on the switch using the CLI.
5. Enable DAD by using the **dhcp auto-deployment enable** command, and enter **yes** when prompted to reboot the system.
6. After the new switch is rebooted, DHCP auto-download process downloads the DAD configuration file to get the VCS mode, VCS ID, and RBridge ID. The RBridge ID should be configured the same as the previous node in the cluster.
7. The DHCP auto-download process sets the VCS ID and RBridge ID for a switch in logical chassis cluster mode. No reboot is triggered.
8. The DHCP auto-download process invokes a firmware download if new firmware is detected. Firmware download completes successfully and the switch comes up with the new firmware and configuration settings.

NOTE

The DAD process will abort if any error is detected.

9. When the new switch comes up, it will join the cluster with the same configuration as the previous switch.
10. Use the **show dadstatus** command to view the current DAD configuration.

Configuring the DAD process for fabric cluster switches

For fabric cluster switches, the DHCP Automatic Deployment (DAD) process is applicable to both cluster upgrades and node replacement.

The following procedure configures DAD on switches in fabric cluster mode.

1. Establish a DAD environment for the new switch. (Make sure DHCP is enabled on the management interface.)
 - a) The management interface of the switch must be set up as DHCP. After setting up the management interface on a switch in fabric cluster mode, you must use the **copy running-config startup-config** command for the configuration to take effect.
 - b) The DHCP server must have the FTP server IP address and configuration file path.
 - c) The configuration file is on FTP server and it contains the firmware path, new configuration file path, VCS ID, VCS mode, and RBridge ID.
 - d) The DHCP server and FTP server must be up and running.
 - e) DAD must be enabled on the switch by means of the CLI.
2. Enable DAD by using the **dhcp auto-deployment enable** command, and enter **yes** when prompted to reboot the system. During system bootup, the DHCP auto-download process downloads the DAD configuration file and obtains the VCS ID and RBridge ID settings for a switch in fabric cluster mode. The switch configuration is retrieved from the FTP server if one is set up.

NOTE

The DAD process will abort if any error is detected.

3. If node configuration needs to be downloaded, set up the new configuration as the startup configuration so it will be applied automatically.
The DHCP auto-download process invokes a firmware download if new firmware is detected. After the firmware download completes successfully, the switch comes up with the new firmware and configuration settings.
4. Use the **show dadstatus** command to view the current DAD configuration.

Telnet and SSH considerations and limitations

- Access to the switch is not allowed for both inbound Telnet and SSH connections from both IPv4 and IPv6 addresses when the Telnet or SSH server are disabled.
- Outgoing Telnet or SSH connections from the switch to any remote device is not affected by disabling or enabling the Telnet or SSH server in the switch.
- No RASLog or auditlog messages are reported when the Telnet or SSH server is disabled or enabled.

Ethernet management interfaces

The Ethernet network interface provides management access, including direct access to the Network OS CLI. You must configure at least one IP address using a serial connection to the CLI before you can manage the system with other management interfaces. You can either configure static IP addresses, or you can use a Dynamic Host Configuration Protocol (DHCP) client to acquire IP addresses automatically. For IPv6 addresses, both static IPv6 and stateless IPv6 autoconfiguration are supported.

ATTENTION

Setting static IPv4 addresses and using DHCP are mutually exclusive. If DHCP is enabled, remove the DHCP client before you configure a static IPv4 address. However, this does not apply to IPv6 addresses.

Brocade VDX Ethernet interfaces

The Brocade VDX Top-of-Rack (ToR) switches have a single configurable Ethernet interface, Eth0, which can be configured as a management interface.

The modular chassis, the Brocade VDX 8770-8 and the Brocade VDX 8770-4, have two redundant management modules, MM1 and MM2. Each management module can communicate with each of the line cards (interface modules) through an Ethernet connection. Each management module has two Ethernet interfaces, Eth0 and Eth2. These interfaces are also known as Out of Band (OoB) management interfaces.

Eth0 is the management interface and can be configured with an IP address. Eth2 provides connectivity to the other management module and the line cards in the chassis. The Eth2 IP addressing scheme uses default IP addresses to communicate between the modules; these addresses are not user-configurable.

To set a virtual IP or IPv6 address for the chassis, use the **chassis virtual-ip** or **chassis virtual-ipv6** command in RBridge ID configuration mode.

Lights-out management

Lights-out management (LOM) is the ability for a system administrator to monitor and manage servers by a LOM remote control program.

A complete LOM system consists of a hardware component called the LOM module and a program that facilitates the continuous monitoring of variables such as microprocessor temperature and utilization. The program also allows for such remote operations as rebooting, shutdown, troubleshooting, alarm setting, fan-speed control, and operating system reinstallation.

The modular chassis, the Brocade VDX 8770-8 and the Brocade VDX 8770-4, have two redundant management modules, MM1 and MM2. Each management module can communicate with each of the line cards (interface modules) through an Ethernet connection. Each management module has two Ethernet interfaces, Eth0 and Eth2. These interfaces are also known as Out of Band (OOB) management interfaces and support LOM programs.

Stateless IPv6 autoconfiguration

IPv6 allows the assignment of multiple IP addresses to each network interface. Each interface is configured with a link local address in almost all cases, but this address is only accessible from other hosts on the same network. To provide for wider accessibility, interfaces are typically configured with at least one additional global scope IPv6 address. IPv6 autoconfiguration allows more IPv6 addresses, the number of which is dependent on the number of routers serving the local network and the number of prefixes they advertise.

When IPv6 autoconfiguration is enabled, the platform will engage in stateless IPv6 autoconfiguration. When IPv6 autoconfiguration is disabled, the platform will relinquish usage of any autoconfigured IPv6 addresses that it may have acquired while IPv6 autoconfiguration was enabled. This same enabled and disabled state also enables or disables the usage of a link local address for each managed entity (though a link local address will continue to be generated for each switch) because those link local addresses are required for router discovery.

The enabled or disabled state of autoconfiguration does not affect any static IPv6 addresses that may have been configured. Stateless IPv6 autoconfiguration and static IPv6 addresses can coexist.

Switch attributes

A switch can be identified by its IP address, World Wide Name (WWN), switch ID or RBridge ID, or by its host name and chassis name. You can customize the host name and chassis name with the **switch-attributes** command.

- A host name can be from 1 through 30 characters long. It must begin with a letter, and can contain letters, numbers, and underscore characters. The default host name is "sw0." The host name is displayed at the system prompt.
- Brocade recommends that you customize the chassis name for each platform. Some system logs identify the switch by its chassis name; if you assign a meaningful chassis name, logs are more useful. A chassis name can be from 1 through 30 characters long, must begin with a letter, and can contain letters, numbers, and underscore characters. The default chassis names are based on the switch models, such as is VDX 8770-4 or VDX 8770-8.

Switch types

The *switchType* attribute is a unique device model identifier that is displayed when you issue the **show chassis** or the **show rbridge-id** commands. When you are gathering information for your switch support provider, you may be asked for the Brocade product name. Use the information in the following table to convert the switchType identifier to a Brocade product name.

TABLE 3 Mapping switchType to Brocade product names

switchType	Brocade product name	Description
112	Management Module	Internal component on the switch
113	Switch Fabric Module	Internal component on the switch
131	VDX 6740	48 10-GbE SFP+ ports and 4 40-GbE QSFP+ ports
137	VDX 6740T	48 10-GbE 10BASE-T ports and 4 40-GbE QSFP+ ports
151	VDX 6740T-1G	Same as VDX 6740T, but ships with ports set to 1 GbE
1000.x	VDX 8770-4	4 I/O slot chassis supporting 48x1 GbE, 48x10 GbE, 48x10G-T, 12x40 GbE, 27x40 GbE, or 6x100 GbE line cards
1001.x	VDX 8770-8	8 I/O slot chassis supporting 48x1 GbE, 48x10 GbE, 48x10G-T, 12x40 GbE, 27x40 GbE, or 6x100 GbE line cards

Operational modes

Network OS supports the following operational modes for Brocade VDX switches.

The three operational modes are:

- *Logical chassis cluster mode* — One of two types of "VCS" modes for a switch. This mode requires Network OS 4.0.0 or later. In this mode, both the data and configuration paths are distributed. The entire cluster is configured from the principal node. Refer to [Logical chassis cluster mode](#) on page 38 for more information.
- *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 [Fabric cluster mode](#) on page 41 for more information.

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

ATTENTION

The generic term *VCS mode* in this manual applies to both fabric cluster mode and logical chassis cluster mode unless otherwise stated.

Logical chassis cluster mode

Logical chassis cluster mode is defined as a fabric in which both the data and configuration paths are distributed. The entire cluster must be globally configured from the principal node. Logical chassis cluster mode requires Network OS 4.0.0 or later.

Logical chassis cluster mode support for platforms

The following platforms support logical chassis cluster mode and is used in any combination:

- Brocade VDX 2740
- Brocade VDX 6740
- Brocade VDX 6740T
- Brocade VDX 6740T-1G
- Brocade VDX 8770-4

- Brocade VDX 8770-8

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 for the Brocade VDX 2740, 6740, 6740T, 6740T-1G, and 8770.
- 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.

Command blocking in logical chassis cluster mode

In logical chassis cluster mode, some commands cannot be run while other commands or events are processing.

If one of the following CLI command types or events is in progress in the cluster, then any one of the CLI command types in the following list will be rejected until the current command or event has finished.

1. **Copy file running-config** commands
2. HA failover commands
3. VCS ID/RbridgeID change commands
4. Cluster mode change from logical chassis cluster to fabric cluster
5. **Copy default-config startup-config** commands
6. Configuration updates by individual commands
7. Cluster formation events, such as initial cluster formation or secondary joining or rejoining of a cluster

These commands and events are considered to be blocked from occurring simultaneously. However, if the principal node changes during one of these operations or HA failover occurs on principal switch, the new principal will not retain the information that the commands or events not in progress are in a blocked state.

In the case of commands being blocked, the following messages are some of the error messages that could result:

- Cluster formation is in progress. Please try again later.
- User Configuration update is in progress. Please try again later.
- Configuration file replay is in progress. Please try again later.

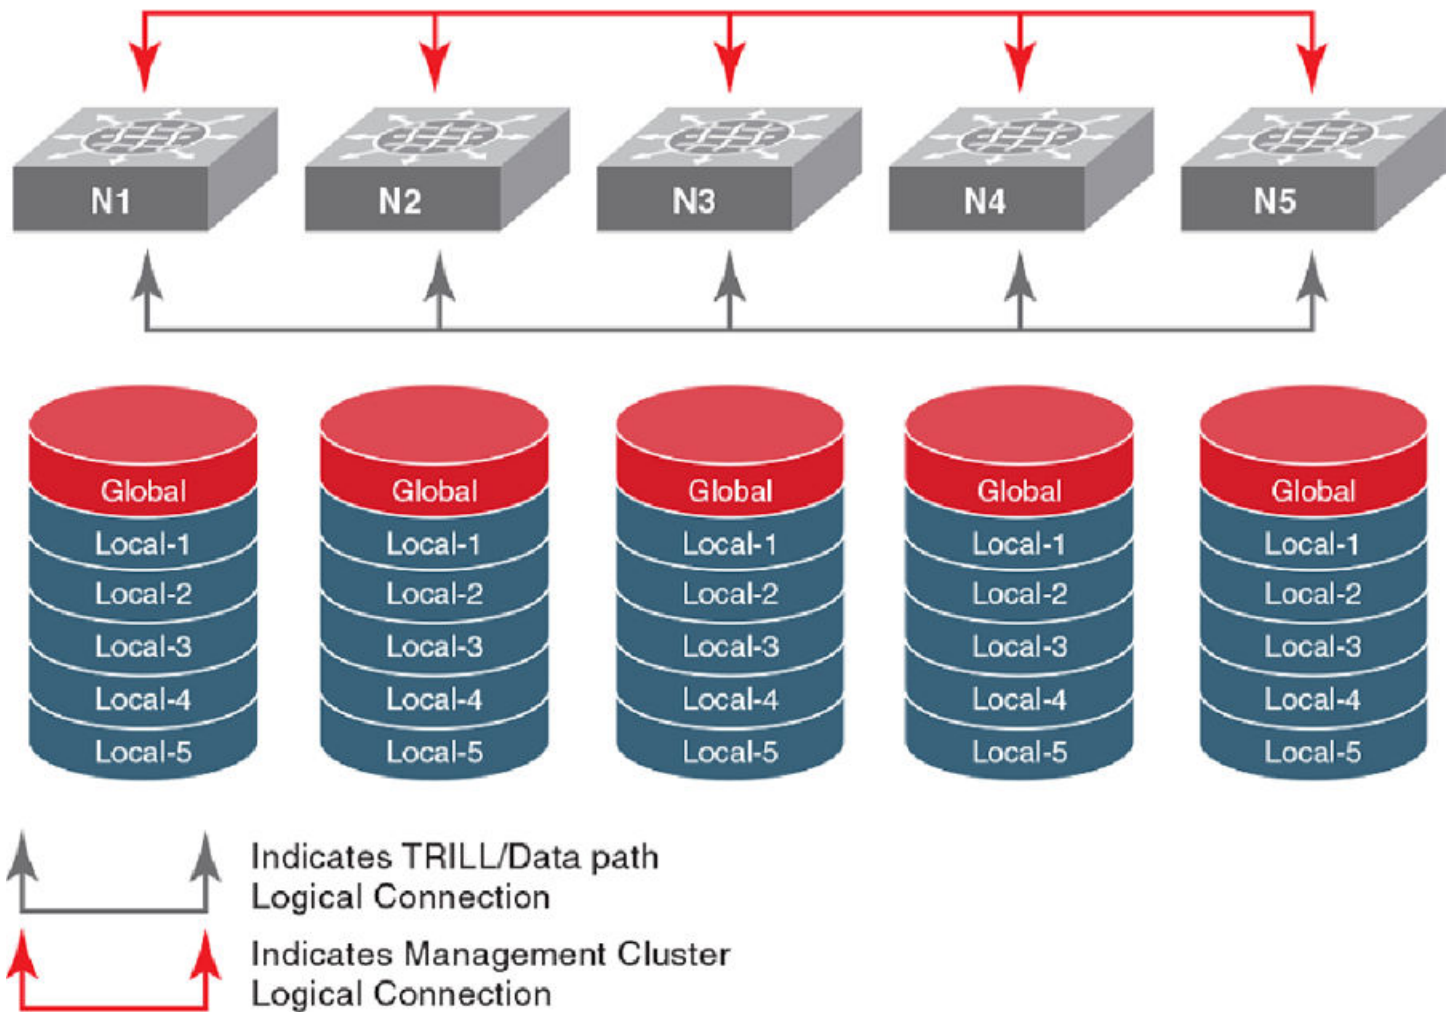
- HA failover is in progress in the cluster. Please try again later.
- VCS Config change is in progress in the cluster. Please try again later.
- Copy default-config startup-config is in progress. Please try again later.

Logical chassis cluster mode configuration

In logical chassis cluster mode, any operation that results in writing to the configuration database gets automatically distributed. There are no exceptions.

Each node in the logical chassis cluster maintains an individual copy of the configuration to enable high availability of the cluster. The following figure illustrates nodes in a logical chassis cluster. Each node has its own databases, and the databases kept by each node are identical at all times.

FIGURE 11 Configuration database in a logical chassis cluster



Network OS switches contain both global and local configuration. In a logical chassis cluster, a single global configuration exists across all cluster members, while each individual member has its own local configuration. (Conversely, in fabric cluster mode, each cluster member can have its own unique global configuration.)

Global configuration is required for cluster-wide operations, whereas local configuration is specific to the operation of an individual node. For more information and examples of each type of configuration, refer to [Examples of global and local configurations](#) on page 64.

Fabric cluster mode

Fabric cluster mode is defined as a fabric in which the data path for nodes is distributed, but the configuration path is not distributed. Each node keeps its configuration database independently.

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

- Brocade VDX 8770-4
- Brocade VDX 8770-8
- Brocade VDX 6740
- Brocade VDX 6740T
- Brocade VDX 6740T-1G

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.

NOTE

In fabric cluster mode, the **all** keyword to the **rbridge-id** command is not available, and a remote RBridge cannot be addressed by means of the **rbridge-id rbridge-id** command.

Modular platform basics

The Brocade VDX 8770 platform features two redundant management modules, three or six switch fabric modules, and four or eight line cards depending on the switch model. The Brocade VDX 8770-4 supports four line cards and the Brocade VDX 8770-8 supports eight line cards.

The following table lists the modules supported on each platform.

TABLE 4 Modules supported on the Brocade VDX 8770 platform

Type	Module ID	Slot numbers	Slot numbers	Description
MM		VDX 8770-4	VDX 8770-8	
	0x70 = 112	M1, M2	M1, M2	Management module (an 8-core 1.5-GHz Control Processor)
SFM	0x71 = 113	S1 - S3	S1 - S6	Switch fabric module (core blade)
LC48X10G	0x72 = 114	L1 - L4	L1 - L8	48-port 10-GbE line card
LC12X40G	0x7F = 127	L1 - L4	L1 - L8	12-port 40-GbE line card
LC48X1G	0x83 = 131	L1 - L4	L1 - L8	48-port 1-GbE line card
LC48X10G-T	0x97 = 151	L1 - L4	L1 - L8	48-port 10 Gbps Base-T line card
LC27X40G	0x96 = 150	L1 - L4	L1 - L8	27-port 40-GbE line card
LC6X100G	0x95 = 149	L1 - L4	L1 - L8	6-port 100-GbE line card

Management modules

Two management modules (MMs) provide redundancy and act as the main controller on the Brocade VDX 8770-4 and VDX 8770-8 chassis. The management modules host the distributed Network OS that provides the overall control plane management for the chassis. You can install a redundant management module in slot M1 or M2 in any of the Brocade VDX 8770 chassis. By default, the system considers the module in slot M1 the active management module and the module in slot M2 the redundant, or standby, management module. If the active module becomes unavailable, the standby module automatically takes over management of the system.

Each management module maintains its own copy of the configuration database. The startup configuration is automatically synchronized with the other management module.

Brocade recommends that each management module (primary and secondary partition) should maintain the same firmware version. For more information on maintaining firmware, refer to the “Installing and Maintaining Firmware” section of the *Network OS Upgrade Guide*.

Each management module has two Ethernet interfaces, Eth0 and Eth2. Eth0 is the management interface and can be configured with an IP address. For more information on configuring the management interface, refer to [Connecting to a switch](#) on page 31.

HA failover

Warm-recovery High Availability (HA) failover is supported for both fabric cluster mode and logical chassis cluster mode.

Warm recovery includes the following behaviors:

- No data path disruption results for Layer 2, Layer 3 and FCoE traffic.
- All Layer 2 and Layer 3 control protocol states are retained.
- The topology state and interface state are retained.
- All running configuration is retained (including the last accepted user configuration just before HA failover).
- During a warm recovery, the principal switch in a logical chassis cluster remains the principal switch. After warm recovery, the principal switch reestablishes cluster management layer connection with other switches and reforms the cluster.
- A secondary switch in a logical chassis cluster reestablishes cluster management layer connection with the principal switch and rejoins the cluster after warm recovery.
- If you run a **reload** command on an active MM, the principal switch in a logical chassis cluster goes into cold recovery and comes back up as a secondary switch.
- HA behavior during In-service software upgrades is the same as for warm-recovery failover.

NOTE

The **ha failover** command is supported only on a dual-management-module chassis system.

Support for in-service software upgrades

In-service software upgrades (ISSUs) are supported in Network OS 4.0.0 and later. Refer to the release notes for upgrade-path information. An ISSU allows a dual management module system or Top of Rack switches to be upgraded non-disruptively and is invoked by entering the **firmware download** command from the active management module.

High Availability behavior during ISSUs is the same as that of warm recovery described in [HA failover](#) on page 42. For more information, refer to the “Upgrading firmware on a modular chassis” of the *Network OS Upgrade Guide*.

Switch fabric modules

The switch fabric modules play a dual role in the fabric connectivity between line cards, providing both the data-plane connectivity and the control-plane connectivity needed for end-to-end credit management in each of the line cards.

In each chassis model, two slots are designated for supporting the control-plane connectivity. In the Brocade VDX 8770-4, the slots S1 and S2 are the designated control-plane slots. In the Brocade VDX 8770-8, the slots S3 and S4 are the designated control-plane slots. At least one of the control-plane slots must be populated to maintain operation. If you remove the switch fabric modules from both the control-plane slots, all line cards will be faulted and the chassis is no longer operational.

Line cards

The following line cards provide I/O ports for network Ethernet protocols:

- LC48x1G - forty eight 1-GbE/10-GbE SFP+ front ports.
- LC48x10G - forty eight 1-GbE/10-GbE SFP+ front ports.
- LC12x40G - twelve 40-GbE QSFP front ports.
- LC48x10G-T - forty eight 10 Gbps Base-T front ports.
- LC27x40G - twenty seven 40-GbE QSFP front ports.
- LC6x100G - six 100-GbE front ports.

Supported interface modes

All interfaces in Brocade VDX chassis come online as Fabric Inter-Switch Links ("Fabric ISLs") by default and will attempt to form a Brocade VCS fabric. If the ISL formation fails, the interfaces come up as "Edge ports".

Slot numbering and configuration

The slot number specifies the physical location of a module in a switch or router, and the number of available slots of each type (interface, management, or switch fabric) depends on the router. Slot configuration is done on a slot-by-slot basis, and the configurations are stored in a persistent database on the switch.

Slot numbering

The slot numbering on the Brocade VDX 8770 chassis is based on the module type. The slot numbers for the line card are numbered L1 through L4 on the Brocade VDX 8770-4, and L1 through L8 on the Brocade VDX 8770-8. The slots for the management modules are numbered M1 and M2. The slots for the switch fabric modules are numbered S1 through S3 on the Brocade VDX 8770-4, and S1 through S6 on the Brocade VDX 8770-8.

Slot configuration

Line cards are registered with the system by type, and the slot must be configured with the correct type before you can install an line card in that slot. When you install a new line card, the system checks whether or not a previous configuration is associated with the slot. The following rules apply when you install or replace an line card:

- When you install an line card and boot it up to an online state in a slot that was never occupied or configured, the module type information is automatically detected and saved to the database. No special configuration is required.
- If you install an line card in a slot that was previously occupied by an line card of the same type and the slot is configured for that same type, you can hot-swap the modules without powering off the line cards. No slot configuration changes are required.

- If the slot was previously configured for a different type of line card, the installation fails and the module is faulted with a "Type mismatch" error. A RASLog error message is generated. You must power off the line card and clear the slot configuration with the **no linecard** command before you can configure the slot for a new line card.

The slot configuration persists in the database even after the line card is physically removed, powered off, or faulted since it first came online. All configuration data associated with the slot is automatically preserved across reboot or hot-swap of the line card with the same type.

Connecting to a switch

You can connect to your switch through a console session on the serial port. Also, you can use SSH or telnet to connect to the management port. You can also use SSH and telnet to connect to an IP inband interface configured on an Ethernet port, a VE interface, or loopback in the Management VRF. You can use any account login present in the local switch database or on a configured authentication, authorization, and accounting (AAA) server for authentication. For initial setup procedures, use the preconfigured administrative account that is part of the default switch configuration.

The switch must be physically connected to the network. If the switch network interface is not configured or the switch has been disconnected from the network, use a console session on the serial port.

- Refer to the Brocade VDX hardware reference manuals for information on connecting through the serial port.
- Refer to [Configuring Ethernet management interfaces](#) on page 49 for details on configuring the management interface.

Establishing a physical connection for a Telnet or SSH session

1. Connect through a serial port to the switch.
2. Verify that the switch's network interface is configured and that it is connected to the IP network through the RJ-45 Ethernet port.
3. Log off the switch's serial port.
4. From a management station, open a Telnet or SSH connection using the management IP address of the switch to which you want to connect.

For more information on setting the management IP address, refer to [Connecting to a switch](#) on page 31.

5. Enter the password.

Brocade recommends that you change the default account password when you log in for the first time. For more information on changing the default password, refer to the *Brocade VDX Hardware Reference* manuals.

6. Verify that the login was successful.

The prompt displays the host name followed by a pound sign (#).

```
switch# login as: admin
admin@10.20.49.112's password:*****
-----
SECURITY WARNING: The default password for at least
one default account (root, admin and user) have not been changed.
Welcome to the Brocade Network Operating System Software
admin connected from 10.110.100.92 using ssh on VDX 6740-48
```

Telnet services

You can use the Telnet service to connect to a switch using either IPv4 or IPv6 protocol.

Establishing a Telnet connection

A Telnet session allows you to access a switch remotely using port 23. However, it is not secure. If you need a secure connection, use SSH.

1. To establish a Telnet session connection, enter **telnet** followed by the switch IP address.

```
switch# telnet 10.17.37.157
```

If the switch is active and the Telnet service is enabled on it, a display similar to the following will appear.

```
Trying 10.17.37.157...
Connected to 10.17.37.157.
Escape character is '^]'.
Network OS (sw0)
switch login:
```

2. Once you have established the Telnet connection, you can log in normally.

NOTE

You can override the default port by using the **telnet ip_address** command with the optional **port** operand (range 0-65535). However, the device must be listening on that port for the connection to succeed.

The following example overrides the default port.

```
switch# telnet 10.17.37.157 87
Trying 10.17.37.157...
Connected to 10.17.37.157.
Escape character is '^]'.
Network OS (sw0)
switch# login:
```

Shutting down the Telnet service

Shutting down the Telnet service will forcibly disconnect all Telnet sessions running on a switch.

You must be in global configuration mode to shut down the Telnet service on a switch.

The Telnet service runs by default.

To shut down the Telnet service on a switch, enter **telnet server shutdown**.

```
switch(config)# telnet server shutdown
switch(config)#
```

All Telnet sessions are immediately terminated, and cannot be re-established until the service is re-enabled.

```
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# telnet server shutdown
```

Re-enabling the Telnet service

Re-enabling the Telnet service permits Telnet access to a switch.

You must be in global configuration mode to shut down the Telnet service on a switch.

To re-enable the Telnet service on a switch enter **no telnet server shutdown**.

```
switch(config)# no telnet server shutdown
```

NOTE

If you are in VCS mode, you must enter RBridge ID configuration mode before issuing the command.

```
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# no telnet server shutdown
```

Connecting with SSH

Connecting to a switch using the SSH (Secure Socket Handling) protocol permits a secure (encrypted) connection.

For a listing and description of all configuration modes discussed here, refer to [Operational modes](#) on page 38.

Establishing an SSH connection

An SSH (Secure Socket Handling) connection allows you to securely access a switch remotely.

You must be in privileged EXEC mode to make an SSH connection to a switch.

1. To establish an SSH connection with default parameters, enter **ssh -l** followed by the *username* you want to use and the *ip_address* of the switch.

```
switch# ssh -l admin 10.20.51.68
```

2. Enter **yes** if prompted.

```
The authenticity of host '10.20.51.68 (10.20.51.68)' can't be established.
RSA key fingerprint is ea:32:38:f7:76:b7:7d:23:dd:a7:25:99:e7:50:87:d0.
Are you sure you want to continue connecting (yes/no)? yes
```

```
Warning: Permanently added '10.20.51.68' (RSA) to the list of known hosts.
admin@10.20.51.68's password: *****
```

```
SECURITY WARNING: The default password for at least
one default account (root, admin and user) have not been changed.
Welcome to the Brocade Network Operating System Software
```

```
admin connected from 10.20.51.66 using ssh on C60_68F
```

NOTE

You can use the **-m** and **-c** options to override the default encryption and hash algorithms.

```
switch# ssh -l admin -m hmac-md5 -c aes128-cbc 10.20.51.68
```

Importing an SSH public key

Importing an SSH public key allows you to establish an authenticated login for a switch.

You must be in privileged EXEC mode to import an SSH public key to a switch.

1. **NOTE**

The following example allows you to import the SSH public key for the user "admin" from a remote host using the credentials shown.

To import an SSH public key, enter **certutil import sshkey**, followed by **user Username host IP_Address directory File_Path file Key_filename login Login_ID**.

```
switch# certutil import sshkey user admin host 10.70.4.106 directory /users/home40/bmeenaks/.ssh
file id_rsa.pub login fvt
```

2. Enter the password for the user.

```
Password: *****
```

```
switch# 2012/11/14-10:28:58, [SEC-3050], 75,, INFO, VDX6740-48, Event: sshutil, Status: success,
Info: Imported SSH public key from 10.70.4.106 for user 'admin'.
```

NOTE

If you are in VCS mode, you must enter RBridge ID configuration mode before issuing the command.

```
switch# certutil import sshkey user admin host 10.70.4.106 directory /users/home40/bmeenaks/.ssh file
id_rsa.pub login fvt rbridge-id 3
```

Deleting an SSH public key

Deleting an SSH public key from a switch prevents it from being used for an authenticated login.

You must be in privileged EXEC mode to delete an SSH public key from a switch.

To delete an SSH public key, enter **no certutil sshkey user Username** followed by either **rbridge-id rbridge-id** or **rbridge-id all**.

```
switch# no certutil sshkey user admin rbridge-id all
```

Specifying a specific RBridge ID removes the key from that RBridge ID; specifying all removes it from all RBridge IDs on the switch.

Shutting down the SSH service

Shutting down the SSH (Secure Socket Handling) service will forcibly disconnect all SSH sessions running on a switch.

You must be in global configuration mode to shut down the SSH service on a switch.

The SSH service runs by default.

To shut down the SSH service on a switch, enter **ssh server shutdown**.

```
switch(config)# ssh server shutdown
switch(config)#
```

All SSH sessions are immediately terminated, and cannot be re-established until the service is re-enabled.

NOTE

If you are in VCS mode, you must enter RBridge ID configuration mode before issuing the command.

```
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# ssh server shutdown
switch(config-rbridge-id-3)#
```

Re-enabling the SSH service

Re-enabling the SSH (Secure Socket Handling) service permits SSH access to a switch.

You must be in global configuration mode to shut down the SSH service on a switch.

To re-enable the SSH service on a switch enter **no ssh server shutdown**.

```
switch(config)# no ssh server shutdown
switch(config)#
```

NOTE

If you are in VCS mode, you must enter RBridge ID configuration mode before issuing the command.

```
switch(config)# rbridge-id 3
switch(config-rbridge-id-3)# no ssh server shutdown
```

Using the management VRF

Virtual Routing and Forwarding (VRF) is a technology that controls information flow within a network, isolating the traffic by partitioning the network into different logical VRF domains. Prior to Network OS release 5.0.0, routers were managed through the "default" VRF; any port that was part of the default VRF could be used for router management.

ATTENTION

Beginning with Network OS release 5.0.0, the default VRF and other user-configured (nondefault) VRFs can no longer be used for router management. Inband management over ports that are part of the default VRF or another user-configured nondefault VRF are no longer supported. Support is now provided for the "management" VRF; this is a dedicated, secure VRF instance that allows users to manage the router inband on switched virtual interfaces (SVIs) and physical interfaces, and that is allowed only on management VRF ports. Services such as Telnet, FTP, SNMP, SSH, SCP, and NetConf are available only through the management VRF. However, Layer 3 routing protocols (such as OSPF, VRRP), including dynamic routing, are not supported. For details, as well as examples of configuring the management VRF and using a variety of **show** commands, refer to the "Understanding and using the management VRF" section in the *Network OS Layer 3 Routing Configuration Guide*.

Configuring and managing switches

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

Configuring Ethernet management interfaces

The Ethernet network interface provides management access, including direct access to the Network OS CLI. You must configure at least one IP address using a serial connection to the CLI before you can manage the system with other management interfaces. You can either configure static IP addresses, or you can use a Dynamic Host Configuration Protocol (DHCP) client to acquire IP addresses automatically. For IPv6 addresses, both static IPv6 and stateless IPv6 autoconfiguration are supported.

ATTENTION

Setting static IPv4 addresses and using DHCP are mutually exclusive. If DHCP is enabled, remove the DHCP client before you configure a static IPv4 address. However, this does not apply to IPv6 addresses.

NOTE

You must connect through the serial port to set the IP address if the network interface is not configured already. Refer to the *Brocade VDX Hardware Reference* manual for your specific product for information on connecting through the serial port.

Configuring static IP addresses

Use static Ethernet network interface addresses in environments where the DHCP service is not available. To configure a static IPv4 or IPv6 address, you must first disable DHCP. Refer to [Configuring IPv4 and IPv6 addresses with DHCP](#) on page 50 for more information.

Configuring a static IPv4 Ethernet address

1. Connect to the switch through the serial console.
2. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
3. Enter the **interface management** *rbridge-id/port* command to configure the management port.

This command enters a management interface configuration mode where you can choose configuration parameters for IPv4 and IPv6 addresses.

- A Top-of-Rack (ToR) switch has a single management port, and the port number for the management port is always 0.
- On a modular switch with two redundant management modules, you can configure two management ports. The port numbers are 1 and 2.

4. Enter the **no ip address dhcp** command to disable DHCP.
5. Enter the **ip address** *IPv4_address/prefix_length* command.
6. Use the **ip route 0.0.0.0/0 gw-ip** command to configure the gateway address.

NOTE

The **ip gateway-address** command is not available on the Brocade VDX series if the L3 or Advanced license is installed. In that case, use the following command sequence:

```
switch(config)# rbridge-id 1
switch(config-rbridge-id-1)# ip route 0.0.0.0/0 default_gateway_address
```

- Verify the configuration with the **do show running-config interface management** command.

NOTE

Specifying an IPv4 address with a subnet mask is not supported. Instead, enter a prefix number in Classless Inter-Domain Routing (CIDR) notation. To enter a prefix number for a network mask, type a forward slash (/) and the number of bits in the mask immediately after the IP address. For example, enter, "209.157.22.99/24" for an IP address that has a network mask with 24 leading 1s in the network mask, representing 255.255.255.0.

```
switch(config-Management-1/0)# do show running-config interface management
interface Management 1/0
no ip address dhcp
ip address 10.24.85.81/20
r-bridge-id1
ip route 0.0.0.0/0 10.24.80.1
no ipv6 address autoconfig
```

- Apart from the two IP addresses on the management modules, modular switches also supports a chassis virtual IP address. Using this virtual IP address, you can login to the switch. The VCS virtual IP address binds to the active MM automatically.

```
switch(config)# rbridge-id 1
switch(config-rbridge-id-1)# chassis virtual-ip 10.24.85.90/20
```

NOTE

In DHCP mode, the chassis IP address is obtained by means of DHCP.

Configuring a static IPv6 Ethernet address

- In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
- Enter the **interface management rbridge-id/port** command.

This command enters a management interface configuration mode where you can choose configuration parameters for IPv4 and IPv6 addresses.

- A Top-of-Rack (ToR) switch has a single management port, and the port number for the management port is always 0.
- On a modular switches with two redundant management modules, you can configure two management ports. The port numbers are 1 and 2.

- Enter the **ipv6 address IPv6_address/prefix_length** command.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface management 1/0
switch(config-Management-1/0)# ipv6 address fd00:60:69bc:832:e61f:13ff:fe67:4b94/64
```

- Apart from the two IP addresses on the management modules, modular switches also support a chassis virtual IP address. Using this virtual IP address, you can log in to the switch. The VCS virtual IP address binds to the active MM automatically.

```
switch(config)# rbridge-id 1
switch(config-rbridge-id-1)# chassis virtual-ipv6 2001:db8:8086:6502/64
```

Configuring IPv4 and IPv6 addresses with DHCP

By default, DHCP is disabled. You must explicitly enable the service. Use the **ip address dhcp** command to enable DHCP for IPv4 addresses, and the **ipv6 address dhcp** command to enable DHCP for IPv6 addresses. The Network OS DHCP clients support the following parameters:

- External Ethernet port IP addresses and prefix length

- Default gateway IP address

NOTE

When you connect the DHCP-enabled switch to the network and power on the switch, the switch automatically obtains the Ethernet IP address, prefix length, and default gateway address from the DHCP server. The DHCP client can only connect to a DHCP server on the same subnet as the switch. Do not enable DHCP if the DHCP server is not on the same subnet as the switch.

The following example enables DHCP for IPv4 addresses.

```
switch(config)# interface management 1/1
switch(config-Management-1/1)# ip address dhcp
```

The following example enables DHCP for IPv6 addresses.

```
switch(config)# interface management 1/1
switch(config-Management-1/1)# ipv6 address dhcp
```

The **show running-config interface management** command indicates whether DHCP is enabled. The following example shows a switch with DHCP enabled for IPv4 addresses.

```
switch# show running-config interface management
interface Management 2/0
ip address dhcp
ip route 0.0.0.0/0 10.24.80.1
ip address 10.24.73.170/20
no ipv6 address autoconfig
```

NOTE

Enabling DHCP removes all configured static IP addresses.

NOTE

Refer to the Network OS Layer 3 Routing Configuration Guide for more information on configuring IP DHCP relay.

Configuring IPv6 autoconfiguration

Refer also to [Stateless IPv6 autoconfiguration](#) on page 37.

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. Take the appropriate action based on whether you want to enable or disable IPv6 autoconfiguration.
 - Enter the **ipv6 address autoconfig** command to enable IPv6 autoconfiguration for all managed entities on the target platform.
 - Enter the **no ipv6 address autoconfig** command to disable IPv6 autoconfiguration for all managed entities on the target platform.

NOTE

On the Brocade VDX 8770, the **autoconfig** command can be issued only on the interface *rbridge-id /1*. However, this operation enables auto-configuration for the entire chassis.

Displaying the network interface

If an IP address has not been assigned to the network interface, you must connect to the Network OS CLI using a console session on the serial port. Otherwise, connect to the switch through Telnet or SSH. Enter the **show interface management** command to display the management interface.

The following example shows the management interface on a Brocade VDX Top-of-Rack (ToR) switch.

```
switch# show interface management
interface Management 9/0
 ip address 10.24.81.65/20
 ip gateway-address 10.24.80.1
 ipv6 ipv6-address [ ]
 ipv6 ipv6-gateways [ fe80::21b:edff:fe0f:bc00 fe80::21b:edff:fe0c:c200 ]
 line-speed actual "1000baseT, Duplex: Full"
 line-speed configured Auto
```

The following example shows the management interfaces on a Brocade VDX 8770-4. IPv6 autoconfiguration is enabled for the entire chassis, and, as a result, a stateless IPv6 address is assigned to both management interfaces.

```
switch# show interface management
interface Management 110/1
 ip address 10.20.238.108/21
 ip gateway-address 10.24.80.1
 ipv6 ipv6-address [ "stateless fd00:60:69bc:85:205:33ff:fe78:7d88/64 preferred" ]
 ipv6 ipv6-gateways [ fe80::21b:edff:fe0b:7800 fe80::21b:edff:fe0b:2400 ]
 line-speed actual "1000baseT, Duplex: Full"
 line-speed configured Auto
interface Management 110/2
 ip address 10.20.238.109/21
 ip gateway-address 10.24.80.1
 ipv6 ipv6-address [ "stateless fd00:60:69bc:85:205:33ff:fe78:be14/64 preferred" ]
 ipv6 ipv6-gateways [ fe80::21b:edff:fe0b:7800 fe80::21b:edff:fe0b:2400 ]
 line-speed actual "1000baseT, Duplex: Full"
 line-speed configured Auto
```

Configuring the management interface speed

By default, the speed of the interface is set to autoconfiguration, which means the interface speed is optimized dynamically depending on load and other factors. You can override the default with a fixed speed value of 10 Mbps full duplex or 100 Mbps full duplex.

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.

```
switch# configure terminal
Entering configuration mode terminal
```

2. Enter the **interface management** command followed by *rbridge-id/O*.

This command places you in the management interface subconfiguration mode.

```
switch(config)# interface management 1/0
switch(config-Management-1/0)#
```

3. Enter the **speed** command with the selected speed parameter. The valid values are **10**, **100**, and **auto**.

```
switch(config-Management-1/0)# speed auto
```

4. Enter the **do show interface management** command followed by *rbridge-id/O* to display the new settings.

```
switch(config-Management-1/0)# do show interface management 1/0
interface Management 1/0
ip address 10.24.81.65/20
ip route 0.0.0.0/0 10.24.80.1
ipv6 ipv6-address [ ]
ipv6 ipv6-gateways [fe80::21b:edff:fe0f:bc00 fe80::21b:edff:fe0c:c200]
line-speed actual "1000baseT, Duplex: Full"
line-speed configured Auto
```

5. Save the configuration changes by using the **copy running-config startup-config** command.

```
switch(config-Management-1/0)# do copy running-config startup-config
```

Configuring a switch banner

A banner is a text message that displays on the switch console. It can contain information about the switch that an administrator may want users to know when accessing the switch.

The banner can be up to 2048 characters long. To create a multi-line banner, enter the **banner login** command followed by the **Esc-m** keys. Enter **Ctrl-D** to terminate the input.

If you are in logical chassis cluster mode, the configuration is applied to all nodes in the cluster.

Complete the following steps to set and display a banner.

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. Enter the **banner login** command and a text message enclosed in double quotation marks (" ").
3. Enter the **do show running-config banner** command to display the configured banner.

```
switch# configure terminal

Entering configuration mode terminal
switch(config)# banner login "Please do not disturb the setup on this switch"

switch(config)# do show running-config banner

banner login "Please do not disturb the setup on this switch"
```

Use the **no banner login** command to remove the banner.

Configuring switch attributes

Refer also to:

- [Switch attributes](#) on page 37.
- [Switch types](#) on page 37.

Setting and displaying the host name

1. In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
2. If Telnet is not activated on the switch, enter the **no telnet server disable** command to activate Telnet.
3. Enter the **switch-attributes** command, followed by a question mark (?) to determine the local RBridge ID.
4. Enter the **switch-attributes** command, followed by the RBridge ID.
5. Enter the **host-name** operand, followed by the host name.

- Save the configuration changes by using the **do copy running-config startup-config** command.

NOTE

This step is used for switches in fabric cluster mode only. If you are using logical chassis cluster mode, startup configurations are not maintained by the cluster; each node preserves its running configuration. For more information about logical chassis cluster mode, refer to [Logical chassis cluster mode](#) on page 38.

- Verify the configuration with the **do show running-config switch-attributes rbridge-id** command.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no telnet server disable
switch(config)# switch-attributes ?
Possible completions: <NUMBER:1-239> Specify the rbridge-id 1
switch(config)# switch-attributes 1
switch(config-switch-attributes-1)# host-name lab1_vdx0023
switch(config-switch-attributes-1)# exit
switch(config)# do copy running-config startup-config
switch(config)# do show running-config switch-attributes 1
switch-attributes 1
  chassis-name VDX 6740-48
  host-name lab1_vdx0023
```

Setting and displaying the chassis name

- In privileged EXEC mode, issue the **configure terminal** command to enter global configuration mode.
- Enter the **switch-attributes** command, followed by a question mark (?) to determine the local RBridge ID.
- Enter the **switch-attributes** command, followed by the RBridge ID.
- Enter the *chassis-name* operand, followed by the chassis name.
- Save the configuration changes using the **do copy running-config startup-config** command.

NOTE

This step is used for switches in fabric cluster mode only. If you are using logical chassis cluster mode, startup configurations are not maintained by the cluster; each node preserves its running configuration. For more information about logical chassis cluster mode, refer to [Logical chassis cluster mode](#) on page 38.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# switch-attributes ?
Possible completions: <NUMBER:1-239> Specify the rbridge-id 1
switch(config)# switch-attributes 1
switch(config-switch-attributes-1)# chassis-name lab1_vdx0023
switch(config)# do copy running-config startup-config
switch(config)# do show running-config switch-attributes 1
switch-attributes 1
  chassis-name lab1_vdx0023
  host-name lab1_vdx0023
```

Viewing switch types

The switchType attribute is a unique device model identifier that allows you to identify the model of a switch from the command line.

In this example, the number 1000 is the value of the switchType attribute. An optional number (.x) indicates the revision of the motherboard.

Refer also to [Switch types](#) on page 37.

Enter **show chassis**.

```
switch# show chassis
Chassis Family: VDX 87xx
Chassis Backplane Revision: 1
switchType: 1000 <== Use table to convert this parameter
(output truncated)
```

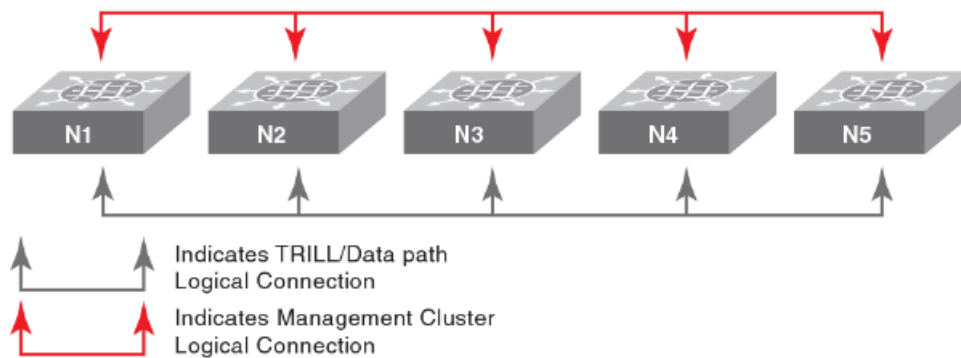
Configuring a switch in logical chassis cluster mode

Refer to [Logical chassis cluster mode](#) on page 38.

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.

FIGURE 12 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          : 44
VCS GUID        : bcab366e-6431-42fe-9af1-c69eb67eaa28
Total Number of Nodes      : 3
-----
```

Rbridge-Id	WWN	Management IP	VCS Status	Fabric Status	HostName
144	10:00:00:27:F8:1E:3C:8C	10.18.245.143	Offline	Unknown	sw0
152	>10:00:00:05:33:E5:D1:93*	10.18.245.152	Online	Online	cz41-h06-
m-r2					
158	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 154 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 page 77.
- 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](#) on page 60.

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 4.0.0

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

Be sure to take the following precautions.

- 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.

- Log in to one switch that you are converting from fabric cluster mode to logical chassis cluster mode.
- 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.

- 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.
- 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, R Bridge ID 1 is the principal.

- 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 R Bridge ID configuration mode for any R Bridge 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](#) on page 60.

Converting a fabric cluster while preserving configuration

There is no specific command that can convert a fabric cluster to a logical chassis cluster while preserving current configurations, but you can accomplish this task as follows:

- Be sure that all nodes are running the same firmware version. Logical chassis cluster functionality is supported in Network OS 4.0 and later.
- Make 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.

- Determine which node contains the global configuration you want to use on the logical chassis cluster, and make a backup of this configuration by running the **copy global-running-config** command and saving the configuration to a file on a remote FTP, SCP, SFTP, or USB location:

NOTE

If you need to combine the global configurations of two or more nodes, manually combine the required files into a single file which will be replayed after the transition to logical chassis cluster mode by using the **copy global-running-config location_config_filename** command. Refer to the section [Taking precautions for mode transitions](#) on page 56

- Back up the local configurations of all individual nodes in the cluster, by running the **copy local-running-config** command on each node and saving the configuration to a file on a remote ftp, scp, sftp, or usb location:

copy local-running-config location_config_filename

- Perform the mode transition from fabric cluster to logical chassis cluster by running the **vcs logical-chassis enable rbridge-id all default-config** command, as shown in [Converting a fabric cluster to a logical chassis cluster](#) on page 57.

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

- 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.
- 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 RBridge ID with the arrow pointing to the WWN is the cluster principal. In this example, RBridge ID 1 is the principal.

- While logged on to the principal node in the logical chassis cluster, copy the saved global configuration file from the remote location to the principal node as follows:

copy location_config_filename running-config

- Verify that the global configuration is available by running the **show global-running-config** command.
- While logged on to the principal node in the logical chassis cluster, copy each saved local configuration file from the remote location to the principal node as follows:

copy location_config_filename running-config

NOTE

You must run this command for each local configuration file you saved (one for each node).

The configuration file is automatically distributed to all nodes in the logical chassis cluster. Each node will contain the same global configuration after the previous steps are performed. Each node will also contain the local configuration information of all the other nodes.

- Verify that the local configurations are available by running the **show local-running-config** command.

- 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](#) on page 60.

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.

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:

- 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.
- Log in to the principal node on the logical chassis cluster.
- 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 logical-chassis enable rbridge-id rbridge-id default-config**.

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

- 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.

Converting to a fabric cluster while preserving configuration

There is no specific command that can convert a logical chassis cluster to a fabric cluster while preserving current configurations, but you can accomplish this task as follows:

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. Back up the configurations of all nodes in the cluster by running the **copy rbridge-running-config rbridge-id** command on each node and saving the configuration to a file on a remote FTP, SCP, SFTP, or USB location:

```
copy rbridge-running-config rbridge-id rbridge-id location_configfilename
```

This command copies both the global and local configurations for the specified RBridge ID.

3. From the principal node of the logical chassis cluster, transition the entire cluster to fabric cluster mode (using the default configuration) by running the following command:

```
no vcs logical-chassis enable rbridge-id all 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.
5. Restore the global and local configurations on each individual node for which you backed up these configurations by running the following command on each node:

```
copy location_configfilename running-config
```

6. To cause this downloaded configuration to be persistent for a node, run the **copy running-config startup-config** command.

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.

The following figure shows the cluster after node N5 has been removed. Nodes N1 through N4 remain in the cluster, and N5 is an island. There is no data path or management path connectivity between the two islands.

FIGURE 13 Removal of Node N5 from the logical chassis 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
```

Merging two logical chassis clusters

You can merge two logical chassis clusters that have the same VCS ID. Follow these steps:

1. Make all required physical connections between the two independent clusters.
2. Decide which cluster should retain the configuration after the merge. Only one configuration can be retained.

3. On the cluster whose configuration will not be retained, issue the **copy default-config startup-config** command so that the nodes in this cluster will reboot with the default configuration.
4. Reboot all nodes in each cluster. The logical chassis cluster whose configuration is being retained recognizes the nodes from the other cluster as new nodes and adds them accordingly.
5. Re-apply the configuration to the cluster whose configuration was not retained.

Changing an RBridge ID on a switch within a fabric

It may become necessary to change the RBridge ID number on a switch that rebooted and has become orphaned from the cluster.

1. Backup the global configuration before changing the RBridge ID, because the local configuration will be reset to default values. Refer to [Backing up configurations](#) on page 84.
2. On the rebooted switch, execute the **chassis disable** command.

```
switch# chassis disable
```

3. From the fabric principal switch, execute the **no vcs enable rbridge-id rbridge-id** command, where *rbridge-id* is the switch that was orphaned.

```
switch# no vcs enable rbridge-id 3
```

4. On the rebooted switch, execute the **vcs rbridge-id rbridge-id** command, where *rbridge-id* is the RBridge you want to use.
5. The VCSID should already be set, if it's not set it with the **vcs rbridge-id rbridge-id**.
6. Reboot the orphaned switch.

The following behavior will take effect after the switch reboots:

- All interfaces will be in *shutdown* state. You must perform a **no shutdown** command on ISL interfaces before the switch will rejoin the cluster.
 - The original configuration will be lost and the switch will have a default configuration when it rejoins the cluster with the new RBridge ID.
7. Use the **show vcs detail** command to verify that the switch is in the fabric.

```
switch# show vcs detail
Config Mode : Local-Only
VCS ID : 1
Total Number of Nodes : 6
Node :1
Serial Number : BKN2501G00R
Condition : Good
Status : Connected to Cluster
VCS Id : 1
Rbridge-Id : 38
Co-ordinator : NO
WWN : 10:00:00:05:33:52:2A:82
Switch MAC : 00:05:33:52:2A:82
FCF MAC : 0B:20:B0:64:10:27
Switch Type : BR-VDX6720-24-C-24
Internal IP : 127.1.0.38
Management IP : 10.17.10.38
Node :2
Serial Number : BZA0330G00P
```

Examples of global and local configurations

The following table provides examples of global and local configuration commands that are available under the respective configuration modes. These settings can be viewed respectively by means of the **show global-running-config** command and the **show local-running-config** command.

TABLE 5 Global and local configuration commands

Global	Local
Interface vlan	switch-attributes
interface port-channel	interface management
port-profile	interface ve
mac access-list	diag post
ip access-list	dpod
sflow	switch-attributes
snmp-server	fabric route mcast
protocol lldp	rbridge-id
zoning	ip route
cee-map	linecard
username	router ospf
	router bgp
	protocol vrrp
	vrrp-group
	interface management
	interface gigabitethernet
	interface tengigabitethernet
	interface fortygigabitethernet
interface fcoe	

Use the **copy snapshot** commands if you need to upload or download configuration snapshot files to and from an ftp or scp server. You may need to use these commands if you took a snapshot of a configuration on a node that was disconnected from the cluster.

Refer to the *Network OS Command Reference* for detailed information about these and other logical chassis server commands.

Configuring a switch in fabric cluster mode

Refer also to [Fabric cluster mode](#) on page 41. When you issue the **show vcs** command to display the VCS configuration for the chassis, the command output shows a single-node VCS with a VCS ID of 1 and an RBridge ID of 1. Use the **vcs** command to change the default values.

```
switch0# show vcs
Config Mode   : Local-Only
VCS ID       : 1
Total Number of Nodes : 1
Rbridge-Id WWN                               Management IP VCS Status Fabric Status HostName
-----
1          >10:00:00:05:33:51:63:42*          10.17.37.154 Online   Online   switch0
                                                2607:f0d0:1002:ff51:ffff:ffff:ffff:fff5
```


Displaying switch interfaces

Interfaces on the VDX 8770 platform are identified by the RBridge ID, slot number, and port number, separated by forward slashes (/). For example, the notation 9/2/8 indicates port 8 located in slot 2 on a chassis with the RBridge ID of 9.

Enter the **show running-config interface** *interface_type* command to display the interfaces and their status.

```
switch# show running-config interface tengigabitethernet
interface tengigabitethernet 1/1/1
fabric isl enable
fabric trunk enable
no shutdown
!
interface tengigabitethernet 1/1/2
fabric isl enable
fabric trunk enable
no shutdown
!
interface tengigabitethernet 1/1/3
fabric isl enable
fabric trunk enable
no shutdown
!
interface tengigabitethernet 1/1/4
fabric isl enable
fabric trunk enable
no shutdown
!
interface tengigabitethernet 1/1/5
fabric isl enable
fabric trunk enable
no shutdown
!
interface tengigabitethernet 1/1/6
fabric isl enable
fabric trunk enable
no shutdown
```

Enter the **show interface** *interface_type rbridge_id/slot/port* command to display the configuration details for the specified interface.

```
switch# show interface tengigabitethernet 1/1/9
tengigabitethernet 1/1/9 is up, line protocol is up (connected)
Hardware is Ethernet, address is 0005.3315.df5a
Current address is 0005.3315.df5a
Pluggable media present
Interface index (ifindex) is 4702109825
MTU 9216 bytes
LineSpeed Actual : 10000 Mbit
LineSpeed Configured : Auto, Duplex: Full
Flowcontrol rx: off, tx: off
Priority Tag disable
Last clearing of show interface counters: 04:12:03
Queueing strategy: fifo
Receive Statistics:
1580 packets, 140248 bytes
Unicasts: 0, Multicasts: 1580, Broadcasts: 0
64-byte pkts: 0, Over 64-byte pkts: 1561, Over 127-byte pkts: 17
Over 255-byte pkts: 2, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
Over 1518-byte pkts(Jumbo): 0
Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
Errors: 0, Discards: 0, TrillportCtrlFrames: 1564
Transmit Statistics:
1583 packets, 140120 bytes
Unicasts: 0, Multicasts: 1583, Broadcasts: 0
Underruns: 0
Errors: 0, Discards: 0, TrillportCtrlFrames: 1583
Rate info:
Input 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Output 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:15:53
```

Refer also to [Slot numbering and configuration](#) on page 43.

Displaying slots and module status information

Use the **show slots** command to display information for all slots in the chassis. The following example shows slot information for the Brocade VDX 8770-8.

```
switch# show slots
Slot  Type      Description                               ID      Status
-----
M1    MM           Management Module                         112     ENABLED
M2    MM           Management Module                         112     ENABLED
S1    SFM          Switch Fabric Module                      113     ENABLED
S2                                VACANT@
S3    SFM          Switch Fabric Module                      113     ENABLED#
S4    SFM          Switch Fabric Module                      113     ENABLED#
S5    SFM          Switch Fabric Module                      113     ENABLED
S6    SFM          Switch Fabric Module                      113     ENABLED
L1                                VACANT
L2                                VACANT
L3    LC48X10G    48-port 10GE card                        114     DIAG RUNNING POST1
L4    LC48X10G    48-port 10GE card                        114     ENABLED
L5                                VACANT
L6                                VACANT
L7    LC48X1G     48-port 1GE card                         114     ENABLED
L7                                VACANT
L8                                VACANT
# = At least one enabled SFM in these slots is required.
@ = The SFM Optical Switch is open.
```

Alternatively, you can use the following commands to display slots per module type:

- Use the **show mm** command to display information for the management modules.
- Use the **show sfm** command to display information for the switch fabric modules.
- Use the **show linecard** command to display information for the line cards.

To make the slot configuration persistent across a chassis reboot (which involves reloading the management modules), you must save the configuration persistently by issuing the **copy running-config startup-config** command after the line card reaches the online state and before the system reboots.

Replacing a line card

You can remove a line card without powering it off. However, doing so will not remove the configuration. When you replace a card with a different type, you must first remove the configuration and then reconfigure the slot for the new line card type.

Install a new line card only if it is supported by the firmware running in the chassis. Inserting a line card into a chassis running firmware that does not support the line card may result in unexpected behavior.

Complete the following steps to replace a line card.



CAUTION

Removing the configuration requires the card to be powered off.

1. Power off the line card by issuing the **power-off linecard** command followed by the slot number.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **rbridge-id rbridge-id** command to enter RBridge ID configuration mode.
4. Enter the **no linecard slot_number** command to clear the slot configuration.

5. Remove the line card.
6. Enter the **linecard slot_number** command followed by a question mark (?) to display the line card menu.
7. Select a line card type and enter the **linecard slot_number linecard_type** command.
8. Enter the **exit** command twice to return to privileged EXEC mode.
9. Insert the new line card into the configured slot.
10. Enter the **power-on linecard** command to power on the line card.
11. Save the configuration persistently by issuing the **copy running-config startup-config** command after the line card reaches the online state.
12. Verify the configuration with the **show running-config linecard linecard** command.

```
switch# power-off linecard 4
switch# configure terminal
Entering configuration mode terminal
switch(config)# rbridge-id 1
switch(config-rbridge-id-1)# no linecard 4
switch(config-rbridge-id-1)# linecard 4 ?
Possible completions:
LC12x40G  12X40G linecard
LC48x1G   48X1G linecard
LC48x10G  48X10G linecard
LC72x1G   72X1G linecard
LC48x10GT 48X10G Base-T linecard
LC27X40G  27X40G linecard
LC6X100G  6X100G linecard
switch(config-rbridge-id-1)# linecard 4 LC48x10G
Creating new linecard configuration was successful.
switch(config-rbridge-id-1)# exit
switch(config)# exit
switch# copy running-config startup-config
switch# show running-config rbridge-id 4 linecard
rbridge-id 1
linecard 1 LC48x10G
linecard 4 LC48x10G
```

Configuring High Availability

The following sections provide you with information on configuring High Availability (HA) support on Brocade switches.

Using HA commands

A variety of High Availability (HA) commands are available on the switch in privileged EXEC mode.

- **show ha** displays the management module status.

```
switch# show ha
Local (M2): Active, Cold Recovered
Remote (M1): Standby, Healthy
HA enabled, Heartbeat Up, HA State synchronized
```

- **ha failover** forces the active management module to fail over. The standby management module will take over as the active management module. This command is only available in a modular chassis system.
- **reload system** reboots the entire chassis. This command is supported only on the active management module. This command is not supported on the standby management module. Both management modules must be in sync for the HA reboot operation to succeed. In logical chassis cluster mode, this command can be issued from the principal node to reset one remote node or all of the remote nodes by specifying either the individual *rbridge-id* or **all**.
- **ha sync start** enables HA state synchronization after an **ha sync stop** command has been invoked.

- **show ha all-partitions** displays details for all line cards and the MM HA state.

NOTE

For additional HA commands and related commands, refer to the *Network OS Command Reference*.

Understanding expected behaviors for reload and failover

The following tables identify expected behaviors that result from controlled and uncontrolled reload and failover conditions.

NOTE

When MMs are out of sync, the **reload** command does not work. Use the **reload system** command to reboot the switch in this case.

TABLE 6 Expected behaviors for controlled reload and failover

Command syntax	Behavior in fabric cluster and logical chassis cluster	Behavior in compact switches
reload	Cold failover to standby management module (MM).	Reloads the switch
reload standby	Reboot the standby MM.	Not available
reload system	Reboot both MMs. MMs will retain the HA roles.	Reloads the switch
ha failover	Warm failover to standby MM.	Not available

TABLE 7 Expected behaviors for uncontrolled failover

Command syntax	Behavior in fabric cluster and logical chassis cluster	Behavior in compact switches
Panic	Warm failover to standby MM.	Reloads the switch
MM removal	Warm failover to standby MM.	Not available
Power cycle	MMs will retain the HA roles upon booting up.	Resets the switch

Disabling and enabling a chassis

The chassis is enabled after power is turned on, and diagnostics and switch initialization routines have finished. All interfaces are online. You can disable and re-enable the chassis as necessary.

- Use the **chassis disable** command if you want to take all interfaces offline. If the switch was part of an Ethernet fabric, the fabric reconfigures.
- Use the **chassis enable** command to bring the interfaces back online. All interfaces that were enabled before the chassis was disabled are expected to come back online. If the switch was part of an Ethernet fabric, it rejoins the fabric.

NOTE

Disabling the chassis is a disruptive operation. Use the **shutdown** command to disable or enable a few selected interfaces only. Refer to the *Network OS Command Reference* for more information on this command.

Rebooting a switch

Network OS provides several commands to reboot your system: **reload**, **fastboot**, and **reload system**.



CAUTION

All reboot operations are disruptive, and the commands prompt for confirmation before executing. When you reboot a switch connected to a fabric, all traffic to and from that switch stops. All ports on that switch remain inactive until the switch comes back online.

Rebooting a Top-of-Rack switch

- The **reload** command performs a "cold reboot" (power off and restart) of the control processor (CP). If the power-on self-test (POST) is enabled, POST is executed when the system comes back up.
- The **fastboot** command performs a "cold reboot" (power off and restart) of the control processor (CP), bypassing POST when the system comes back up. Bypassing POST can reduce boot time significantly.



CAUTION

Do not perform a reload command between a chassis disable command and a chassis enable command. Your ports will be closed.

Rebooting a modular chassis

A chassis reboot brings up the system in sequential phases. First, software services are launched on the management modules and brought up to the active state. Then, the line cards are powered on and initialized. Software services are launched on the line cards and brought up to the active state. When the line card initialization reaches the final state, the chassis is ready to accept user commands from the CLI interface.

During the boot process system initialization, configuration data (default or user-defined) are applied to the switch through configuration replay. For more information, refer to [Managing configurations across redundant management modules](#) on page 86.

- On a modular chassis, the **reboot** and the **fastboot** commands only reboot the management module on which the command is executed. If you log in to the switch IP address and execute one of these commands, only the active management module reboots and POST is bypassed.
- The **reload system** command performs a "cold reboot" (power off and restart) of the entire chassis. If the power-on self-test (POST) is enabled, POST is executed when the system comes back up.

Troubleshooting switches

This section presents an overview of a variety of techniques for capturing data and system messages, which can be helpful in interactions with technical support.

Capturing and managing supportSave data

If you are troubleshooting a production system, you will have to capture data for further analysis or send the data to your switch service provider. The **copy support** command provides a mechanism for capturing critical system data and uploading the data to an external host or saving the data to an attached USB device.

Uploading supportSave data to an external host

To upload supportSave data interactively, enter the **copy support-interactive** command and provide input as prompted. Specifying an IPv6 address for the server requires Network OS v3.0.0 or later. For a non-interactive version of the command, refer to the *Network OS Command Reference*.

```
switch# copy support-interactive

Server Name or IP Address: 10.38.33.131
Protocol (ftp, scp): ftp
User: admin
Password: *****
Directory: /home/admin/support
VCS support [y/n]? (y): n
Module timeout multiplier [Range: 1 to 5. Default: 1]: 1

copy support start
Saving support information for chassis:sw0, module:RAS...(output truncated)
```

Saving supportSave data to an attached USB device

You can use a Brocade-branded USB device to save the support data. The Brocade-branded USB device comes with factory-configured default directories and interacts with the Network OS CLI.

1. Enter the **usb on** command to enable the USB device.
2. Enter the **usb dir** command to display the default directories.
3. Enter the **copy support usb directory** command.

```
switch# usb on
USB storage enabled
switch# usb dir

firmwarekey\ 0B 2010 Aug 15 15:13
support\ 106MB 2010 Aug 24 05:36
support1034\ 105MB 2010 Aug 23 06:11
config\ 0B 2010 Aug 15 15:13
firmware\ 380MB 2010 Aug 15 15:13
Available space on usbstorage 74%
switch# copy support usb directory support
```

If you are in logical chassis cluster mode, you can use the **rbridge-id all** option to invoke supportSave on all nodes at the same time. The **copy support rbridge-id all** command is a blocking command. The Telnet session from which the command is issued will be blocked until supportSave is completed on all nodes in the cluster; however, users can again Telnet into the same node or any other nodes in the cluster. When the command is in progress, output messages from all nodes are shown that include the respective node RBridge IDs. The **copy support** command, when executed with USB as the protocol option, will collect support files to the USB device that is connected to the respective nodes. All USB devices connected to each of the nodes should be enabled before the **copy support usb** command is executed.

The following example shows the **copy support** command with the **rbridge-id all** option.

```
switch# copy support ftp host 10.1.2.30 user fvt password pray4green directory /support rbridge-id all
switch 100: copy support start
switch 117: Saving support information for chassis:sw0, module:RAS...
switch 100: Saving support information for chassis:sw, module:RAS...
switch 117: Saving support information for chassis:sw0, module:CTRACE_OLD...
.....
switch 100: copy support completed
switch 117: copy support completed
2011/04/07-18:03:07, [SS-1000], 2752,, INFO, VDX6720-24, copy support has uploaded support information to the
host with IP address 10.70.4.101.
```

Displaying the status of a supportSave operation

Enter the **show copy-support status** command.

```
switch# show copy-support status
Slot Name      SS type      Completion Percentage
# # # # # # # # # # # # # # # # # # # # # # # # # # # # # #
M1             NORMAL      [100%]
L1/0           NORMAL      [100%]
L1/1           NORMAL      [100%]
L2/0           NORMAL      [100%]
L2/1           NORMAL      [100%]
L4/0           NORMAL      [100%]
L4/1           NORMAL      [100%]
```

Configuring automatic uploading of supportSave data

You can configure a switch to upload first-fault data capture (FFDC) and trace data files automatically to a remote server that is specifically set up for collecting information that results from the **supportSave** command. To enable this feature, you must configure a dedicated server, then invoke the **autoupload-param** command to set the parameters, followed by the **support autoupload enable** command to enable the configurations.

```
switch(config)# support autoupload-param hostip 10.31.2.27 username supportadmin directory /users/support/
ffdc_autoupload protocol ftp password (<string>): *****
```

Displaying the autoupload configuration

Enter the **show running-config support autoupload-param** command to display the autoupload configuration on the local switch.

```
switch(config)# do show running-config support autoupload-param

support autoupload-param hostip 10.31.2.27 username supportadmin directory /users/support/ffdc_autoupload
protocol ftp password "3iTYxJWEUHp9axZQt2tbvw==\n"
```

Using additional supportSave commands

Use the following commands to configure additional supportSave data collection parameters:

- Use the **show support** command to display a list of core files on the switch.
- Use the **clear support** command to erase support data on the switch.

Refer to the *Network OS Command Reference* for more information on these commands.

Logging error messages

Network OS provides several mechanisms for logging error messages including syslog, RASLog, and audit log. The types of message logging available and the setup procedures are documented in the "Introduction to Brocade Error Message Logging" chapter of the *Network OS Message Reference Manual*.

Configuring policy-based resource management

The policy-based resource management feature allows users to make better use of hardware resources. In particular, pre-made profiles are provided that optimize ASIC resources for route profiles and ternary content-addressable memory (TCAM) profiles. The profiles are enabled by keywords available under the **hardware-profile** command in RBridge ID configuration mode. The profile configuration is local to an RBridge within a VCS Fabric.

NOTE

In order for the last update of the profile configuration to take effect on a switch, the switch has to be rebooted. In Logical Chassis mode, use the **reload system** command. In Fabric Cluster mode, run the **copy running-config startup-config** command, followed by the **reload system** command.

The following table describes the available command options (keywords) to optimize route profiles, available under the **route-table** keyword. Refer also to the **hardware-profile** command in the *Network OS Command Reference*.

TABLE 8 Options for optimizing route profiles

Keyword	Optimizes resources for . . .
default	IPv4/IPv6 dual-stack operations
ipv4-max-route	Maximum number of IPv4 routes
ipv4-max-arp	Maximum number of IPv4 ARP entries
ipv4-min-v6	IPv4 routes in dual-stack configurations
ipv6-max-route	Maximum number of IPv6 routes
ipv6-max-nd	Maximum number of IPv6 Neighbor Discovery entries

The following table describes the available command options (keywords) to optimize TCAM profiles, available under the **tcam** keyword.

TABLE 9 Options for optimizing TCAM profiles

Keyword	Optimizes resources for . . .
default	Basic support for all applications
l2-ipv4-acl	Layer 2 and IPv4 ACLs
ipv4-v6-pbr	IPv4 and IPv6 ACLs and policy-based routing tables
ipv4-v6-qos	IPv4 and IPv6 ACLs and QoS
ipv4-v6-mcast	Multicast
l2-acl-qos	Layer 2 ACLs and QoS

Note the following conditions for TCAM profiles:

- TCAM profiles affect only ACLs, policy-based routing (PBR), flow-based QoS, and multicast entries, without affecting other features, protocols, or hardware resources.
- The TCAM profile options (listed in the table) are not customizable or configurable, and they may not be appropriate to all network designs.
- The following QoS features are optimized by TCAM profiles:
 - Flow-based QoS and flow-based policing for Layer2/Layer 3 ingress and egress
 - System Qos (VLAN-based) for Layer2/Layer 3 ingress and egress
 - Auto NAS
 - Storm control
 - Flow-based SPAN and RSPAN, including VXLAN based
 - Flow-based Sflow, including VXLAN based
- The following QoS features are not affected by TCAM profiles:
 - All port-based QoS features (RED; PFC and legacy flow control; CoS mutation, DSCP CoS, DSCP traffic class, DSCP mutation; scheduling, shaping, and port-based policing)
 - Port-based SPAN and RSPAN
 - Port-based Sflow

Configuring hardware profiles

The following examples illustrate the application of the **hardware-profile** command, which is executed in the RBridge ID configuration mode. The options for the **route-table** and **tcam** keywords are as listed in the tables in the previous section.

NOTE

To apply the most recent profile configuration update, you must reboot (reload) the switch. Before reloading, the current profile is in effect and functioning.

The following example selects a route table profile to optimize resources for the maximum number of IPv6 Neighbor Discovery entries:

```
switch(config)# rbridge-id 10
switch(config-rbridge-id-10)# hardware-profile route-table ipv6-max-nd
%Warning: To activate the new profile config, please run 'reload system' on the target switch.
```

NOTE

When you use the **hardware-profile route-table ?** command to see the available options, the currently applied hardware profile is at the top of the list and is enclosed by square brackets ([]).

The following example selects a TCAM profile to optimize resources for the maximum number of IPv4/IPv6 multicast entries:

```
switch(config)# rbridge-id 10
switch(config-rbridge-id-10)# hardware-profile tcam ipv4-ipv6-mcast
%Warning: To activate the new profile config, please run 'reload system' on the target switch.
```

NOTE

When you use the **hardware-profile tcam ?** command to see the available options, the currently applied TCAM profile is at the top of the list and is enclosed by square brackets ([]).

Guidelines for changing hardware profiles

Note the following guidelines for changing hardware profiles:

- In fabric cluster and logical chassis cluster mode, you must reload the target switch after changing a profile for the new profile to take effect.
 - In logical chassis mode, when a secondary switch rejoins the cluster with a default profile configuration while the profile configuration for the secondary switch is "nondefault" on the principle switch, you must reload the secondary switch again after it has rejoined the cluster for the nondefault profile to take effect. The TCAM and LPM hardware profiles are persistent across "copy default-config startup-config" operation.
 - In fabric cluster mode, you must use the **copy running-config startup-config** command first, before reloading the switch. If the VCS ID changes to move one RBridge ID out of the cluster, the system reboots with the default running configuration and boots with the "default" TCAM and routing hardware profile. The previously applied TCAM-based hardware profile is not persistent.
- After Netinstall or a firmware upgrade from to *Network OS v5.0.0.*, the default profiles are automatically set for both TCAM and route tables in the running configuration. Also, the profile configuration defaults after changing a switch VCS ID or Rbridge ID. However, the current profile persists after the **copy default-config startup-config** command completes and the switch reboots. Additionally, the current profile persists after a VCS mode change.
- There is no "no" option for the hardware-profile command, because hardware profiles always exist, with either the default or one of the non-default configurations.
- When you change a hardware profile, the supported scale numbers remain the same with respect to the configuration even if hardware may not be able to fulfill them. This ensures that the same protocol and interface information remain valid with all hardware profile settings.

Using hardware profile show commands

The following **show** commands can be used to verify the status of hardware profiles. For details, refer to the *Network OS Command Reference*.

TABLE 10 Network OS show commands

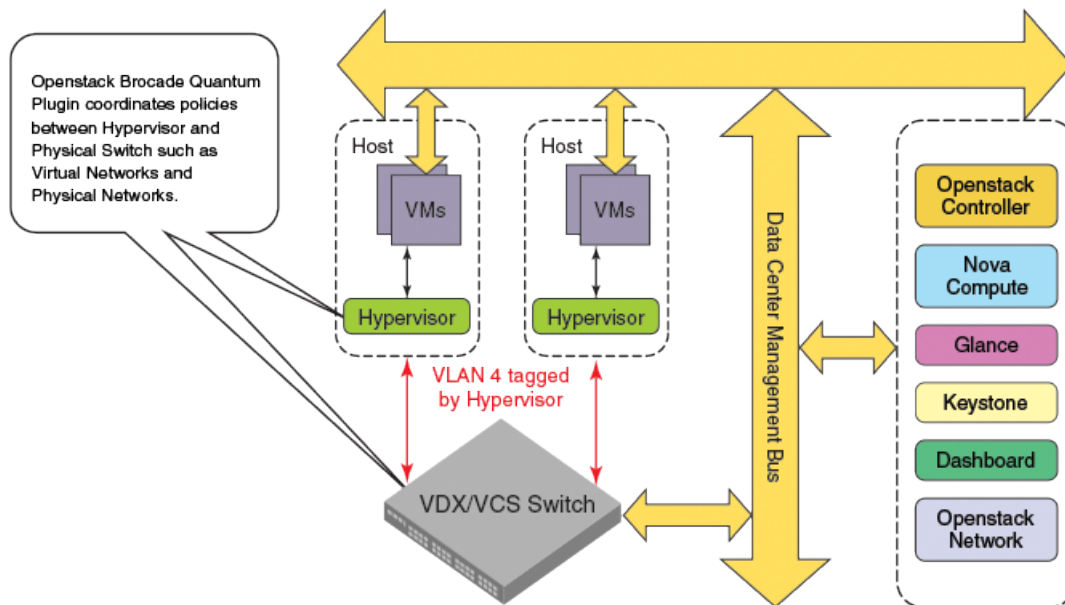
Command	Description
show hardware-profile	Displays the current active profile information and subtype details for each profile type and RBridge ID on local switch or specified RBridge ID or all switches in LC cluster. For complete details on the show hardware-profile command, refer to the <i>Network OS Command Reference</i> .
show running-config rbridge-id hardware-profile	Displays the enabled route table and TCAM profiles in the running configuration for all RBridge IDs, or a specific enabled RBridge ID.

Brocade support for OpenStack

OpenStack is an open source Infrastructure as a Service (IaaS) initiative for creating and managing large groups of virtual private servers in a cloud computing environment. The Brocade Neutron Plugin for VDX/VCS provides a means to interface with OpenStack's networking to orchestrate the Brocade physical switches.


In cloud environments where Virtual Machines (VMs) are hosted by physical servers, the VMs see a new virtual access layer provided by the host machine. This new access layer can be created using many mechanisms, such as Linux Bridges or Virtual Switches. The policies of the virtual access layer (virtual network), must be coordinated with the policies set in the hardware switches. The Brocade Neutron Plugin helps in coordinating this behavior automatically without any intervention from the administrator.

FIGURE 14 Virtual and physical network orchestration



The Brocade Neutron ML2 Plugin communicates with the Brocade Mechanism Driver, which uses NETCONF on the back-end to configure the Brocade switches. The OpenStack feature supports VLANs only.

FIGURE 15 OpenStack configuration path



For the latest version of the OpenStack Neutron plugin, go to the <https://wiki.openstack.org/wiki/Neutron> site and search under plugins for the Brocade plugin. The OpenStack community can be contacted at <http://www.openstack.org/>.

Configuring OpenStack to access Network OS

You must configure the Neutron Plugin and Brocade configuration to activate OpenStack access.

The Brocade ML2 drivers have been certified on Redhat and Ubuntu.

1. Modify the physical switch configuration parameters and the Brocade-specific database configuration in the `brocade.ini` configuration file.

```
% cat /etc/neutron/plugins/brocade/brocade.ini

[SWITCH]
username = admin
password = password
address = <switch mgmt ip address>
ostype = NOS

[DATABASE]
sql_connection = mysql://root:pass@localhost/brcd_Neutron?charset=utf8
```

2. Modify the `m12_config.ini` file to enable the Brocade ML2 driver.

```
% cat /etc/neutron/plugins/m12/m12_conf.ini

[m12]
tenant_network_types = vlan
type_drivers = local,flat,vlan,gre,vxlan
mechanism_drivers = openvswitch,brocade
```

Mixed-version fabric cluster support

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

The following limitations and considerations apply to mixed node support:

- Mixed-versions are supported only in fabric cluster mode and are not supported in logical chassis cluster mode.

- For a fabric cluster running Network OS v4.1.3 and Network OS v5.0.0, the cluster supports the Network OS v4.1.3 feature set. For a list of features that the nodes running Network OS v5.0.0 support, refer to the Network OS v5.0.0 release notes.

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

- Upgrade all Brocade VDX 2740, 6740, 6740T, 6740T-1G, and 8770 units to Network OS v5.0.0.
- Update all other Brocade hardware to Network OS v4.1.3.

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

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

TABLE 11 Feature support for mixed-version fabric clusters

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

Using Network Time Protocol

- [Network Time Protocol overview.....](#) 77
- [Configuring NTP.....](#) 77

Network Time Protocol overview

Network Time Protocol (NTP) maintains uniform time across all switches in a network. The NTP commands support the configuration of an external time server to maintain synchronization between all local clocks in a network.

To keep the time in your network current, it is recommended that each switch have its time synchronized with at least one external NTP server. External NTP servers should be synchronized among themselves in order to maintain fabric-wide time synchronization.

All switches in the fabric maintain the current clock server value in nonvolatile memory. By default, this value is the local clock server of the switch.

Date and time settings

Brocade switches maintain the current date and time inside a battery-backed real-time clock (RTC) circuit. Date and time are used for logging events. Switch operation does not depend on the date and time; a switch with incorrect date and time settings can function correctly. However, because the date and time are used for logging, error detection, and troubleshooting, you should set them correctly.

Time zone settings

You can set the time zone by specifying a geographic region and city by name. You can choose one of the following regions: Africa, America, Pacific, Europe, Antarctica, Arctic, Asia, Australia, Atlantic, and Indian.

The time zone setting has the following characteristics:

- The setting automatically adjusts for Daylight Savings Time.
- Changing the time zone on a switch updates the local time zone setup and is reflected in local time calculations.
- By default, all switches are in the Greenwich Mean Time (GMT) time zone (0,0). If all switches in a fabric are in one time zone, it is possible for you to keep the time zone setup at the default setting.
- System services that have already started will reflect the time zone changes only after the next reboot.
- Time zone settings persist across failover for high availability.
- Time zone settings are not affected by NTP server synchronization.

Configuring NTP

The following sections discuss how to correctly configure the Network Time Protocol for Brocade switches.

Configuration considerations for NTP

Network time synchronization is guaranteed only when a common external time server is used by all switches. If you are in VCS mode, when an **ntp server** command is invoked on one switch in a cluster, the configuration is applied to all switches in the cluster.

The **ntp server** command accepts up to five server addresses in IPv4 or IPv6 format. When you configure multiple NTP server addresses, the **ntp server** command sets the first obtainable address as the active NTP server. If there are no reachable time servers, then the local switch time is the default time until a new active time server is configured.

Setting the date and time

The **clock set** command sets the local clock date and time. Valid date and time values must be in the range between January 1, 1970 and January 19, 2038. If a time zone is not configured, the time zone defaults to Greenwich Mean Time (GMT). If an active NTP server is configured for the switch, it overrides the local time settings.

Enter the **clock set CCYY-MM-DDTHH:MM:SS** command.

The variables represent the following values:

- *CCYY* specifies the year; the valid range is 1970 through 2038.
- *MM* specifies the month; the valid range is 01 through 12.
- *DD* specifies the day; the valid range is 01 through 31.
- *HH* specifies the hour; the valid range is 00 through 23.
- *MM* specifies the minutes; the valid range is 00 through 59.
- *SS* specifies the seconds; the valid range is 00 through 59.

If you are in VCS mode, setting the time and date is done using the RBridge ID of the node.

Here is an example of setting and displaying the date and time in VCS mode:

```
switch# clock set 2013-06-06T12:15:00 rbridge-id all
switch# show clock
rbridge-id all: 2013-06-06 12:15:05 Etc/GMT+0
```

Setting the time zone

Use the **clock timezone** command to set the time zone for a switch. You must use the command for all switches for which a time zone must be set. However, you only need to set the time zone once on each switch, because the value is written to nonvolatile memory.

Setting the time and date can be done in Privileged EXEC mode by using the RBridge ID of the node. (Setting the date and time can also be done in RBridge ID configuration mode, but must be done on a per-node basis in this mode.) Refer to the **clock timezone** command in each mode in the *Network OS Command Reference*.

Refer to [Using Network Time Protocol](#) on page 77 for a complete list of configurable regions and cities.

Enter the **clock timezone region/city** command.

```
switch# clock timezone America/Los_Angeles rbridge-id all
```

NOTE

Upgrade considerations: The existing timezone of the system is retained after a firmware upgrade, and it will be updated in configuration settings.

Downgrade Considerations: Existing timezone of system will be retained after firmware downgrade and the respective entry will be removed from configuration settings.

Displaying the current local clock and time zone

The **show clock** command returns the local time, date, and time zone.

NOTE

This command is currently supported on the local switch.

This example shows the local switch clock time:

```
switch# show clock
rbridge-id 1: 2012-05-04 16:01:51 America/Los Angeles
```

This example shows the clock time for all switches in the cluster (logical chassis cluster mode only):

```
switch# show clock rbridge-id all
rbridge-id 1: 2013-06-06 12:15:05 Etc/GMT+0
rbridge-id 5: 2013-06-06 12:15:05 Etc/GMT+0
rbridge-id 10: 2013-06-06 12:15:05 Etc/GMT+0
```

This example shows the clock time for the switch with rbridge-id 16:

```
switch# show clock rbridge-id 16
rbridge-id 16: 2012-05-04 18:18:51 America/Los Angeles
```

Removing the time zone setting

Use the **no clock timezone** command to remove the time zone setting for the local clock. This operation returns the local time zone to the default value (GMT). When using the **no** operand, you do not need to reference a timezone setting.

Enter the **no clock timezone** command.

```
switch# no clock timezone rbridge-id 5
```

NOTE

The **clock timezone** command can be run in privileged EXEC mode, as shown in the previous example, or in RBridge ID configuration mode. Refer to the *Network OS Command Reference* for descriptions of this command in each of these modes.

Synchronizing the local time with an external source

Use the **ntp server** command to synchronize the local switch time with an NTP server. You can configure up to five IP address. At least one IP address in the list must be a reachable, configured NTP server or the request will fail.

The following example synchronizes the time on the local switch with the ntp server at 192.168.10.1.

Enter the **ntp server ip_address** command.

```
switch(config)# ntp server 192.168.10.1
```

Displaying the active NTP server

Use the **show ntp status** command to display the current active NTP server IP address. If an NTP server is not configured or the server is unreachable, the output displays LOCL (for local switch time). Otherwise, the command displays the NTP server IP address. The command displays the local NTP server configuration only.

If the RBridge ID parameter is not provided, status results default to the local switch. If **rbridge-id all** is specified, the command displays the status for all switches in the cluster. If the RBridge ID is specified, the command displays that node's NTP status.

This example shows the local switch NTP status when an NTP server is not configured:

```
switch# show ntp status
rbridge-id 1: active ntp server is LOCL
```

This example shows the configured NTP server:

```
switch# show ntp status
rbridge-id 1: active ntp server is 10.31.2.81
```

This example shows NTP status for all switches in a cluster.

```
switch# show ntp status rbridge-id all
rbridge-id 7: active ntp server is LOCL
```

Removing an NTP server IP address

To remove an NTP server IP address from the list of server IP addresses on a switch, enter **no ntp server** followed by the server IP address.

The following example removes the NTP server at 192.168.10.1 from the local server IP address database.

```
switch(config)# no ntp server 192.168.10.1
switch# show ntp status
rbridge-id 1: active ntp server is LOCL
```

IMPORTANT



At least one IP address in the remaining list must be for a reachable and configured NTP server; if there is not one the remove request will fail.

Configuration Management

- Configuration management overview..... 81
- Displaying configurations..... 82
- Saving configuration changes..... 83
- Backing up configurations..... 84
- Configuration restoration..... 85
- Managing configurations on a modular chassis..... 86
- Managing configurations in Brocade VCS Fabric mode..... 87
- Rejoining an offline node to a logical chassis cluster..... 88
- Managing flash files..... 89

Configuration management overview

Maintaining consistent configuration settings among switches in the same fabric is an important part of switch management and minimizes fabric disruptions. As part of standard maintenance procedures, it is recommended that you back up all important configuration data for every switch on an external host for emergency reference.

Typical configuration management tasks include the following actions:

- Saving the running configuration to the startup configuration file ([Saving configuration changes](#) on page 83).
- Uploading the configuration files to a remote location ([Backing up configurations](#) on page 84).
- Restoring a configuration file from a remote archive ([Configuration restoration](#) on page 85).
- Archiving configuration files for all your switches to a remote location ([Managing configurations in Brocade VCS Fabric mode](#) on page 87).
- Downloading a configuration file from a remote location to multiple switches ([Managing configurations in Brocade VCS Fabric mode](#) on page 87).

Configuration file types

Brocade Network OS supports three types of configuration files. The table below lists the standard configuration files and their functions.

TABLE 12 Standard switch configuration files

Configuration file	Description
Default configuration <ul style="list-style-type: none"> • defaultconfig.novcs • defaultconfig.vcs 	Part of the Network OS firmware package. The default configuration is applied, if no customized configuration is available.
Startup configuration <ul style="list-style-type: none"> • startup-config 	Configuration effective on startup and after reboot.
Running configuration <ul style="list-style-type: none"> • running-config 	<p>Current configuration active on the switch. Whenever you make a configuration change, it is written to the running configuration. For fabric cluster mode, the running configuration does not persist across reboot, unless you copy it to the startup configuration.</p> <p>However, when the switch is in logical chassis cluster mode, the running-config file is saved automatically and it does not need to be copied.</p>

Configuration management follows a transaction model. When you boot up a switch for the first time, the running configuration is identical to the startup configuration. As you configure the switch, the changes are written to the running configuration. To save the changes, you must save the currently effective configuration (the running configuration) as the startup configuration. When the switch reboots, the configuration changes become effective.

Default configuration

Default configuration files are part of the Network OS firmware package and are automatically applied to the startup configuration under the following conditions:

- When the switch boots up for the first time and no customized configuration is available.
- When you restore the default configuration.

You cannot remove, rename, or change the default configuration.

Startup configuration

NOTE

There is no startup configuration for logical chassis cluster mode. Switches in a logical chassis cluster always preserve their running configuration.

The startup configuration is persistent. It is applied when the system reboots.

- When the switch boots up for the first time, the switch uses the default configuration as the startup configuration, depending on the mode.
- The startup configuration always matches the current Brocade VCS Fabric mode. It is deleted when you change modes, unless you make a backup copy.
- When you make configuration changes to the running configuration and save the changes to the startup configuration with the **copy** command, the running configuration becomes the startup configuration.

Running configuration

The configuration currently effective on the switch is referred to as the running configuration. Any configuration change you make while the switch is online is made to the running configuration.

- The running configuration is nonpersistent.
- To save configuration changes, you must copy the running configuration to the startup configuration. If you are not sure about the changes, you can copy the changes to a file, and apply the changes later.

Displaying configurations

The following examples illustrate how to display the default, startup, and running configurations, respectively.

Displaying the default configuration

To display the default configuration, enter the **show file filename** command in privileged EXEC mode.

```
switch# show file defaultconfig.novcs
switch# show file defaultconfig.vcs
```

Displaying the startup configuration

To display the contents of the startup configuration, enter the **show startup-config** command in privileged EXEC mode.

```
switch# show startup-config
```

Displaying the running configuration

To display the contents of the running configuration, enter the **show running-config** command in the privileged EXEC mode.

```
switch# show running-config
```

Saving configuration changes

Configuration changes are nonpersistent and are lost on reboot unless you save them permanently. You have two options for saving configuration changes:

- Copy the running configuration to the startup configuration. The changes become effective upon reboot.
- Copy the running configuration to a file, and apply it at some later date.

NOTE

Always make a backup copy of your running configuration before you upgrade or downgrade the firmware.

Saving the running configuration

To save the configuration changes you made, copy the running configuration to the startup configuration. The next time the switch reboots, it uses the startup configuration and the changes you made earlier become effective.

NOTE

When the switch is in logical chassis cluster mode, the running-config file is saved automatically and it does not need to be copied.

Enter the **copy running-config startup-config** command in privileged EXEC mode.

```
switch# copy running-config startup-config
copy running-config startup-config
This operation will modify your startup configuration. Do you want to continue? [Y/N]: y
```

Saving the running configuration to a file

If you want to save the changes you made to the configuration, but you do not want the changes to take effect when the switch reboots, you can save the running configuration to a file. You can apply the changes at some later time.

1. Enter the **copy running-config file** command in privileged EXEC mode. Specify the file name as the file URL.

```
switch# copy running-config flash://myconfig
```

2. Verify the transaction by listing the directory contents.

```
switch# dir
total 32
drwxr-xr-x  2 root    sys      4096 Feb 17 17:50 .
drwxr-xr-x  3 root    root     4096 Jan  1 1970 ..
-rwxr-xr-x  1 root    sys       417 Oct 12 2010 defaultconfig.novcs
-rwxr-xr-x  1 root    sys       697 Oct 12 2010 defaultconfig.vcs
-rw-r--r--  1 root    root     6777 Feb 17 17:50 myconfig
-rw-r--r--  1 root    root     6800 Feb 13 00:37 startup-config
```

Applying previously saved configuration changes

When you are ready to apply the configuration changes you previously saved to a file, copy the file (*myconfig* in the example) to the startup configuration. The changes take effect after the switch reboots.

Enter the **copy filename startup-config** command in privileged EXEC mode. Specify the file name as the file URL.

```
switch# copy flash://myconfig startup-config
This operation will modify your startup configuration. Do you want to continue? [Y/N]: y
```

Backing up configurations

Always keep a backup copy of your configuration files, so you can restore the configuration in the event the configuration is lost or you make unintentional changes.

NOTE

This operation is not supported in logical chassis cluster mode, because the running-config will be auto-synced to the startup-config.

The following recommendations apply:

- Keep backup copies of the startup configuration for all switches in the fabric.
- Upload the configuration backup copies to an external host or to an attached Brocade-branded USB device.
- Avoid copying configuration files from one switch to another. Instead restore the switch configuration files from the backup copy.

Uploading the startup configuration to an external host

Enter the **copy startup-config destination_file** command in privileged EXEC mode.

In the following example, the startup configuration is copied to a file on a remote server by means of FTP.

```
switch# copy startup-config ftp://admin:*****@122.34.98.133//archive/startup-config_vdx24-08_20101010
```

Backing up the startup configuration to a USB device

When you make a backup copy of a configuration file on an attached USB device, the destination file is the file URL on the USB device. You do not need to specify the target directory. The file is automatically recognized as a configuration file and stored in the default configuration directory.

1. Enable the USB device.

```
switch# usb on
USB storage enabled
```

2. Enter the **copy startup-config destination_file** command in privileged EXEC mode.

```
switch# copy startup-config usb://startup-config_vdx24-08_20101010
```

Configuration restoration

Restoring a configuration involves overwriting a given configuration file on the switch by downloading an archived backup copy from an external host or from an attached USB device.

A typical scenario for configuration restoration is:

- [Restoring the default configuration](#) on page 85.

All interfaces remain online. The following parameters are unaffected:

- Interface management IP address
- Software feature licenses installed on the switch
- Virtual IP address

NOTE

Configuration files that were created using Brocade Network OS 2.x should not be loaded onto a system running Brocade Network OS 3.x or later. The ACL and VLAN configuration information has changed in Brocade Network OS 3.x or later, and the affected lines of configuration are skipped when loading a Brocade Network OS 2.x configuration file.

Restoring the default configuration

This restoration procedure resets the configuration to the factory defaults. The default configuration files are always present on the switch and can be restored with the **copy** command.

To restore the default configuration, perform the following procedure in privileged EXEC mode.

1. Enter the **copy source_file destination_file** command to overwrite the startup configuration with the default configuration.

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

2. Confirm that you want to make the change by entering Y when prompted.

```
This operation will modify your startup configuration. Do you want to continue? [Y/N]: y
```

3. Reboot the switch.

```
switch# reload system
```

Managing configurations on a modular chassis

NOTE

When the switch is in logical chassis cluster mode, the running-config file is saved automatically and does not need to be copied. There is no startup configuration for logical chassis cluster mode; therefore, the information about startup configuration does not apply to logical chassis cluster mode.

The configuration data on a modular chassis are managed in a distributed fashion. The Brocade VDX 8770-4 and VDX 8770-8 chassis maintain two types of configuration data, global configuration parameters and slot configuration parameters. The global configuration, such as the VLAN configuration, applies to the entire chassis. The slot configuration includes specific parameters that apply only to the line cards.

The startup configuration is maintained at the chassis level and includes both chassis-wide and slot-specific configuration parameters.

Managing configurations on line cards

When an line card (interface module) boots up in a slot which was never occupied previously or is not configured, the module type is automatically saved in the configuration database. The type configuration associated with a given slot persists in the database even after the line card is physically removed, powered off, or faulted. This mechanism ensures that all configuration data associated with a given slot is automatically preserved across reboots or hot swaps with the same type of line card.

If you insert an line card in a slot that was previously occupied by a module of a different type, the line card will be faulted with a "type mismatch" error. Before you replace an line card with a different type, you must clear the existing type configuration from the database. Refer to [Replacing a line card](#) on page 66 for more information.

NOTE

The line card configuration is non-persistent. You must issue the **copy running-config startup-config** command after the line card comes online. Otherwise, all configuration data associated with the slot along with line module type will be lost after a chassis reboot.

Managing configurations across redundant management modules

In modular switches with redundant management modules, the VCS configuration, the startup configuration, and the startup database are synchronized and shared between the two management modules. The initial configuration synchronization occurs when the system boots up. After the initial synchronization has been completed successfully, synchronization can be triggered during the following events:

- When a failover occurs from the active management module to the standby management module. Unsaved configuration changes made on the active management module are lost after a failover. Issue the **copy running-config startup-config** command on the active management module to preserve the running configuration across a management module failover.
- When you insert a standby management module into a chassis after the active management module is already fully initialized.
- When you change the startup configuration by issuing the **copy running-config startup-config** command on the active management module.
- When you restore the default configuration by issuing the **copy default-config startup-config** command on the active management module.
- When you change the VCS configuration (VCS mode, RBridge ID, or VCS ID), the configuration change is synchronized with the standby management module and saved persistently. This event triggers a chassis reboot after the synchronization is complete.
- When you initiate a firmware download. Refer to [Configuration Management](#) on page 81 for more information.

Managing configurations in Brocade VCS Fabric mode

With the exception of a few parameters, configuration changes you make to a single switch in a Brocade VCS Fabric are not automatically distributed. When configuring Ethernet fabric parameters and software features on multiple switches, you must configure each switch individually. To simplify the procedure, you can upload a configuration file from one switch and download it to the other switches in the fabric, provided the switches are of the same type.

NOTE

The switches must be of the same model to share a configuration file. For example, downloading a configuration file from a Brocade VDX 6740-48 to a Brocade VDX 6740-64 or to a switch with a different firmware version may cause the switch to misapply the configuration and lead to unpredictable behavior.

To determine the switch type, issue the **show system** command. To map the switch type to the Brocade switch model name, refer to [Switch types](#) on page 37.

If you need to reset affected switches, restore the default configuration as described in [Restoring the default configuration](#) on page 85.

Automatic distribution of configuration parameters

A few configuration parameters are fabric-wide in fabric cluster mode. This means they are automatically distributed to all switches in a VCS fabric when you configure one or more of these parameters on a single RBridge that is part of a VCS fabric. These parameters include the following:

- Zoning configuration
- vCenter parameters
- Virtual IP address

NOTE

In logical chassis cluster mode, all configuration is automatically distributed.

The **show running configuration** command displays the same configuration for these features on all RBridges in the VCS fabric. Copy operations from any RBridge include all fabric-wide configuration parameters.

Downloading a configuration to multiple switches

NOTE

This section does not apply to logical chassis cluster mode because, in that mode, configuration is automatically distributed.

1. Configure one switch.
2. Copy the running configuration to the startup configuration as described in [Saving the running configuration](#) on page 83.
3. Upload the configuration to an external host ([Uploading the startup configuration to an external host](#) on page 84) or to an attached USB device as described in [Backing up the startup configuration to a USB device](#) on page 85.
4. Download the configuration file to each of the target switches. Refer to [Configuration restoration](#) on page 85 for more information.

Rejoining an offline node to a logical chassis cluster

If a node goes into offline state, and configuration changes are made to online nodes while the cluster is in a degraded state, database mismatches can occur when the offline node tries to rejoin the cluster.

Situations that can cause a node to go offline include:

- ISLs are shutdown, isolating a node from the fabric and cluster.
- A node reboots.

The following list describes possible scenarios that can occur when an offline node tries to rejoin its cluster, and what actions, if any, you must take. For more information about commands that are referenced, refer to the *Network OS Command Reference*.

- If local-only configurations have been added, updated, or deleted for online nodes after another node has temporarily left the cluster, the offline node automatically rejoins the cluster. Any non-default configurations specific to the rejoining node are then pushed back to the cluster configuration. If the rejoining node has the default configuration, and if no local-only configuration changes were made while this node was offline, then the cluster's configuration for the rejoining node is pushed onto the rejoining node.
- If the offline node has local-only configuration changes and its global configuration is non-default and matches the global configuration of the cluster, then the offline node is allowed to rejoin the cluster. The local-only configuration changes that were made while the node was offline are preserved.
- If global configurations have been added, updated, or deleted after another node has temporarily left the cluster, the global configuration differences between the cluster and the rejoining node result in a configuration database mismatch and node segmentation. To rejoin the cluster, issue the **copy default-config startup-config** on the segmented node.
- If the rejoining node's global configuration is the default configuration, and both of the following are true: 1) local-only configuration changes have been made to the rejoining node while it was offline, and 2) the cluster contains a different set of local configurations for the rejoining node, then a configuration database mismatch occurs. This situation can occur if the user issues a **copy default to start** command on a segmented node and issues local-only configurations on the segmented node before it rejoins the cluster. To rejoin the cluster, issue the **copy default to startup** command on the segmented node or cluster island.

The following table lists commands that change a local configuration for a node but also affect the global configuration. If these commands are run on online nodes in the cluster, be aware that global configuration in the cluster will also change.

TABLE 13 Local-configuration commands that affect global configuration

Command (Local Configuration)	Description
flexport <RBridge-ID/slot/port>	<p>Done in Hardware configuration mode, this command, which converts an Ethernet interface to a Fibre-Channel interface, can cause global configuration changes because the Ethernet interface affects L2Sys, SPAN, IGMPs, and MLD configuration settings.</p> <p>NOTE Only when you issue the type fibre-channel command from the flexport rbridge-id/slot/port prompt does this local-only change also change the global configuration.</p>
vrf <name>	<p>Done in RBridge-ID configuration mode, the creation of a VRF on an RBridge is a local-only VRF configuration exception and also changes the global configuration.</p>

Managing flash files

Brocade Network OS provides a set of tools for removing, renaming, and displaying files you create in the switch flash memory. You can use the display commands with any file, including the system configuration files. The **rename** and **delete** commands only apply to copies of configuration files you create in the flash memory. You cannot rename or delete any of the system configuration files.

Listing the contents of the flash memory

To list the contents of the flash memory, enter the **dir** command in privileged EXEC mode.

```
switch# dir
drwxr-xr-x  2 root    sys      4096 Feb 13 00:39 .
drwxr-xr-x  3 root    root     4096 Jan  1 1970 ..
-rwxr-xr-x  1 root    sys      417 Oct 12 2010 defaultconfig.novcs
-rwxr-xr-x  1 root    sys      697 Oct 12 2010 defaultconfig.vcs
-rw-r--r--  1 root    root     6800 Feb 13 00:37 startup-config
```

Deleting a file from the flash memory

To delete a file from the flash memory, enter the **delete file** command in privileged EXEC mode.

```
switch# delete myconfig
```

Renaming a flash memory file

To rename a file in the flash memory, enter the **rename source_file destination_file** command in privileged EXEC mode.

```
switch# rename myconfig myconfig_20101010
```

Viewing the contents of a file in the flash memory

To investigate the contents of a file in the flash memory, enter the **show file file** command in privileged EXEC mode.

```
switch# show file defaultconfig.novcs
!
no protocol spanning-tree
!
vlan dot1q tag native
!
cee-map default
remap fabric-priority priority 0
remap lossless-priority priority 0
priority-group-table 1 weight 40 pfc on
priority-group-table 2 weight 60 pfc off
priority-group-table 15.0 pfc off
priority-table 2 2 2 1 2 2 2 15.0
!
interface Vlan 1
shutdown
!
port-profile default
vlan-profile
switchport
switchport mode trunk
switchport trunk allowed vlan all
!
protocol lldp
!
end
!
```

NOTE

To display the contents of the running configuration, use the **show running-config** command. To display the contents of the startup configuration, use the **show startup-config** command.

Configuring SNMP

- [Simple Network Management Protocol overview](#).....91
- [SNMP configuration](#).....92

Simple Network Management Protocol overview

Simple Network Management Protocol (SNMP) is a set of protocols for managing complex networks. SNMP protocols are application layer protocols. Using SNMP, devices within a network send messages, called protocol data units (PDUs), to different parts of a network. Network management using SNMP requires three components:

- SNMP Manager
- SNMP Agent
- Management Information Base (MIB)

SNMP Manager

The SNMP Manager can communicate to the devices within a network using the SNMP protocol. Typically, SNMP Managers are network management systems (NMS) that manage networks by monitoring the network parameters, and optionally, setting parameters in managed devices. Normally, the SNMP Manager sends read requests to the devices that host the SNMP Agent, to which the SNMP Agent responds with the requested data. In some cases, the managed devices can initiate the communication, and send data to the SNMP Manager using asynchronous events called traps.

SNMP Agent

The SNMP agent is a software that resides in the managed devices in the network, and collects data from these devices. Each device hosts an SNMP Agent. The SNMP Agent stores the data, and sends these when requested by an SNMP Manager. In addition, the Agent can asynchronously alert the SNMP Manager about events, by using special PDUs called traps.

Management Information Base (MIB)


SNMP Agents in the managed devices store the data about these devices in a database called Management Information Base (MIB). The MIB is a hierarchical database, which is structured on the standard specified in the RFC 2578 [Structure of Management Information Version 2 (SMIv2)].

The MIB is a database of objects that can be used by a network management system to manage and monitor devices on the network. The MIB can be retrieved by a network management system that uses SNMP. The MIB structure determines the scope of management access allowed by a device. By using SNMP, a manager application can issue read or write operations within the scope of the MIB.

Basic SNMP operation


Every Brocade device carries an *agent* and management information base (MIB), as shown in the next figure. The agent accesses information about a device and makes it available to an SNMP network management station.

FIGURE 16 SNMP structure



When active, the management station can "get" information or "set" information when it queries an agent. SNMP commands, such as **get**, **set**, **getnext**, and **getresponse**, are sent from the management station, and the agent replies once the value is obtained or modified as shown in the next figure. Agents use variables to report such data as the number of bytes and packets in and out of the device, or the number of broadcast messages sent and received. These variables are also known as managed objects. All managed objects are contained in the MIB.

FIGURE 17 SNMP query



The management station can also receive *traps*, unsolicited messages from the switch agent if an unusual event occurs as shown in the next figure.

FIGURE 18 SNMP trap



The agent can receive queries from one or more management stations and can send traps to up to six management stations.

SNMP configuration

The following sections discuss configuring the Simple Network Management Protocol on Brocade devices. This includes configuring SNMP community strings, SNMP server hosts, SNMP server contexts, password encryption for SNMPv3 users, and displaying SNMP configurations.

Configuring SNMP community strings

SNMP versions 1 and 2c use community strings to restrict access to the switch. There is support for a total of 256 SNMP communities.

Adding an SNMP community string

The **snmp-server community** command sets the community string and associates it with the user-defined group to restrict the access of MIBs for SNMPv1 and SNMPv2c requests. You execute this command in global configuration mode.

1. Enter the **configure** command to access global configuration mode.

```
switch# configure
```

2. Enter the **snmp-server community** command.

```
switch(config)# snmp-server community comm1 groupname accGroup1
switch(config)#
```

The following example also applies an IPv4 ACL and an IPv6 ACL.

```
switch(config)# snmp-server community comm1 groupname accGroup1 ipv4-acl standV4ACL1 ipv6-acl
standV6ACL1
switch(config)#
```

NOTE

For the entire flow of implementing SNMP ACLs, refer to [Implementation flow of ACLs under SNMP](#) on page 99.

NOTE

When creating a new community string without specifying a group name, there is no group name associated with the community string. You must associate the community string with any nonexistent or existing group name to be able to contact the switch using SNMPv1/v2c.

Removing an SNMP community string

The following example removes the community string "public" and its associated group name "user".

1. Enter the **configure** command.
2. Enter the **no snmp-server community string** command.

```
switch(config)# no snmp-server community public
```

The following example removes the associated IPv4 and IPv6 ACLs from the SNMP community.

```
switch(config)# no snmp-server community public ipv4-acl
switch(config)# no snmp-server community public ipv6-acl
```

Configuring SNMP server hosts

The **snmp-server host** command sets the trap destination IP addresses, SNMP version, community string for SNMPv1 and SNMPv2c, the destination port for the SNMP server host, and the severity level.

To configure SNMP trap hosts associated with community strings, you must create the community string using the **snmp-server community** command before configuring the host.

The SNMP agent supports six trap-recipient severity levels. The default value for each attribute is as follows: host = 0.0.0.0; udp-port = 162; severity-level = none. The length of the community string must be from 2 through 16 characters.

Setting the SNMP server host

You can execute SNMP server commands in global configuration mode.

1. Enter the **configure** command.
2. Enter the **snmp-server host ipv4_host | ipv6_host | dns_host community-string [version { 1 | 2c }] [udp-port port] [severity-level | none | debug | info | warning | error | critical] [use-vrf { mgmt-vrf | default-vrf }]** command.
 - The *ipv4_host | ipv6_host | dns_host* variable specifies the IP address of the host.
 - The *community-string* variable sets the community string.
 - The **version** option specifies either SNMPv1- or SNMPv2c-related configuration parameters. These parameters include the community string. The default SNMP version is 1.
 - The **udp-port port** option specifies the UDP port where SNMP traps will be received. The default port is 162. The acceptable range of ports is from 0 through 65535.
 - The **severity sev_level** option provides the ability to filter traps based on severity level on both the host and the v3host. Only RASLog (swEvent) traps can be filtered based on severity level. If the severity level of None is specified, all traps are filtered and no RASLog traps are received. If the severity level of Critical is specified, no traps are filtered and all traps are received by the host.
Severity level options include None, Debug, Info, Warning, Error, and Critical.
 - The **use-vrf** option configures SNMP to use the selected VRF to communicate with the host. This parameter is optional. The VRF name can be only two alphanumeric strings, "mgmt-vrf" and "default-vrf". The default option is "mgmt-vrf".

The following example sets up "commaccess" as a read-only community string and sets 10.32.147.6 as a trap recipient with SNMPv2c on target port 162, with default-vrf.

```
switch(config)# snmp-server host 10.32.147.6 commaccess version 2c udp-port 162 severity warning use-  
vrf default-vrf
```

Removing the SNMP server host

The **no snmp-server host host community-string string version 2c use-vrf default-vrf** command brings version 2c down to version 1 and removes the default-vrf mapping with the host.

The **no snmp-server host host community-string string** command or the **no snmp-server v3host host** command removes the SNMP server host from the switch configuration altogether.

Configuring the SNMP system group

The following tasks allow you to configure the system contact and system location objects for the SNMP system group.

Setting the SNMP server contact

Use the **snmp-server contact** command to set the SNMP server contact string. The default contact string is "Field Support." The number of characters allowed is from 4 through 255.

1. Enter the **configure** command.
2. Enter the **snmp-server contact string** command.

```
switch(config)# snmp-server contact "Operator 12345"
```

The example changes the default contact string to "Operator 12345." You must enclose the text in double quotes if the text contains spaces.

Removing the SNMP server contact

The `no snmp-server contact string` command restores the default contact information (Field Support).

Setting the SNMP server location

Use the `snmp-server location` command to set the SNMP server location string. The default SNMP server location string is "End User Premise." The number of characters allowed is from 4 through 255.

1. Enter the `configure` command.
2. Enter the `snmp-server location string` command.

```
switch(config)# snmp-server location "Building 3 Room 214"
```

3. Enter the `no snmp-server location` command to remove the location.

The example changes the default location string to "Building 3 Room 214." You must enclose the text in double quotes if the text contains spaces.

Setting the SNMP server description

Use the `snmp-server sys-descr` command to set the SNMP server description string. The default SNMP server description string is "Brocade-VDX-VCS <vcsid>." The number of characters allowed is from 4 through 255.

1. Enter the `configure` command.
2. Enter the `snmp-server sys-descr string` command.

```
switch(config)# snmp-server sys-descr "Brocade-VDX Test Bed"
```

3. Enter the `no snmp-server sys-descr` command to remove the location.

The example changes the default location string to "Brocade-VDX Test Bed." You must enclose the text in double quotes if the text contains spaces.

Configuring multiple SNMP server contexts

A single SNMP agent can be supported by the multiple instances of the same MIB module by mapping the context name to a virtual routing and forwarding (VRF) instance as described below. The SNMP agent supports 256 contexts to support context-to-VRF mapping. Do the following to set the SNMP server context, using the `snmp-server context` command to map a context to the name of a VRF instance.

1. Enter the `configure` command.
2. Enter the `snmp-server context context_name vrf-name vrf_name` command.

```
switch(config)# snmp-server context mycontext vrf-name myvrf
```

The example maps the context name "mycontext" to the VRF name "myvrf."

Configuring SNMP server views

SNMP views are named groups of MIB objects that can be associated with user accounts to allow limited access for viewing and modification of SNMP statistics and system configuration.

Use the following procedure to create or remove a view entry with MIB objects to be included or excluded for user access.

1. Enter the **configure** command.
2. Enter the **snmp-server view** *view-name mib_tree* **{included | excluded}** command.

```
switch(config)# snmp-server view view1 1.3.6.1.2.1.1.3 excluded
```

Enter the **no** form of the command to remove the configured SNMP server view entry.

NOTE

The maximum number of views supported with MIB tree entries is 10.

Configuring SNMP server groups

SNMP groups map the SNMP user for the version v3 and the community for the versions v1 and v2c to SNMP views. Each SNMP group can be configured with a read view, a write view, a notify view, or all of the above.

Use the following procedure to configure or remove a specified SNMP group.

1. Enter the **configure** command.
2. Enter the **snmp-server group** *groupname* **{v1 | v2c | v3 {auth | noauth | priv}}** **[read viewname] [write viewname] [notify viewname]** command.

```
switch(config)# snmp-server group group1 v3 auth read myview write myview notify myview
```

Enter the **no** form of the command to remove the configured SNMP server group.

NOTE

The maximum number of SNMP groups supported is 10.

Configuring SNMP server users

The **snmp-server user** command configures a SNMPv3 user and allows the configured user to be associated with user-defined SNMP groups. You execute this command in global configuration mode and RBridge ID configuration mode.

1. Enter the **configure** command to access global configuration mode.

```
switch# configure
```


2. Enter the **snmp-server user** command.

The following example configures the SNMPv3 users "user1" and "user2" associated with user-defined group "group1" under global configuration mode.

```
switch(config)# snmp-server user user1 groupname group1
switch(config)# snmp-server user user2 groupname group1 auth md5 auth-password password priv DES
priv-password
```

The following example configures the SNMPv3 users "user1" and "user2" associated with user-defined group "group1" under global configuration mode. It also applies an IPv4 ACL and an IPv6 ACL to "user1."

```
switch(config)# snmp-server user user1 groupname group1 ipv4-acl standV4ACL1 ipv6-acl standV6ACL1
switch(config)# snmp-server user user2 groupname group1 auth md5 auth-password password priv DES
priv-password
```

NOTE

You can associate only global SNMPv3 users with ACLs. For the entire flow of implementing SNMP ACLs, refer to [Implementation flow of ACLs under SNMP](#) on page 99.

The following example configures the SNMPv3 user "snmpadmin1" under RBridge ID configuration mode.

```
switch(config-rbridge-id-1)# snmp-server user snmpadmin1 groupname snmpadmin auth sha auth-password
privatel123 priv DES priv-password public123
```

NOTE

When creating a new SNMPv3 user without the **groupname** option, by default there is no group name mapped with the SNMPv3 user. You must map the configured SNMPv3 user with a group name available in the group CLI configuration to contact the switch through SNMPv3.

NOTE

The behavior of this command in RBridge ID configuration mode is same as in global configuration mode. If the user name is configured to be the same in both global and RBridge ID configurations, the RBridge ID configuration takes precedence. The encrypted password generated in global configuration mode can be used for another global user to modify the passwords. The encrypted passwords generated in global configurations cannot be used in RBridge ID configurations, and vice versa.

Configuring password encryption for SNMPv3 users

For SNMPv3 users, the passwords for **auth-password** and **priv-password** are encrypted. You can configure either with a plain-text password or an encrypted password. In both cases, the passwords are shown in the **show running-config** command as encrypted.

The following example shows a plain-text password configuration:

```
switch(config)# snmp-server user snmpadmin1 auth md5 auth-password privatel123 priv DES priv-password
public123
```

The following example shows an encrypted password configuration:

```
switch(config)# snmp-server user snmpadmin2 groupname snmpadmin auth md5 authpassword "MvB
+360X3kcfBzug5Vo6dQ==\n" priv DES priv-password "ckJFoHbzVvhR0xFRPjsMTA== \n" encrypted
```

NOTE

This command may not be successful where encrypted passwords are generated by third-party or open-source tools.

Configuring SNMP server v3hosts

The `snmp-server v3host` command configures a SNMPv3 host by associating with the SNMP users. You execute this command in global configuration mode and RBridge ID configuration mode.

Use the following procedure to configure a SNMPv3 host.

1. Enter the `configure` command.
2. Enter the `snmp-server v3host [host { ipv4_host | ipv6_host | dns_host}] user_name[notifytype {traps | informs}] engineid engine-id udp-port port_number [severity-level | {none | debug | info | warning | error | critical}] [use-vrf { mgmt-vrf | default-vrf }]` command.

```
switch(config)# snmp-server v3host 1050::5:600:300c:326b snmpadmin2 severity-level Info use-vrf
default-vrf
```

This example configures the SNMPv3 trap IPv6 host 1050:0:0:0:5:600:300c:326b associated with SNMP user "snmpadmin2" under global configuration mode with default vrf.

```
switch(config-rbridge-id-1)# snmp-server v3host 10.26.3.166 snmpuser2 severity-level Info udp-port
4425
```

This example configures the SNMPv3 trap host 10.26.3.166 associated with SNMP user "snmpuser2" under RBridge ID configuration mode.

The global SNMPv3 host can be configured by associating with only global SNMPv3 users and the local SNMPv3 host can be configured by associating with only local SNMPv3 users. You cannot create a SNMPv3 host in global configuration by associating with the local SNMPv3 users and vice versa.

Managing SNMP access rights using ACLs

Access lists (ACLs) enable you to permit or deny SNMP access by IP address.

SNMP server groups enable you to specify read, write, and notify permissions for the following entities:

- Community, under SNMPv1 and SNMPv2c
- User, under SNMPv3

For SNMP packets that pass community/user validation, access lists (ACLs) offer an additional permit/deny level, filtered by IP addresses that you specify.

If SNMP ACLs are applied, the validation order is as follows:

1. SNMP-server validation (community/user string). If not validated, the SNMP packet is dropped.
2. Server-ACL validation
 - If there is a **deny** match—including an explicit or implicit `deny any` rule—the packet is dropped.

NOTE

Unless you include an explicit `permit any` rule, an implicit `deny any` rule is automatically applied for IP addresses not explicitly permitted.

- If there is a **permit** match—including a `permit any` rule—validation continues.
3. Server-group validation, the concluding step of the validation flow

Implementation flow of ACLs under SNMP

The implementation flow for ACLs under SNMP is as follows:

1. Create access lists (ACLs) that permit or deny specified IP addresses. For details, refer to [Creating standard ACLs for SNMP](#) on page 99.
2. Define server groups with the needed combination of Read or Write; and Notify permissions. For details, refer to [Adding an SNMP community string](#) on page 93.
3. For SNMPv1 or SNMPv2c, implement [Adding an SNMP community string](#) on page 93.
4. For SNMPv3, implement [Configuring SNMP server users](#) on page 96.

Creating standard ACLs for SNMP

Use these procedures to create access lists (ACLs) that contain rules permitting or denying access from specified IP addresses.

Creating an IPv4 ACL for SNMP

A standard ACL permits or denies traffic according to source address only. SNMP supports only standard ACLs.

1. Enter **configure** to access global configuration mode.

```
switch# configure
```

2. Enter the **ip access-list standard** command to create the access list.

```
switch(config)# ip access-list standard stdACL3
```

3. For each ACL rule, enter a **seq** command, specifying the needed parameters.

```
switch(config-ipacl-std)# seq 5 permit host 10.20.33.4
switch(config-ipacl-std)# seq 15 deny any
```

The following example does the following, under SNMPv3:

1. Creates a IPv4 standard ACL named "test".
2. Defines rules that permits packets from a specified host and denies packets from any other host.
3. Configures the SNMP server user "user1", including application of the "test" IPv4 ACL.

```
switch(config)# ip access-list standard test
switch(conf-ipacl-std)# permit host 10.1.1.1
switch(conf-ipacl-std)# deny any
switch(conf-ipacl-std)# exit
switch(config)# snmp-server user user1 groupname snmpadmin auth sha auth-password private123 priv DES priv-
password public123 ipv4-acl test
```

Creating an IPv6 ACL for SNMP

A standard ACL permits or denies traffic according to source address only. SNMP supports only standard ACLs.

1. Enter **configure** to access global configuration mode.

```
switch# configure
```

2. Enter the **ipv6 access-list standard** command to create the access list.

```
switch(config)# ipv6 access-list standard std_V6_ACL4
```

- For each ACL rule, enter a **[seq] {permit | deny | hard-drop}** command, specifying the needed parameters.

```
switch(config-ip6acl-std)# seq 5 permit host 2001:db8::1:2
switch(config-ip6acl-std)# seq 15 deny any
```

The following example does the following, under SNMPv1 or SNMPv2c:

- Creates an IPv6 standard ACL named "stdv6acl".
- Defines rules that permits packets from a specified host and denies packets from any other host.
- Configures the SNMP server community "c1", including application of the "stdv6acl" IPv6 ACL.

```
switch(config)# ipv6 access-list standard stdv6acl
switch(conf-ip6acl-std)# permit fe::/24
switch(conf-ip6acl-std)# deny any
switch(conf-ip6acl-std)# exit
switch(config)# snmp-server community c1 groupname admin ipv6-acl stdv6acl
```

Displaying SNMP configurations

Use the **show running-config snmp-server** command to display the current SNMP configurations for the SNMP host, community string, contact, and location, as well as other SNMP configuration options such as SNMPv3 host address, context, VRF mapping, and applied ACLs.

Enter the **show running-config snmp-server** command.

```
switch# show running-config snmp-server

snmp-server contact "Field Support."
snmp-server location "End User Premise."
snmp-server sys-descr "Brocade VDX Switch."
snmp-server enable trap
snmp-server community ConvergedNetwork groupname user
snmp-server community OrigEquipMfr groupname admin
snmp-server community "Secret C0de" groupname admin
snmp-server community c1 groupname admin ipv4-acl test ipv6-acl stdv6acl
snmp-server community c2 groupname g1 ipv4-acl 1 ipv6-acl 2
snmp-server community com1
snmp-server community common groupname user
snmp-server community private groupname admin
snmp-server community public groupname user
snmp-server host 10.20.53.161 private
  severity-level Info
  use-vrf mgmt-vrf
!
snmp-server user snmpadmin1 groupname user
snmp-server user snmpadmin2 groupname snmpadmin
snmp-server user snmpadmin3 groupname snmpadmin
snmp-server user snmpuser1 groupname snmpuser
snmp-server user snmpuser2 groupname snmpuser
snmp-server user snmpuser3 groupname snmpuser
snmp-server user user1 groupname admin auth sha auth-password "ggc+pJR30+ORXd6OULEU6GaUFN4=\n" priv DES
priv-password "4qrmMZ5wR1W9MQMKxrQWFpaAgl8=\n" encrypted ipv4-acl test ipv6-acl stdv6acl
snmp-server view All 1 included
snmp-server group admin v1 read All write All notify All
snmp-server group admin v2c read All write All notify All
snmp-server group snmpadmin v3 notify All
snmp-server group snmpuser v3 notify All
snmp-server group user v1 read All notify All
snmp-server group user v2c read All notify All
```

Configuring Brocade VCS Fabrics

- [Fabric overview](#)..... 101
- [Configuring a Brocade VCS Fabric](#)..... 104

Fabric overview

The Brocade VCS Fabric Ethernet fabric is defined as a group of switches that exchange information between each other to implement distributed intelligence. The Brocade Ethernet fabric uses Transparent Interconnection of Lots of Links (TRILL) protocol, designed for the sole purpose of scaling Ethernet networks by allowing a set of devices, called routing bridges (RBridges), to connect with each other.

A link state dynamic routing protocol, rather than Spanning Tree Protocol, determines how the traffic is forwarded between the interconnected RBridges. Link state routing in Brocade VCS Fabric-based TRILL networks is performed using Fabric Shortest Path First (FSPF) protocol.

TRILL enables Layer 2 networks to behave like routed Layer 3/IP networks. TRILL also defines native support for forwarding both unicast and multicast traffic, and therefore unifies support for both of these different classes of applications over a single transport.

Brocade VCS Fabric formation

Brocade VCS Fabric technology uses RBridge identifiers (IDs) to discover fabric creation problems, such as duplicate IDs. The RBridge ID of a cluster unit is equal to the domain ID of an FC switch. RBridge ID assignment is implemented by leveraging the domain ID assignment protocols in the FC SANs. Request for Domain ID (RDI) and Domain ID Assignment (DIA) protocols ensure that a single switch (the principal switch) is centrally allocating the domain IDs for every RBridge in the fabric and detecting any domain ID conflicts in the fabric. In case of conflict, the conflicting node is segmented from the fabric. You must take action to resolve the conflict

NOTE

Network OS 5.0.1 and later support a maximum of 48 RBridges in a single Brocade VCS Fabric, depending on the hardware mix. However, Brocade recommends using only 24 RBridges per fabric.

The following sequence of events describes the Brocade VCS Fabric formation process:

- Each Brocade VCS Fabric is identified by a VCS ID.
- All Brocade VCS Fabric-capable switches are configured with a factory default VCS ID of 1.
- The switch software searches for the value for the "VCS enable" attribute setting and verifies it is set to "enabled".
- Assuming the switch is Brocade VCS Fabric-enabled, the switch software invokes a series of protocols:
 - Brocade Link Discovery Protocol (BLDP) attempts to discover if a Brocade VCS Fabric-capable switch is connected to any of the edge ports. Refer to [Neighbor discovery](#) on page 102 for more information.
 - BLDP attempts to merge the adjacent Brocade switch into the Brocade VCS Fabric environment at the link level.
- A series of FC fabric formation protocols (RDI, DIA, and FSPF) are initiated once a link level relationship has been established between two neighbor switches. Refer to [Fabric formation](#) on page 103 for more information.
- The "Merge and Join" protocol invokes a merge of switch configuration between the cluster units once the fabric has successfully formed.

How RBridges work

RBridges find each other by exchanging FSPF Hello frames. Like all TRILL IS-IS frames, Hello frames are transparently forwarded by RBridges and are processed by RBridge Inter-Switch Link (ISL) ports. Using the information exchanged in the Hello frames, the RBridges on each link elect the designated RBridge for that link.

The RBridge link state includes information such as VLAN connectivity, multicast listeners, and multicast router attachment, claimed nicknames, and supported ingress-to-egress options. The designated RBridge specifies the appointed forwarder for each VLAN on the link (which could be itself) and the designated VLAN for inter-RBridge communication. The appointed forwarder handles native frames to and from that link in that VLAN.

The Ingress RBridge function encapsulates frames from the link into a TRILL data frame. The Egress RBridge function decapsulates native frames destined for the link from the TRILL data frames. TRILL data frames with known unicast destinations are forwarded by RBridge next hop. Multi-destination frames (broadcast, unknown unicast, and multicast) are forwarded on a tree rooted at the multicast root RBridge.

- Unicast forwarding is handled by combining domain routing generated by FSPF and MAC-to-RBridge learning generated by MAC learning and a distributed MAC database.
- Multicast forwarding usually uses one tree that is rooted at the RBridge with the lowest RBridge ID. However, there are several rules for Multicast root tree selection. It is not always the lowest RBridge ID.

If a duplicated RBridge ID is found while the links are still coming up, the links are segmented. Both sides recognize the error and segment the link. If the RBridge ID overlap cannot be found at ISL link bringup time (in the case where a new switch is brought from an offline state into the fabric) it will be found during the fabric build and the conflicting switch is isolated.

An RBridge requests a specific RBridge ID from the coordinator switch. If the coordinator switch detects that this RBridge ID is already used, it returns the next unused RBridge ID. The requesting RBridge is not allowed to take another RBridge ID and it segments itself from the fabric. In this case, you cannot boot the ISLs. The ISLs have to be explicitly disabled and then enabled again in order for the RBridge with the overlapping RBridge ID to be removed.

Neighbor discovery

Brocade VCS Fabric-capable neighbor discovery involves the following steps:

- Discover whether the neighbor is a Brocade switch.
- Discover whether the Brocade neighbor switch is Brocade VCS Fabric-capable.

Only Brocade VCS Fabric-capable switches with the same VCS ID can form a virtual cluster switch. The default settings for Brocade Network OS switches are Brocade VCS Fabric capable and a VCS ID of "1."

Brocade trunks

Network OS 4.0.0 and later supports Brocade trunks (hardware-based link aggregation groups, or LAGs). These LAGs are dynamically formed between two adjacent switches. The trunk formation is controlled by the same Fibre Channel Trunking protocol that controls the trunk formation on FC switches. As such, it does not require user intervention or configuration except enabling or disabling, which instructs the switch software to form a trunk at the global level or not. All ISL ports connected to the same neighbor Brocade switch will attempt to form a trunk. Refer to [Enabling a fabric trunk](#) on page 107 for instructions.

Ports groups have been established on supported standalone switches and on line cards in chassis systems for trunking. For a successful trunk formation, all ports on the local switch must be part of the same port group and must be configured at the same speed.

Following are the number of ports allowed per trunk from port groups in supported platforms. For details on how port groups are arranged on these platforms, refer to the switch or chassis system Hardware Reference Manual.

- VDX 8770 switches - up to six port groups of eight ports each per blade (1, 10, or 40 GbE)
- VDX 6740 switches - up to 16 ports per trunk.
- VDX 2740 - up to 16 ports per trunk.
- 48x10 GbE line card - up to eight ports per trunk.
- 48x10G-T line card - up to 16 ports per trunk.
- 12x40 GbE line card - up to two 40-GbE ports are allowed per trunk, and these ports must be configured in breakout mode.
- 27x40 GbE line card - up to two 40-GbE ports are allowed per trunk, and these ports must be configured in breakout mode. Note that breakout mode is only allowed on the first two ports in port groups that are configured in Performance operating mode. Refer to the *Brocade 8770 Hardware Reference Manual* for more information on line card operating modes.

The following additional rules apply to Brocade trunks:

- On Brocade VDX 6740 switches, low-volume traffic below certain thresholds may not be evenly distributed on all links. This threshold can be as low as 64 K.
- The trunk is turned on by default.
- Trunks are not supported between the Brocade 8000 and the Brocade VDX 8770.

Fabric formation

Brocade VCS Fabric technology leverages proven FC Fabric protocols to build a TRILL fabric. The main functions of the fabric formation protocols are as follows:

- Assign the Brocade VCS Fabric-wide unique RBridge IDs (Domain ID Assignment).
- Create the network topology database using link state routing protocol (Fabric Shortest Path First, or FSPF). FSPF calculates the shortest path routes to a destination RBridge.
- Distribute fabric multicast traffic.

Principal switch selection

Every Brocade VCS Fabric-enabled switch, upon boot-up and after the Fabric port formation, declares itself to be a principal switch and advertises this intent on all fabric ports. The intent includes a priority and its switch WWN. If all switches boot up at the same time, the default priority is the same and all switches will compare their mutual intents. The switch with the lowest Switch WWN becomes the principal switch. The WWN is an industry-standard burnt-in switch identifier, similar to the Bridge-MAC except it is 8 bytes. The role of the principal switch is to decide whether a new RBridge joining the fabric conflicts with any of the RBridge IDs already present in the fabric. At the end of the principal switch selection process, all the switches in the cluster have formed a tree with the principal switch at the root.

NOTE

Brocade VDX Data Center switches are shipped with factory-programmed world wide names (WWNs) that are unique.

NOTE

In a logical chassis cluster, you can select the principal node by using the command line interface. For more information, refer to [Selecting a principal node for the cluster](#) on page 60.

RBridge ID allocation

RBridge ID assignment is implemented by leveraging proven Domain ID assignment protocols from FC SANs. Request for Domain ID (RDI) and Domain ID Assignment (DIA) protocols ensure that a single switch (the principal switch) centrally allocates the domain IDs for every RBridge in the fabric and detects and resolves any domain ID collisions in the fabric. A Brocade VCS Fabric supports up to 48 RBridge IDs. RBridge IDs can range from 1 through 239.

Only the principal switch can allocate RBridge IDs (domain IDs) for all other switches in the fabric. The principal switch starts the allocation process by allocating an RBridge ID for itself (using the ID value supplied by the user), and initiates the DIA messages on all ports.

Other switches, which are now in subordinate mode, upon receiving the DIA frames respond with an RDI message towards the principal switch. The process continues until all the switches in the fabric have been allocated a unique ID.

Fabric routing protocol

After a RBridge ID is assigned to a switch, the Fabric Shortest Path First (FSPF) link state routing protocol begins to form adjacencies and collects topology and inter-connectivity information with its neighbors. Brocade VCS Fabric uses FSPF to calculate and elect a loop-free multicast tree rooted at the multicast root RBridge. The multicast tree is calculated after the unicast routes are computed.

Configuring a Brocade VCS Fabric

Refer to the following tables for commands and examples used in configuring a Brocade VCS Fabric. For command details, refer to the *Network OS Command Reference*.

The following table lists command examples for enabling Brocade VCS logical chassis cluster mode:

TABLE 14 Command examples for enabling logical chassis cluster mode

Command	Command Behavior
<code>switch# vcs logical-chassis enable</code>	The VCS ID becomes the default value of 1, the RBridge ID is not changed, and Brocade VCS logical chassis cluster mode is enabled.
<code>switch# vcs vcsid 22 rbridge-id 15 logical-chassis enable</code>	The VCS ID is changed to 22, the RBridge ID is changed to 15, and Brocade VCS logical chassis cluster mode is enabled. RBridge IDs range from 1 through 239.
<code>switch# vcs vcsid 11 logical-chassis enable</code>	The VCS ID is changed to 11, the RBridge ID is not changed, and Brocade VCS logical chassis cluster mode is enabled.
<code>switch# vcs rbridge-id 6 logical-chassis enable</code>	The VCS ID becomes the default value of 1, the RBridge ID is changed to 6, and Brocade VCS logical chassis cluster mode is enabled.

The following table lists command examples for enabling Brocade VCS fabric cluster mode:

TABLE 15 Command examples for enabling fabric cluster mode

Command	Command Behavior
<code>switch# vcs vcsid 55 rbridge-id 19 enable</code>	The VCS ID is changed to 55, the RBridge ID is changed to 19, and Brocade VCS fabric cluster mode is enabled.
<code>switch# vcs vcsid 73 enable</code>	The VCS ID is changed to the value of 73, the RBridge ID is not changed, and Brocade VCS fabric cluster mode is enabled.
<code>switch# vcs rbridge-id 10 enable</code>	The VCS ID becomes the default value 1, the RBridge ID is changed to 10, and Brocade VCS fabric cluster mode is enabled.

The following table lists command examples for switches that are already in either fabric cluster mode or logical chassis cluster mode.

TABLE 16 Command examples for when one of the VCS modes is already enabled:

Command	Command Behavior
<code>switch# vcs vcsid 44 rbridge-id 22</code>	The VCS ID is changed to 44 and the RBridge ID is changed to 22.
<code>switch# vcs vcsid 34</code>	The VCS ID is changed to 34.

Adding a new switch into a fabric

Complete the following configuration steps to add a new switch into a fabric.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **vcs rbridge-id *rbridge-id* enable** command.

The switch remembers its RBridge ID once it has been assigned. The **vcs rbridge-id *rbridge-id* enable** command also sets the insistent RBridge ID property on the switch.

3. Reboot the system.

After the required reboot the switch participates in the RBridge ID allocation protocol, which insists that the same value that was manually configured prior to reboot be allocated after reboot.

The switch is not allowed into the fabric if there is a conflict; for example, if another switch with the same ID exists and is operational in the fabric. You have the opportunity to select a new RBridge ID by using the same CLI.

Once an ID has been assigned by the fabric protocol, these IDs are then numerically equated to RBridge IDs and are treated as such after that.

Use the **vcs** command to configure the Brocade VCS Fabric parameters, VCS ID, and the switch RBridge ID, and to enable Brocade VCS Fabric mode (also called *VCS mode*).

VCS mode encompasses two mode types:

- *Fabric cluster mode* — The data path for nodes is distributed, but the configuration path is not distributed. Each node keeps its configuration database independently.
- *Logical chassis cluster mode* — Both the data and configuration paths are distributed. The entire cluster can be configured from the principal node. Logical chassis cluster mode requires Network OS 4.0 or later.

The generic term *VCS mode* in this manual applies to both fabric cluster mode and logical chassis cluster mode unless otherwise stated.

You can set the Brocade VCS Fabric parameters and enable VCS mode at the same time, or you can enable VCS mode and then perform the ID assignments separately. Refer to [Configuring a Brocade VCS Fabric](#) on page 104 for details.

After configuring the Brocade VCS Fabric parameters, the switch applies the changes and reboots.

The switch disable is not saved across a reboot, so if the switch was disabled prior to the reboot, the switch returns to the enabled state when it finishes the boot cycle.

Configuring fabric interfaces

A physical interface in a virtual switch cluster can either be an edge port or a fabric port, but not both. Similar to a switch-port configuration on a physical interface, you can also change a fabric-port configuration on its physical interface by using the **fabric ISL enable** and **fabric trunk enable** commands, described below.

Enabling a fabric ISL

The **fabric isl enable** command controls whether an ISL should be formed between two cluster members. With the default setting of ISL discovery to **auto** and the ISL formation mode to **enable**, an ISL automatically forms between two cluster switches.

Performing a **fabric isl enable** command on an operational ISL has no effect. However, performing a **no fabric isl enable** command on an interface toggles its link status and subsequently disables ISL formation. In addition, the **no fabric isl enable** command triggers the switch to inform its neighbor that the local interface is ISL disabled. Upon receiving such information, a neighbor switch stops its ISL formation activity regardless of its current interface state.

NOTE

After you repair any segmented or disabled ISL ports, toggle the fabric ISL in order to propagate the changes.

NOTE

A **shutdown** command on an operating ISL interface not only brings down the physical link but also its FSPF adjacency. The main difference between a **shutdown** command and a **no fabric isl enable** command is that the link stays up after a **no fabric isl enable**, while the link stays down after a shutdown.

NOTE

Upon a fabric reconvergence that due to a topology change involving the ECMP fabric-isl path, there may be sub-second flooding of known unicast traffic.

Disabling a fabric ISL

The **no fabric isl enable** command takes this interface out of the trunk group if this interface happens to be currently part of the trunk. If you know and would like to fix the edge and fabric port assignments on a switch, then this command allows you to completely turn off ISL formation logic and shorten any link bring-up delays on edge ports.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **no fabric isl enable** command.

Enabling a fabric trunk**NOTE**

Trunks are not supported between the Brocade 8000 and the Brocade VDX 8770.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **fabric trunk enable** command.

Disabling a fabric trunk

Fabric trunking is enabled by default. Enter the **no fabric trunk enable** command to revert the ISL back to a standalone adjacency between two Brocade VCS Fabric switch.

Configuring broadcast, unknown unicast, and multicast forwarding

All switches in a Brocade VCS Fabric cluster share a single multicast tree rooted at the RBridge with the lowest RBridge ID (domain ID). All broadcast, unknown unicast, and multicast traffic between two edge RBridges is forwarded on this multicast tree inside the Brocade VCS Fabric. The multicast tree includes all RBridges in the Brocade VCS Fabric.

Multicast distribution tree-root selection

Network OS supports the following distribution tree behaviors.

- The root of the distribution tree is the switch with the lowest RBridge ID. The automated selection process does not require any user intervention.
- Each switch in the cluster optionally carries a multicast root priority. This priority setting overrides the automatically-selected multicast root. In deployments where a multicast root is required to be a specific switch that does not have the lowest RBridge ID, then the priority setting on that switch can override the root selection. If there are two switches with the same priority, then the switch with the lower RBridge ID prevails.
- A back-up multicast root is pre-selected, which is the switch with the next lowest RBridge ID. The back-up multicast root is automatically selected by all switches should the current multicast root fail.

Configuring priorities

As stated above, the root of the tree is auto-selected as the switch with the lowest RBridge ID. For example, if you had a cluster with RBridge IDs 5, 6, 7, and 8, then 5 would be the root. If you then added RBridge ID 1 to this fabric, the tree would be re-calculated with 1 as the root.

In order to avoid this behavior, you can set a priority (default is 1). The highest priority overrides the lowest RBridge ID and becomes the root.

For example, to build a fabric with RBridge ID 7 or 8 as the root, set their priority to something higher than 1 (priority values are 1 through 255). If there is a tie in priority, the lower RBridge is still chosen. If RBridge ID 7 and 8 are both set to priority 1, 7 becomes the root.

Changing the priority

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **fabric route mcast rbridge-id** command.

Here is an example of changing an RBridge multicast priority:

```
switch(config)# fabric route mcast rbridge-id 12 priority 10
```

Displaying the running configuration

The **show running-config fabric route mcast** command allows you to display fabric route multicast configuration information. The configuration currently effective on the switch is referred to as the running configuration. Any configuration change you make while the switch is online is made to the running configuration. The running configuration is nonpersistent.

NOTE

To save configuration changes, you must save the running-config file to a file, or you can apply the changes by copying the running configuration to the startup configuration.

1. Connect to the switch and log in using an account assigned to the admin role.
2. Enter the **show running-config fabric route mcast** command.

```
switch# show running-config fabric route mcast priority
fabric route mcast rbridge-id 12 priority 10
```

Configuring VCS virtual IP addresses

A virtual IP address is assigned for each VCS cluster. This virtual IP address is tied to the principal switch in the cluster. The management interface of the principal switch can be accessed by means of this virtual IP address. Because the virtual IP address is a property of the fabric cluster and logical chassis cluster, in the event that the principal switch goes down, the next principal switch is assigned this address.

Virtual IP address can be configured by means of the **vcs virtual ip address** command:

```
switch(config)# vcs virtual ip address 10.0.0.23/24
```

This command can be used in logical chassis cluster and fabric cluster modes only. When the virtual IP address is configured for the first time, the current principal switch in the cluster is assigned this IP address.

Virtual IP configuration is global in nature. All the nodes in the cluster are configured with the same virtual IP address, but address is bound to the current principal switch only. Make sure that the assigned virtual IP address is not a duplicate of an address assigned to any other management port in the cluster or network.

Brocade recommends that you use a /32 address in the same subnet as the IP address of the management interface. For example, if you are using inband management via rbridge interface ve 100 with the ip address 192.168.100.10/24, set your vcs virtual ip address as a /32 address in this subnet by using the command **vcs virtual ip address 192.168.100.1/32**. To display the currently configured virtual IP address, use the **show vcs virtual ip** command:

```
switch# show vcs virtual ip

Virtual IP           : 10.21.87.2/20
Associated rbridge-id : 2
```

To remove the currently configured virtual IP address, use the **no vcs virtual ip address** command.

```
switch(config)# no vcs virtual ip address
switch# show running-config vcs virtual ip address
% No entries found.
```

NOTE

You should not use the **no vcs virtual ip address** command when logged onto the switch through the virtual IP address. Use the management port IP address of the principal switch, or the serial console connection of the principal switch.

If you wish to rebind this virtual IP address to this management interface, remove the currently configured virtual IP address and reconfigure it. This situation can arise when the virtual IP address is not bound to management interface of the principal switch as a result of duplicate address detection.

A separate gateway cannot be configured for virtual IP address. The default gateway is the same as the gateway address for the management port of the same switch.

Virtual IP address configuration scenarios

Virtual IP address may be assigned to a switch whenever it is the principal switch in the cluster. The configuration scenarios that may occur are described below.

TABLE 17 Virtual IP address configuration scenarios

Scenario	Description
First time cluster formation	When the cluster is first being formed, and if the virtual IP address is already configured, the principal switch is assigned the Virtual IP address. If no Virtual IP configuration exists, then the principal switch can be access using the management port IP address.
Virtual IP configuration	When you configure the virtual IP address for a cluster the first time, the address is bound to the management interface of the principal switch.
Principal switch failover	If the principal switch becomes a secondary switch while the virtual IP address is assigned to its management interface, then the virtual IP address is reassigned to the new principal switch.
Principal switch goes down	When the principal switch in the cluster goes down, the virtual IP address is released from its management interface. The virtual IP address will be assigned to the next switch that becomes the principal switch.
Principal switch chassis is disabled	When the chassis disable command is executed on the principal switch, the virtual IP address is released from its management interface. The virtual IP address will be assigned to the next switch that becomes the principal switch.
Virtual IP removal	If you remove the virtual IP address from the configuration, then the address is unbound from management interface of the principal switch. In

TABLE 17 Virtual IP address configuration scenarios (continued)

Scenario	Description
	this case, the principal switch can still be accessed by using the management port's IP address.
Trivial merge	In the event that two clusters merge together, the global configuration of the smaller cluster (Cluster A) is overwritten by the larger cluster (Cluster B). During this time, the virtual IP address is unbound from the principal switch of Cluster A. The virtual IP address of Cluster B can now be used to access the principal of new merged cluster. If the virtual IP address of Cluster B is not configured, there will not be a virtual IP address in the merged cluster.
Cluster reboot	When the cluster reboots, the virtual IP address is persistent and is bound to the new principal switch.
Cluster Islanding	If the ISL link goes down between two or more clusters that are forming, the principal switch in the original cluster retains the virtual IP address. The new principal switch in the second cluster will perform a check to confirm that the virtual IP address is not in use. If it is in use, then the address is not assigned to the switch and an error is logged in RASLog.
Virtual MAC address	Virtual MAC address are not supported by virtual IP addresses.
Management port primary IPv4 address	For a virtual IP address to work correctly, the management port's IPv4 address should be assigned and functional.

Configuring fabric ECMP load balancing

Traffic towards ECMP paths are load-balanced using the following eight fields as the Key: VlanID, MAC DA/SA, L3_ULP, L3 DA/SA, L4 Dst/Src. For some pattern of streams, most of the traffic falls into one ECMP path, and rest of the ECMP paths are underutilized. This results in loss of data traffic, even though more ECMP paths are available to offload the traffic. You can configure the ECMP path selection method within the fabric by using the **fabric ecmp load-balance** command. The operands for this command are listed and described in the following table.

TABLE 18 ECMP load-balancing operands

Operand	Description
dst-mac-vid	Destination MAC address and VID-based load balancing
src-dst-ip	Source and Destination IP address-based load balancing
src-dst-ip-mac-vid	Source and Destination IP and MAC address and VID-based load balancing
src-dst-ip-mac-vid-port	Source and Destination IP, MAC address, VID and TCP/UDP port based load balancing
src-dst-ip-port	Source and Destination IP and TCP/UDP port-based load balancing
src-dst-mac-vid	Source and Destination MAC address and VID-based load balancing
src-mac-vid	Source MAC address and VID-based load balancing

Additionally, you can choose to swap adjacent bits of the hash key using the **fabric ecmp load-balance-hash-swap** command. This is useful in cases where a choice of any of the hash key combinations causes the distribution of traffic to not be uniform.

The **fabric ecmp load-balance-hash-swap** command is used to configure the swapping of the input fields before feeding them to the hash function. The integer is interpreted as a bitwise control of the 212-bit key. Each bit controls whether the two adjacent bits of the key are to be swapped. This 32-bit control value is written to all four hash swap control registers. This value is replicated in 32-bit block to form a 106-bit value. A value of 0x0 does not swap any input fields while a value of 0xffffffff swaps all 106 input bit-pairs.

To configure the ECMP load-balancing feature, perform the following steps in global configuration mode.

1. Enter RBridge ID configuration mode.

```
switch(config)# rbridge-id 2
switch(config-rbridge-id-2)#
```

2. Execute the **fabric ecmp load-balance** command for the stream you want to favor.

This example uses the Destination MAC address and VID-based load balancing flavor.

```
switch(config-rbridge-id-2)# fabric ecmp load-balance dst-mac-vid
```

3. Optional: Use the **fabric ecmp load-balance-hash-swap** command to swap the input fields before feeding them to the hash function.

```
switch(config-rbridge-id-2)# fabric ecmp load-balance-hash-swap 0x4
```

4. Use the **show fabric ecmp load-balance** command to display the current configuration of hash field selection and hash swap.

```
switch# show fabric ecmp load-balance
Fabric Ecmp Load Balance Information
-----
Rbridge-Id          : 2
Ecmp-Load-Balance Flavor : Destination MAC address and VID based load balancing
Ecmp-Load-Balance HashSwap : 0x4
```


Configuring Metro VCS

- [Metro VCS overview.....](#) 113
- [Configuring a long-distance ISL.....](#) 120
- [Configuring interconnected Ethernet Fabrics.....](#) 121

Metro VCS overview

Metro VCS allows you to interconnect different locations and form clusters of data centers over long distance in order to provide disaster protection/recovery and load sharing.

In cases where distances are moderate, within 30km, and where either dedicated fiber or transparent wavelength services are available, the Metro VCS approach is a good and cost-effective Layer 2 interconnect solution. Due to the multi-pathing capabilities of the TRILL-based Metro VCS solution there is no issue with topology loops between multiple DC locations.

For longer distances alternative solutions are available.

- **Interconnecting separate fabrics through Layer 2 point-to-point connectivity** -- Layer 2 point-to-point connectivity is used to interconnect VCS Fabrics using their Edge-Ports. If more than one Layer 2 link is needed for capacity or for redundancy reasons, link aggregation (LAG/vLAG) can be used in order to avoid loops between the VCS Fabric edge ports and allow for active/active protection.
- **Interconnecting separate fabrics through Layer 2-VPN connectivity** -- VCS Fabrics are interconnected through their edge ports using Layer 2-VPN connectivity. This can be implemented with Layer 2-VPN services from Connectivity Service Providers or using VPLS functionality implemented on Brocade MLX routers.

FIGURE 19 Metro VCS configuration example



Both options are distance independent in relation to the speed of the protocols used (for example, LAG/vLAG is a slow protocol) and provide flat Layer 2 interconnection between multiple locations. In the case of more than two DC locations, care needs to be taken in order to avoid any topology loops.

Metro VCS details and configuration

Metro VCS allows for interconnection of different locations and allows to form clusters of data centers over metro distances (<10/30 km) in order to provide disaster protection/recovery and application load sharing.


By using Inter-Switch Links (ISLs) over longer distances (more than the standard 1000 m), Ethernet fabrics can be distributed across data centers located in geographically different locations.

FIGURE 20 Basic Metro VCS configuration



If more complex setups are needed within different locations, then local VCS Fabrics can be used interconnected via a stretched interconnect fabrics.

FIGURE 21 Complex Metro VCS configuration



Standard VCS Fabrics scaling limitations apply if Inter-Switch Links (ISLs) are used over distances of up to 1000 m. Using Inter-Switch Links (ISLs) over longer distances (more than the standard 1000 m) is currently available for 10G, 40G, and 100G ISLs and can be done in two ways:

- **Configuring long-distance ISLs (LD-ISLs)** -- This is supported only on 10G ISLs does not restrict fabric topologies (such as the numbers of nodes and number of locations) beyond the standard VCS Fabric scalability. LD-ISLs can be used over distances up to 10km if lossless services (DCB, FCoE) are needed and over distances of 30km if only standard Ethernet is needed.
- **Using standard ISLs over longer distances** -- This only works for restricted topologies (a maximum of 6 nodes and 3 locations) and for standard Ethernet (lossless Ethernet capabilities, such as no FCoE or lossless iSCSI, cannot be used in this case). The standard 10G ISL can be used over distances up to 30 km, and standard 40G and 100G ISLs can support distances up to 10 km.

Metro VCS using long-distance ISLs

Extending Ethernet Fabrics over distance is accomplished by using long-distance ISLs. The buffer allocation within a single port group is optimized, which extends the supported ISL distance.

Metro VCS supports long-distance ISL ports up to 30 km on the Brocade VDX platforms listed in the following table. Links up to 10 km are lossless.

TABLE 19 Limitations for long-distance Metro VCS

Supported hardware	Extended ISL up to 2 km	Extended ISL up to 5 km	Extended ISL up to 10 km	Extended ISL up to 30 km
Brocade VDX 6740	yes	yes	yes	yes
Brocade VDX 8770 - LC48x10G linecard	yes	yes	yes	yes

The following table lists the conditions on extended ISLs for Network OS hardware.

TABLE 20 Conditions for long-distance Metro VCS

Condition	Extended ISL up to 2 km	Extended ISL up to 5 km	Extended ISL up to 10 km	Extended ISL up to 30 km
Support for lossless FCoE/iSCSI traffic on the Metro VCS port group	yes	yes	yes	no
Layer 2/IP lossy traffic support	yes	yes	yes	yes
Number of Metro VCS long-distance ports supported per port group	1	1	1	1
Number of regular ISLs supported on a port group configured for long distance	1	1	0	0
Trunking support between multiple long-distance ISLs	no	no	no	no
CEE map or FCoE port allowed in same port group	no	no	no	no
eNS Sync (MAC address table sync)	yes	yes	yes	yes
Zoning	yes	yes	yes	yes
HA failover	yes	yes	yes	yes
Node redundancy check	yes	yes	yes	yes
vMotion	yes	yes	yes	yes
Maximum PFCs supported	3 (2 on the Brocade VDX 6740)	3(2 on the Brocade VDX 6740)	3 (2 on the Brocade VDX 6740)	3 (2 on the Brocade VDX 6740)
Long-distance ISL on 40G to 4x10G breakout interfaces	no	no	no	no
Long-distance ISL on 1G and 10G copper interfaces	no	no	no	no

The following table lists the port groups and number of port groups available on each platform for long-distance Metro VCS.

TABLE 21 Long-distance Metro VCS port-group schema

Platform	Port groups	Number of port groups on platform
Brocade VDX 6740	1-32, 33-48 (49-52 are 40G ports and do not support long distance)	2 *
Brocade VDX 8770 (LC48x10G linecard)	1-8, 9-16, 17-24, 25-32, 33-40, 41-48	6 per 10GbE blade

* Not a valid deployment scenario at distances longer than 5 km, as no normal ISLs are allowed if both port groups are configured with long-distance ISLs for 10 km and 30 km.

Guidelines and restrictions for long-distance Metro VCS

Consider the following guidelines and restrictions when configuring long-distance Metro VCS:

- Long-distance-ISLs are only supported on 10G interface links.
- Only one long-distance-ISL is supported within a single port group.
- Long-distance-ISL is not supported on 10G copper RJ-45 interfaces
- Long-distance-ISL is not supported on 40G-to-10G breakout interfaces
- Brocade trunking is not supported with long-distance ISLs, but up to eight 8-link ECMP trunks can be used.
- Edge ports within the same port group where a long-distance ISL is configured cannot be configured with DCB maps.
- Edge ports within the same port group where a long-distance ISL is configured cannot be configured by means of the **fcoeport default** command.
- A maximum of three PFCs can be supported on a Metro VCS configured platform. By default, PFC is enabled by class 3 and 7.
- The Brocade VDX 6740 switches support only two PFCs.
- All the ports in the port group are rebooted when a port is configured for long distance.
- For 2-, 5-, 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.
- A port group containing a long-distance port cannot have a CEE map configuration on any edge port.
- For 2-km and 5-km long-distance ISLs, one additional standard ISL connection is supported on the same long-distance port group.
- Lossless FCoE traffic is supported up to 10 km with long-distance ISL configured.
- Lossy Layer 2/Layer 3 traffic is supported up to 30 km with long-distance ISL configured.

Metro VCS using standard-distance ISLs

In order to deploy Metro VCS using standard-distance ISLs, no configuration is required on the ISL. The default configuration on the 10-, 40-, and 100-Gbps interface by means of the **fabric isl enable** and **fabric trunk enable** commands allows ISL formation with other Brocade VDX switches in the same VCS cluster automatically. BLDP negotiation takes place to form ISLs for distances up to 30 km for 10G and 10 km for 40G and 100G interfaces. (Refer to [Configuring Brocade VCS Fabrics](#) on page 101.)


Metro VCS using standard ISLs is supported on the following platforms:

- Brocade VDX 2740
- Brocade VDX 6740, VDX 6740T, and VDX 6740T-1G
- Brocade VDX 8770 with the LC48X10G line card
- Brocade VDX 8770 with the LC27X40G line card
- Brocade VDX 8770 with the LC12X40G line card
- Brocade VDX 8770 with the LC6X100G line card

The following figure is an example deployment topology supported for interconnecting data centers by extending Brocade VCS Ethernet fabrics using standard-distance ISLs. The local VCS clusters are connected to the Metro VCS clusters by Brocade vLAGs. In this case,

local data-center Ethernet fabrics from both site are not merged while providing seamless Layer 2 extension. For Metro VCS, Brocade standard-distance ISL trunking is supported, with up to a maximum of eight ISLs to form 80G trunks.

FIGURE 22 Typical deployment topology for Metro VCS using standard-distance ISLs



The following table lists the port groups and number of port groups available on each platform for Metro VCS using standard-distance ISLs.

TABLE 22 Standard Metro VCS port-group schema

Platform	Port groups	Number of port groups on platform
Brocade VDX 6740	1-16, 17-32, 33-40, 41-48, 49-50, 51-52	6
Brocade VDX 6740T	49-50, 51-52 (40G interfaces only)	2
Brocade VDX 6740T-1G	49-50, 51-52 (40G interfaces only)	2
Brocade VDX 8770 (LC6X100G)	1-2, 3-4, 5-6	3

TABLE 22 Standard Metro VCS port-group schema (continued)

Platform	Port groups	Number of port groups on platform
Brocade VDX 8770 (LC27X40G)	1-3, 4-6, 7-9, 10-12, 13-15, 16-18, 19-21, 22-24, 25-27	9
Brocade VDX 8770 (LC12X40G)	1-2, 3-4, 5-6, 7-8, 9-10, 11-12	6
Brocade VDX 8770 (LC48x10G line card)	1-8, 9-16, 17-24, 25-32, 33-40, 41-48	6

Guidelines and restrictions for standard-distance Metro VCS

Consider the following guidelines and restrictions when configuring Metro VCS with standard-distance ISLs:

- Only two data-center and three data-center topologies are supported.
- A maximum of two nodes are supported for each site, which provides node redundancy. If more-complex local designs are required, a local VCS sub-fabric design must be used.
- Only standard Ethernet services are supported, lossless Ethernet capabilities, such as no FCoE or lossless iSCSI, cannot be used in this case.
- Brocade trunking with up to 80G (8x10G) or (2x40G) links are supported as follows:
 - VDX 8770: 8x10G
 - VDX 6740: 8x10G and 2x40G
 - VDX 6740T: 2x40G
 - VDX 6740T-1G: 2x40G

Metro VCS combined with vLAGs

Outside of Metro distances, and whenever bit transparency may be a problem, edge-to-edge interconnected fabrics using 10G, 40G, or 100G vLAG over multiple standard Ethernet links can be used. This allows you to connect separate Ethernet fabrics that can be located in different data centers, even if the distance between those locations up to 30 km for 10G interfaces and up to 10 km for 40G and 100G interfaces.


Metro VCS over a long-distance fabric

Whenever standard VCS Fabrics are interconnected through LAG/vLAG, they are not limited beyond LAG/vLAG protocol capabilities.

Metro VCS Fabrics (stretched fabrics) that are interconnected with standard fabrics are supported for distances up to 100 km.

As shown in the following figure, one side can be a Metro VCS over a long-distance fabric.

FIGURE 23 Metro VCS and distributed Ethernet fabrics



In all deployment with interconnects using edge ports, lossless Ethernet traffic (DCB/FCoE) is not supported.

In order to connect two distinct VCS Ethernet fabrics between data centers, a third Metro VCS fabric can be formed, and the distinct local VCS Ethernet fabrics can connect to the Metro VCS fabric by means of Virtual Link Aggregation (vLAG).

Alternatively, the distinct VCS Ethernet fabrics in the respective data centers can be directly connected to each other by means of vLAG over xWDM up to a distance of 10 or 30 km. Wherever bit-transparency is not achievable in xWDM equipment, this solution can be successfully deployed for edge-to-edge interconnectivity (using 10G, 40G, or 100G vLAGs over multiple standard Ethernet links). This deployment is referred to as "Distributed Ethernet Fabrics using vLAG."

This implementation eliminates the need for the creation of a separate Metro VCS fabric to achieve local VCS cluster isolation while providing Layer 2 connectivity. In such a deployment, DCB/FCoE lossless Ethernet traffic is not supported.

NOTE

When a port-channel from a node in one VCS spans across multiple R Bridges in other VCS cluster, a vLAG is formed on the R Bridges in the VCS cluster that are part of the same port-channel. For Distributed Ethernet Fabrics using vLAG over long distances, only LACP-based standard port-channels are supported. For details on how to create port-channels and vLAGs, refer to "Configuring Link Aggregation" chapter of the *Network OS Layer 2 Switching Configuration Guide*.

Supported platforms for Distributed Ethernet Fabrics using vLAG

The following VDX platforms are supported for Distributed Ethernet Fabrics using vLAG:

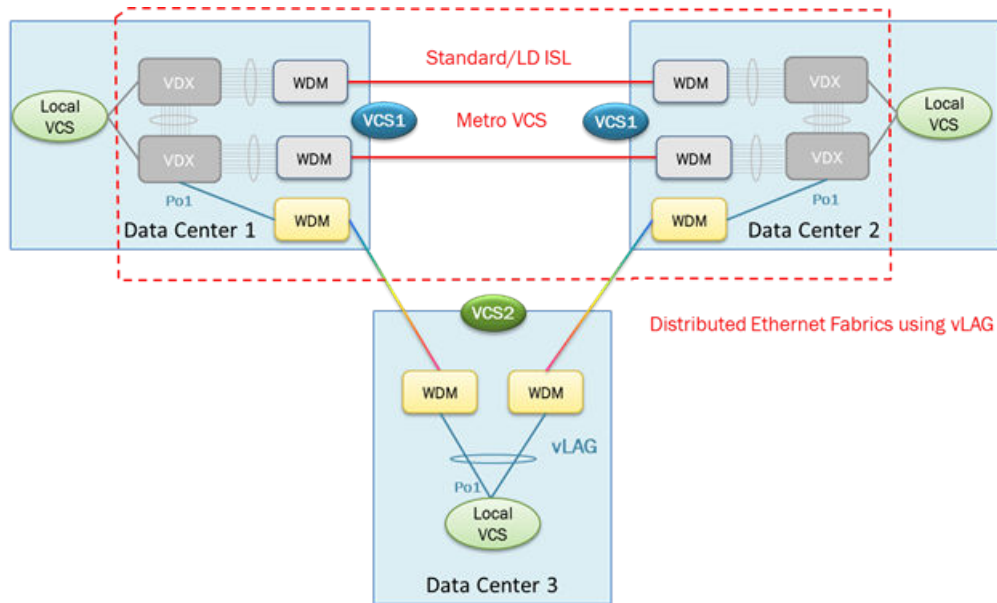
- Brocade VDX 6740, 6740T, and 6740T-1G
- Brocade VDX 8770 with the LC48X10G line card
- Brocade VDX 8770 with the LC27X40G line card
- Brocade VDX 8770 with the LC12X40G line card
- Brocade VDX 8770 with the LC6X100G line card

Topology for Distributed Ethernet Fabrics using vLAG

When a port-channel from a node in one VCS spans across multiple R Bridges in other VCS cluster, a vLAG is formed on the R Bridges in the VCS cluster that are part of the same port-channel. For Distributed Ethernet Fabrics using vLAG over long distances, only LACP-based standard port-channels are supported. For details on how to create port-channels and vLAGs, refer to *Configuring Link Aggregation* in the *Layer 2 Switching Configuration Guide*.

The following figure is a typical deployment topology that uses distributed Ethernet Fabrics using vLAG to interconnect data centers. Nodes from the local VCS cluster are connected by means of xWDM to the nodes in a distant VCS clusters to form a vLAG in between. The distant VCS cluster can be a standard VCS cluster or could be spanned across two data centers over standard-distance or long-distance ISLs, as shown in the figure below. In this case, the vLAG between the two data centers provides VCS fabric isolation while providing seamless Layer 2 connectivity.

FIGURE 24 Connecting local VCS clusters over long distance using vLAG



Guidelines and restrictions for Distributed Ethernet Fabrics using vLAG

Note the following guidelines and restrictions for Distributed Ethernet Fabrics using vLAG.

- Only dynamic vLAG is supported.
- DCB/FCoE lossless Ethernet traffic is not supported.
- The maximum distance between standard VCS Fabrics is not limited beyond the capabilities of the LAG/vLAG protocol.
- The maximum supported distance between a stretched Fabric (Metro VCS) and an additional standard VCS Fabric connected through a vLAG is limited to 100 km.

Configuring a long-distance ISL

To configure a long-distance ISL, perform the following steps in privileged EXEC mode. Each long-distance ISL port of a VCS must be connected to a long-distance ISL port on the remote VCS.

1. Verify that the default standard-distance ISL configuration is correct by using the **show running-config** command.

```
switch# show running-config interface tengigabitethernet 51/0/1
interface TenGigabitEthernet 51/0/1
fabric isl enable
fabric trunk enable
no shutdown
```


- Set the port to support Metro VCS up to 30 km by using the **long-distance-isl** command.

```
switch# interface tengigabit 51/0/1
switch(conf-if-te-51/0/1)# long-distance-isl 30000
```

- Perform the same long-distance ISL configuration on the interface of the peer RBridge on the remote sites of the Metro VCS.
- Verify that the long-distance ISL is correctly formed by using the **show fabric isl** and **show fabric islports** command.

```
switch(conf-if-te-51/0/1)# do show fabric isl
Rbridge-id: 51 #ISLs: 1
  Src      Src      Nbr      Nbr
  Index   Interface  Index   Interface          Nbr-WWN          BW  Trunk  Nbr-Name
-----
  4       Te 51/0/1    4       Te 53/0/1          10:00:00:05:33:65:3B:50  10G Yes  "VCS3-53"
switch(conf-if-te-51/0/1)# do show fabric islports
Name:      VCS3-51
Type:      131.4
State:     Online
Role:      Fabric Principal
VCS Id:    3
Config Mode:Local-Only
Rbridge-id: 51
WWN:      10:00:00:05:33:e5:d0:4b
FCF MAC:   00:05:33:e5:d0:cf
  Index   Interface  State  Operational State
=====
  0       Fo 51/0/49  Down
  1       Fo 51/0/50  Down
  2       Fo 51/0/51  Down
  3       Fo 51/0/52  Down
  4       Te 51/0/1   Up     ISL 10:00:00:05:33:65:3B:50 "VCS3-53" (Trunk Primary)
<Truncated>
```

- Use the **show ip interface brief** command to confirm the configuration, making sure that Status is "up" and Protocol is "LD ISL," as in the following example output.

```
switch# show ip interface brief

Interface                               IP_Address      VRF              Status           Protocol
-----                               -
TenGigabitEthernet 51/0/1              unassigned      default-vrf      up               up (LD ISL)
```

Configuring interconnected Ethernet Fabrics

To deploy interconnected Ethernet Fabrics using vLAG, create a port-channel interface on the R Bridges that are to be connected. Then add the member interfaces to the port-channel and bring them online. Configure switchport and add the VLANs that are to be allowed over the port-channel. After the port-channels on all the R Bridges are online, the vLAG forms automatically on the R Bridge that connects to multiple nodes on the other VCS cluster. In the deployment topology shown in [Topology for Distributed Ethernet Fabrics using vLAG](#) on page 119, the vLAG forms on the R Bridges that are part of port-channel Po1 in Data-Centers 1 and 2 that forms VCS 1.

This configuration must be applied on RBridges that connect the two VCS instances. Perform the following task in global configuration mode.

1. Create a port-channel interface on all RBridges that are directly connected to RBridges in other VCS instances.

NOTE

In logical chassis cluster mode, the port-channel is created only from the principal node and is applied globally.

```
switch(config)# interface port-channel 300
```

2. Verify that the port-channel is created correctly by using the **show running-config** command.

```
switch(config-Port-channel-300)# do show running-config interface port-channel 300

interface Port-channel 300
 vlag ignore-split
 shutdown
```

3. Configure the port-channel interface for the switchport trunk and add the VLANs to be allowed on the trunk interface by using the **switchport** command.

```
switch(config-Port-channel-300)# switchport
switch(config-Port-channel-300)# switchport mode trunk
switch(config-Port-channel-300)# switchport trunk allowed vlan all
```

4. Verify the port-channel interface configuration by using the **show running-config** command.

```
switch(config-Port-channel-300)# do show running-config interface Port-channel 300

interface Port-channel 300
 vlag ignore-split
 switchport
 switchport mode trunk
 switchport trunk allowed vlan all
 switchport trunk tag native-vlan
 spanning-tree shutdown
 shutdown
```

5. Add member interfaces to the port-channel interface by using the **channel-group** command. Do this for all interfaces that must be part of the port-channel.

```
switch(conf-if-te-53/0/31)# channel-group 300 mode active type standard
switch(conf-if-te-53/0/31)# do show running-config interface

TenGigabitEthernet 53/0/31
interface TenGigabitEthernet 53/0/31
 fabric isl enable
 fabric trunk enable
 channel-group 300 mode active type standard
 lacp timeout long
 no shutdown
```

6. Bring the port-channel online in both VCS instances by executing **no shutdown** on the port-channel interface.

```
switch(config-Port-channel-300)# no shutdown
```

7. Verify the port-channel interface configuration by using the **show running-config** command.

```
switch(config-Port-channel-300)# do show running-config interface Port-channel 300

interface Port-channel 300
 vlag ignore-split
 switchport
 switchport mode trunk
 switchport trunk allowed vlan all
 switchport trunk tag native-vlan
 spanning-tree shutdown
 no shutdown
```

8. Verify that console RASLogs indicate the formation of the vLAG by using the **no shutdown** command.

```
switch(config-Port-channel-300)# no shutdown

2013/06/17-16:40:53, [NSM-1023], 224126, DCE, INFO, VCS1-51, RBridge ID 51 has joined Port-channel
300. Port-channel is a vLAG with RBridge IDs 52 51.
```

9. Verify the formation of the port-channel vLAG by using the **show port-channel** command.

```
switch# show port-channel 300
LACP Aggregator: Po 300 (vLAG)
Aggregator type: Standard
Ignore-split is enabled
Member rbridges:
  rbridge-id: 51 (2)
  rbridge-id: 52 (2)
Admin Key: 0010 - Oper Key 0010
Partner System ID - 0x8000,01-e0-52-00-00-02
Partner Oper Key 0010
Member ports on rbridge-id 51:
  Link: Te 51/0/31 (0x19180E801C) sync: 1
  Link: Te 51/0/32 (0x19180F001D) sync: 1
```


Administering Zones

- [Zoning overview.....](#) 125
- [Configuring and managing zones](#) 132

Zoning overview

Zoning is a fabric-based service that enables you to partition your network into logical groups of devices that can access each other and prevent access from outside the group. Grouping devices into zones in this manner not only provides security, but also relieves the network from Registered State Change Notification (RSCN) storms that occur when too many native FCoE devices attempt to communicate with one another.

You can use zoning to partition your network in many ways. For example, you can partition your network into two zones, *winzone* and *unixzone*, so that your Windows servers and storage do not interact with your UNIX servers and storage. You can use zones to logically consolidate equipment for efficiency or to facilitate time-sensitive functions; for example, you can create a temporary zone to back up nonmember devices.

A device in a zone can communicate only with other devices connected to the fabric within the same zone. A device not included in the zone is not available to members of that zone. When zoning is enabled, devices that are not included in *any* zone configuration are inaccessible to all other devices in the fabric.


Zones can be configured dynamically. They can vary in size, depending on the number of fabric-connected devices, and devices can belong to more than one zone.

Example zoning topology

Consider the following figure, which shows three configured zones: Red, Green, and Blue. In this figure the following is true:


- Server 1 can communicate only with the Storage 1 device.
- Server 2 can communicate only with the RAID and Storage 2 devices.
- Server 3 can communicate with the RAID and Storage 1 devices.
- The Storage 3 is not assigned to a zone; no other zoned fabric device can access it.

FIGURE 25 Zoning



Connecting to another network through a Fibre Channel (FC) router, you can create a Logical SAN (LSAN) zone to include zone objects on other fabrics, including Fabric OS networks. No merging takes place across the FC router when you create an LSAN zone. The figure below shows an example in which Server 1, which is connected to switch in a Brocade VCS Fabric cluster, has access to local storage and to RAID storage on a Fabric OS fabric. (For a detailed discussion of LSAN zones, refer to [LSAN zones](#) on page 127.)

FIGURE 26 LSAN zoning

**NOTE**

Zoning in Network OS 4.0.0 and later has the following restrictions:

- Zone objects based on physical port number or port ID (D,I ports) are not supported.
- You cannot access a target on a Network OS fabric from a server on the Fabric OS fabric.

LSAN zones

LSAN zones are distinct from conventional zones. This section details how to define and manage LSAN zones and provides recommendations about LSAN zone naming.

LSAN zones overview

A Logical SAN (LSAN) consists of zones in two or more edge or backbone fabrics that contain the same devices. LSANs essentially provide selective device connectivity between fabrics without forcing you to merge those fabrics. FC routers provide multiple mechanisms to manage inter-fabric device connectivity through extensions to existing switch management interfaces. For details of this FC-FC routing service, refer to the *Fabric OS Administrator's Guide*.

NOTE

A backbone fabric consists of one or more FC switches with configured EX_Ports. These EX_Ports in the backbone connect to edge fabric switches through E_Ports. This type of EX_Port-to-E_Port connectivity is called an "Inter-Fabric Link (IFL)".

The Brocade VCS Fabric connection to the FC router is an ISL that connects an FC port on a Brocade VDX 6740 to an EX_Port on the FC router. Similarly, an FC port on the Fabric OS fabric connects to an EX_Port on the FC router.

You can define and manage LSANs using the same zone management tools as for regular zones. The FC router makes LSAN zoning possible by importing devices in effective zones. For example, consider two devices:

- 11:22:33:44:55:66:77:99 is connected to a switch in a Brocade VCS Fabric cluster.
- 11:22:33:44:55:66:77:88 is connected to a switch in a Fabric OS fabric.

The FC-FC routing service on the FC router that connects the two fabrics presents 11:22:33:44:55:66:77:88 as a phantom device to the Brocade VCS Fabric and also presents 11:22:33:44:55:66:77:99 as a phantom device to the Fabric OS fabric. You can then use the regular zone management tools on the Brocade VCS Fabric cluster to incorporate 11:22:33:44:55:66:77:99 into an LSAN zone on the Brocade VCS Fabric. Similarly, you can use the regular zone management tools in Fabric OS to incorporate 11:22:33:44:55:66:77:88 into an LSAN zone in the Fabric OS fabric. Once both the Brocade VCS Fabric zone and the Fabric OS zone are enabled, the FC router imports devices common to both zones and makes them available to the zones in each fabric.

LSAN naming

Zones that contain hosts and targets that are shared between the two fabrics need to be explicitly coordinated. To share devices between any two fabrics, you must create an LSAN zone in both fabrics containing the WWNs of the devices to be shared. Although you manage an LSAN zone by using the same tools as any other zone on the edge fabric, two behaviors distinguish an LSAN zone from a conventional zone:

- A required naming convention. The name of an LSAN zone begins with the prefix "LSAN_". The LSAN name is case-insensitive; for example, `lsan_` is equivalent to `LSAN_`, `Lsan_`, and so on.
- LSAN zone members in all fabrics must be identified by their WWN. You cannot use the port IDs that are supported only in Fabric OS fabrics.

NOTE

The "LSAN_" prefix must appear at the beginning of the zone name.

To enable device sharing across multiple fabrics, you must create LSAN zones on the edge fabrics (and optionally on the backbone fabric as well), using normal zoning operations to create zones with names that begin with the special prefix "LSAN_", and adding host and target port WWNs from both local and remote fabrics to each local zone as desired. Zones on the backbone and on multiple edge fabrics that share a common set of devices will be recognized as constituting a single multi-fabric LSAN zone, and the devices that they have in common will be able to communicate with each other across fabric boundaries.

Managing domain IDs

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

It is important to ensure that front domain IDs and translate domain IDs presented by the FC router do not overlap routing bridge IDs in the FCoE fabric; otherwise, the connectivity will fail and the Network OS switch with the overlapping routing bridge ID becomes isolated from the fabric. To prevent potential overlap, use the `portCfgExport -d` Fabric OS command on the FC router to apply a unique front domain ID — one that will not be used in the FCoE fabric. Similarly, use the `fcrXlateConfig importedFID exportedFID preferredDomainID` Fabric OS command to set the translate domain ID to a unique value that is also not used as a routing bridge ID.

Refer to the *Fabric OS Command Reference Manual* for details about the `portCfgExport` and `fcrXlateConfig` commands.

Approaches to zoning

The following table lists the various approaches you can take when implementing zoning in a Network OS fabric.

TABLE 23 Approaches to fabric-based zoning

Zoning approach	Description
Recommended approach	
Single HBA	Zoning by single HBA most closely re-creates the original SCSI bus. Each zone created has only one HBA (initiator) in the zone; each of the target devices is added to the zone. Typically, a zone is created for the HBA and the disk storage ports are added. If the HBA also accesses tape devices, a second zone is created with the HBA and associated tape devices in it. In the case of clustered systems, it could be appropriate to have an HBA from each of the cluster members included in the zone; this zoning is equivalent to having a shared SCSI bus between the cluster members and assumes that the clustering software can manage access to the shared devices. In a large fabric, zoning by single HBA requires the creation of possibly hundreds of zones; however, each zone contains only a few members. Zone changes affect the smallest possible number of devices, minimizing the impact of an incorrect zone change. <i>This zoning philosophy is the preferred method.</i>
Alternative approaches	
Application	Zoning by application typically requires zoning multiple, perhaps incompatible, operating systems into the same zones. This method of zoning creates the possibility that a minor server in the application suite could disrupt a major server (such as a Web server disrupting a data warehouse server). Zoning by application can also result in a zone with a large number of members, meaning that more notifications, such as RSCNs, or errors, go out to a larger group than necessary.
Operating system	Zoning by operating system has issues similar to zoning by application. In a large site, this type of zone can become very large and complex. When zone changes are made, they typically involve applications rather than a particular server type. If members of different operating system clusters can detect storage assigned to another cluster, they might attempt to own the other cluster's storage and compromise the stability of the clusters.
Port allocation	Avoid zoning by port allocation unless the administration team has very rigidly enforced processes for port and device allocation in the fabric. It does, however, provide some positive features. For instance, when a storage port, server HBA, or tape drive is replaced, the change of WWN for the new device is of no consequence. As long as the new device is connected to the original port, it continues to have the same access rights. The ports on the edge switches can be pre-associated to storage ports, and control of the fan-in ratio (the ratio of the input port to output port) can be established. With this pre-assigning technique, the administrative team cannot overload any one storage port by associating too many servers with it.
Not recommended	
No zoning	Using no zoning is the least desirable zoning option because it allows devices to have unrestricted access on the fabric and causes RSCN storms. Additionally, any device attached to the fabric, intentionally or maliciously, likewise has unrestricted access to the fabric. This form of zoning should be used only in a small and tightly controlled environment, such as when host-based zoning or LUN masking is deployed.

Zone objects

A zone object can be one of the following types: a zone, a zone member, an alias for one or more zone members, or a zoning configuration.

Zones

A zone is made up of one or more zone members. Each zone member can be a device, a port, or an alias. If the zone member is a device, it must be identified by its Node World Wide Name (node WWN). If it is a port, it must be identified by its Port World Wide Name (port WWN). Port WWNs and node WWNs can be mixed in the same zone. For LSAN zones, only port WWNs can be used.

World Wide Names are specified as 8-byte (16-digit) hexadecimal numbers, separated by colons (:) for example, 10:00:00:90:69:00:00:8a. When a zone object is the node WWN, only the specified device is in the zone. When a zone object is the port WWN name, only the single port is in the zone.

NOTE

You are not restricted from configuring more than 255 zone members. However, that figure is considered a best-practices limit, and exceeding it can lead to unpredictable results.

Zone aliases

A zone alias is a name assigned to a device or a group of devices. By creating an alias, you can assign a familiar name to one or more devices and refer to these devices by that name. Aliases simplify cumbersome data entry by allowing you to create an intuitive naming structure (such as using "NT_Hosts" to define all NT hosts in the fabric).

As a shortcut for zone members, zone aliases simplify the entry and tracking of zone objects that are defined by their WWNs. For example, you can use the name "Eng" as an alias for "10:00:00:80:33:3f:aa:11".

Naming zones for the initiator they contain can also be useful. For example, if you use the alias SRV_MAILSERVER_SLT5 to designate a mail server in PCI slot 5, then the alias for the associated zone is ZNE_MAILSERVER_SLT5. This kind of naming strategy clearly identifies the server host bus adapter (HBA associated with the zone).

Zone configurations

A *zone configuration* is a group of one or more zones. A zone can be included in more than one zone configuration. When a zone configuration is enabled, all zones that are members of that configuration are enabled.

Several zone configurations can reside on a switch at once, and you can quickly alternate between them. For example, you might want to have one configuration enabled during the business hours and another enabled overnight. However, only one zone configuration can be enabled at a time.

The different types of zone configurations are:

- Defined Configuration

The complete set of all zone objects defined in the fabric.

- Enabled Configuration

A single zone configuration that is currently in effect. The enabled configuration is built when you enable a specified zone configuration.

If you disable the enabled configuration, zoning is disabled on the fabric, and default zoning takes effect. When default zoning takes effect, either all devices within the fabric can communicate with all other devices, or no device communicate with any other device, depending on how default zoning is configured. Disabling the configuration does not mean that the zone database is deleted, however, only that no configuration is active in the fabric.

On power-up, the switch automatically reloads the saved configuration. If a configuration was active when it was saved, the same configuration is reinstated on the local switch.

Naming conventions

Naming zones and zone configurations is flexible. You can devise prefixes to differentiate between zones used for production, backup, recovery, or testing. One configuration should be named PROD_*fabricname*, where *fabricname* is the name that the fabric has been assigned. The purpose of the PROD configuration is to easily identify the configuration that can be implemented and provide the most generic services. If you want to use other configurations for specific purposes, you can use names such as "BACKUP_A," "RECOVERY_2," and "TEST_18jun02".

Zoning enforcement

Zone enforcement is by name server. The name server filters queries and RSCNs based on the enabled zoning configuration.

Considerations for zoning architecture

This table lists considerations for zoning architecture.

TABLE 24 Considerations for zoning architecture

Item	Description
Effect of changes in a production fabric	Zone changes in a production fabric can result in a disruption of I/O under conditions when an RSCN is issued because of the zone change and the HBA is unable to process the RSCN fast enough. Although RSCNs are a normal part of a functioning SAN, the pause in I/O might not be acceptable. For these reasons, you should perform zone changes only when the resulting behavior is predictable and acceptable. Ensuring that the HBA drivers are current can shorten the response time in relation to the RSCN.
Allowing time to propagate changes	Zoning commands make changes that affect the entire fabric. When executing fabric-level configuration tasks, allow time for the changes to propagate across the fabric before executing any subsequent commands. For a large fabric, you should wait several minutes between commands.
Confirming operation	After changing or enabling a zone configuration, you should confirm that the nodes and storage can identify and access one another. Depending on the platform, you might need to reboot one or more nodes in the fabric with the new changes.
Use of aliases	The use of aliases is optional with zoning. Using aliases requires structure when defining zones. Aliases aid administrators of zoned fabrics in understanding the structure and context of zoning.

Operational considerations for zoning

Consider the following topics when configuring zoning.

Zoning configuration changes

When you save, enable, or disable a configuration, the changes are automatically distributed to all switches in the VCS Fabric.

Supported firmware for zoning

Zoning is supported only if all R Bridges in the fabric are running Network OS 2.1 or later.

Connecting an R Bridge running Network OS 2.0 to an R Bridge running Network OS 2.1 or later merges the two networks only if the R Bridge running Network OS 2.1 or later is in Brocade VCS Fabric mode and no zone database elements are defined or enabled.

A switch running Network OS v3.0.0 will segment if it is attached to a switch running Network OS v2.0.0 regardless of zoning configuration. A switch running Network OS v3.0.0 will join the fabric with a 2.1.x switch and zones will be merged, but the cluster will not form, so no further zoning commands will be allowed until all switches are upgraded to the same firmware version and the cluster has formed.

The Inter-Switch Links (ISLs) connecting the two R Bridges will segment if the R Bridge running Network OS 2.1 or later has any zone defined or enabled, or the default zone is set to No Access. Any such configuration requires automatic distribution of zoning configuration data, which is not compatible with R Bridges running Network OS 2.0.

Firmware downgrade and upgrade considerations for zoning

A firmware downgrade from Network OS 4.1.0 to Network OS 2.1.x is not permitted under the following conditions:

1. One or more zone aliases are configured on the switch. You must remove all references to zone aliases prior to a firmware downgrade. Use the **no zoning defined-configuration alias** command to delete all zone alias objects. Then issue the **zoning enabled-configuration cfg-action{cfg-save | cfg-disable}** command or the **zoning enabled-configuration cfg-name *cfg_name*** command to commit the operation before re-attempting a firmware download.

- An open zone transaction in progress. You must either commit or abort the current open transaction before re-attempting a firmware download. Use the **zoning enabled-configuration cfg-action {cfg-save | cfg-disable}** command or the **zoning enabled-configuration cfg-name cfg_name** command to commit the current open transaction. Alternately, use the **zoning enabled-configuration cfg-action cfg-transaction-abort** command to abort the open transaction.

You cannot downgrade any switch in a Brocade VCS Fabric to Network OS 2.0 or earlier if any zone definition exists in the defined configuration. Any attempt to do so will fail while attempting to download the earlier firmware. For the downgrade to succeed, you must clear the defined configuration, disable any active configuration, set the default zoning mode to *All Access*, and then try again to download the firmware.

When you upgrade from Network OS 2.1.0 to versions 2.1.1 or later, the zone database is cleared.



CAUTION

Clearing the defined configuration clears the zoning database for the entire fabric. If you want to downgrade just one switch without affecting the rest of the fabric, disconnect the switch from the fabric before deleting the defined configuration.

Configuring and managing zones

Zone configuration management overview

You can perform zoning operations on any RBridge in the VCS Fabric, but they are always executed on the principal RBridge. In Logical Chassis mode, any edits made to the zoning database are allowed only from the principal RBridge, and you can issue **show** commands from non-principal switches in this mode. In Fabric Cluster mode, you can make edits from any RBridge.

Automatic distribution of the zoning configuration ensures that the effects of these operations are shared and instantly visible on all switches in the VCS Fabric. However, these operations are not permanent until a transaction commit operation saves them to nonvolatile memory, which holds the master copy of the zoning database. In fabric cluster mode, any user can commit the transaction on any switch, and the commit operation saves the operations performed by all users. Once the zoning configuration is saved in permanent memory, it persists across reboot operations.

A transaction commit occurs when you or another user initiates any of the following zoning operations:

- Saving the database to nonvolatile memory with the **zoning enabled-configuration cfg-action cfg-save** command.
- Enable a specific zone configuration with the **zoning enabled-configuration cfg-name** command.
- Disabling the currently enabled zone configuration with the **no zoning enabled-configuration cfg-name** command.

Executing the **zoning enabled-configuration cfg-action cfg-transaction-abort** command cancels the currently open transaction.

If the principal RBridge reboots or goes down, Network OS selects a new principal and any pending zoning transaction is rolled back to the last committed transaction, which is the effective zoning configuration saved in nonvolatile memory. Any changes made to the effective configuration prior to an abort operation must be re-entered.

If an RBridge other than the principal reboots or goes down, the ongoing transaction is not backed out. Any zoning operations initiated by the RBridge are still part of the global transaction maintained on the principal RBridge.

If a fabric segments, the newly elected principal RBridge determines whether transaction data are retained. If a segment retains the original principal, it also retains ongoing transaction data. If a segment elects a new principal, the transaction is aborted.

The zone startup configuration is always equal to the running configuration. The running configuration will always be overwritten by the information from the master copy of the zoning database in nonvolatile memory at startup, so you always start up with the previous running configuration. It is not necessary to copy the running configuration to the startup configuration explicitly.

You can save a snapshot of the current running configuration using the **copy running-config file** command. You can add configuration entries from a saved configuration using the **copy file running-config** command. When saving the snapshot you must ensure that the saved running configuration contains no zoning transaction data, otherwise failures will occur when you attempt to restore configuration entries from the saved file. Any transaction data would cause such a failure, including empty configuration definitions or empty zones.

Notes

- When you re-enable the enabled-configuration (using the **zoning enabled-configuration** command) on the principal switch in the cluster, the system propagates the enabled-configuration across the cluster. There is a slight risk of doing this in that the defined-configuration may contain configuration edits that you may not want to enable yet. This feature prevents switches in the cluster from having mismatched enabled-configurations.
- When restoring the running configuration, Brocade recommends copying the file to the running configuration in the absence of any other command line input.
- When you restore a configuration using the **copy** command, the contents of the file are added to the defined configuration; they do not replace the defined configuration. The result is cumulative, is as if the input came from the command line.

Understanding and managing default zoning access modes

The default zoning mode controls device access if zoning is not implemented or if there is no enabled zone configuration. Default zoning has two access modes:

- *All Access* — All devices within the fabric can communicate with all other devices.
- *No Access* — Devices in the fabric cannot access any other device in the fabric.

The default setting is All Access. Changing the default access mode requires committing the ongoing transaction for the change to take effect.

The default zoning mode takes effect when you disable the effective zone configuration. If your default zone has a large number of devices, to prevent RSCN storms from overloading those devices, you should set the default zoning mode to No Access before attempting to disable the zone configuration. If your default zone includes more than 300 devices, the zoning software prevents you from disabling the zoning configuration if the default zoning mode is All Access.

Setting the default zoning mode

1. In privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter one of the following commands, depending on the default access mode you want to configure:
 - To set the default access mode to All Access, enter **zoning enabled-configuration default-zone-access allaccess**.
 - To set the default access mode to No Access, enter **zoning enabled-configuration default-zone-access noaccess**.
3. Enter the **zoning enabled-configuration cfg-action cfg-save** or **zoning enabled-configuration cfg-name** command to commit the ongoing transaction and save the access mode change to nonvolatile memory.

4. Enter the **show running-config zoning enabled-configuration** command to verify the access mode change.

Example of setting the default zoning mode to no access:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration default-zone-access noaccess
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)# do show running-config zoning enabled-configuration
zoning enabled-configuration cfg-name cfg1
zoning enabled-configuration default-zone-access noaccess
zoning enabled-configuration cfg-action cfg-save
```

Understanding and managing zone database size

The maximum size of a zone database is the upper limit for the defined configuration, and it is determined by the amount of memory available for storing the master copy of the defined configuration in flash memory.

Use the following information displayed by the **show zoning operation-info** command to determine whether there is enough space to complete outstanding transactions:

- db-max — Theoretical maximum size of the zoning database kept in nonvolatile memory
- db-avail — Theoretical amount of free space available
- db-committed — The size of the defined configuration currently stored in nonvolatile memory
- db-transaction — The amount of memory required to commit the current transaction

The supported maximum zone database size is 100 KB. If the outstanding transaction data (db-transaction field) is less than the remaining supported space (100 KB minus db-committed), enough space exists to commit the transaction.

NOTE

The db-max field has a theoretical zone database limit of approximately 1 MB. However, performance might become unacceptable if the zoning database exceeds 150 KB.

Viewing database size information

In privileged EXEC mode, enter the **show zoning operation-info** command.

Database and transaction size information is displayed in bytes.

```
switch# show zoning operation-info
db-max 1045274
db-avail 1043895
db-committed 367
db-transaction 373
transaction-token 1
last-zone-changed-timestamp 2011-11-16 16:54:31 GMT-7:00
last-zone-committed-timestamp 2011-11-16 16:23:44 GMT-7:00
```

Managing zone aliases

A zone alias is user-defined name for a logical group of ports or WWNs. You can simplify the process of creating and managing zones by first specifying aliases for zone members. Aliases facilitate tracking and eliminate the need for long lists of individual zone member names. An alias can be a member of a zone, but it cannot be a member of a zoning configuration.

Creating an alias

1. In privileged EXEC mode, enter the **show name-server detail** command to list the WWNs of devices and targets available in the Brocade VCS Fabric.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **zoning defined-configuration alias** command followed by a name for the alias.

A subconfiguration mode prompt appears.

4. Enter the subconfiguration mode **member-entry** command to specify at least one member entry.

The member entry must be specified as a port WWN or a node WWN.

You can add multiple members in one operation by separating each member entry with a semicolon (;). No spaces are allowed after the semicolon.

5. Enter the **exit** command to return to global configuration mode.
6. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the configuration to nonvolatile memory.

The following is an example of creating an alias with one member node WWN:

```
switch# show name-server detail
PID: 013100
Port Name: 20:00:00:00:00:00:00:01
Node Name: 10:00:00:00:00:00:00:01
(output truncated)
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration alias alias1
switch(config-alias-alias1)# member-entry 10:00:00:00:00:00:00:01
switch(config-alias-alias1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

Adding additional members to an existing alias

1. In privileged EXEC mode, enter the **show name-server detail** command to list the WWNs of devices and targets available in the Brocade VCS Fabric.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **zoning defined-configuration alias** command followed the name of an existing zone alias.

A subconfiguration mode prompt appears.

4. Enter the subconfiguration mode **member-entry** command to specify at least one member entry.

The member entry must be specified as a port WWN or a node WWN.

You can add multiple members in one operation by separating each member entry with a semicolon (;). No spaces are allowed after the semicolon.

5. Enter the **exit** command to return to global configuration mode.

- Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the configuration to nonvolatile memory.

Example of adding two member node WWNs to an existing alias:

```
switch# show name-server detail
PID: 013200
Port Name: 20:00:00:00:00:00:02
Node Name: 10:00:00:00:00:00:02
(output truncated)
PID: 013300
Port Name: 20:00:00:00:00:00:03
Node Name: 10:00:00:00:00:00:03
(output truncated)
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration alias alias1
switch(config-alias-alias1)# member-entry 10:00:00:00:00:00:02;10:00:00:00:00:00:03
switch(config-alias-alias1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)#
```

Removing a member from an alias

- In privileged EXEC mode, enter the **show running-config zoning** command to display the alias and its member WWNs.
- Enter the **configure terminal** command to enter global configuration mode.
- Enter the **zoning defined-configuration alias** command followed the name of an existing zone alias.
A subconfiguration mode prompt appears.
- Enter the subconfiguration mode **no member-entry** command to specify the WWN to be removed from the zone alias.
You can only remove one member at a time.
- Enter the **exit** command to return to the global configuration mode.
- Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the configuration to nonvolatile memory.

The following provides an example of removing two members from an alias:

```
switch# show running-config zoning
zoning defined-configuration alias alias1
member-entry 10:00:00:00:00:00:01
member-entry 10:00:00:00:00:00:02
member-entry 10:00:00:00:00:00:03
(output truncated)
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration alias alias1
switch(config-alias-alias1)# no member-entry 10:00:00:00:00:00:02
switch(config-alias-alias1)# no member-entry 10:00:00:00:00:00:03
switch(config-alias-alias1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

Deleting an alias

- In privileged EXEC mode, enter the **show running-config zoning** command to display the alias and its member WWNs.
- Enter the **configure terminal** command to enter global configuration mode.
- Enter the **no zoning defined-configuration alias** command followed by the name of the alias you want to delete.
- Enter the **show running-config zoning** command to verify the change in the defined configuration (optional).

5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the configuration to nonvolatile memory.

Example of deleting an alias:

```
switch# show running-config zoning
zoning defined-configuration alias alias1
  member-entry 10:00:00:00:00:00:01
!
zoning enabled-configuration cfg-name ""
zoning enabled-configuration default-zone-access allaccess
zoning enabled-configuration cfg-action cfg-none
switch#
switch# configure terminal
Entering configuration mode terminal
switch(config)# no zoning defined-configuration alias alias1
switch(config)# do show running-config zoning
zoning enabled-configuration cfg-name ""
zoning enabled-configuration default-zone-access allaccess
zoning enabled-configuration cfg-action cfg-none
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

Creating zones

Consider the following topics when creating zones.

Creating a zone

A zone cannot persist without any zone members. When you create a new zone, the **zoning defined-configuration zone** command places you in a command subconfiguration mode where you can add the first zone member entry. You can specify multiple members by separating each member from the next by a semicolon (;).

NOTE

Zones without any zone members cannot exist in volatile memory. They are deleted when the transaction commits successfully.

The following procedure adds a new zone to the defined configuration.

1. In privileged EXEC mode, enter the **show name-server detail** command to obtain the WWNs of servers and targets available in the Brocade VCS Fabric.
2. Enter the **configure terminal** command to enter global configuration mode.
3. Enter the **zoning defined-configuration zone** command and enter a new zone name to add a new zone.

A subconfiguration mode prompt appears.

4. Enter the subconfiguration mode **member-entry** command to specify at least one member entry.

The member entry must be specified as a port WWN, a node WWN, or an alias. You can mix WWNs and aliases.

Add multiple members in one operation by separating each member entry with a semicolon (;). No spaces are allowed after the semicolon.

5. Enter the **exit** command to return to global configuration mode.

- Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

Example of creating a zone with two members, a WWN and an alias:

```
switch# show name-server detail
PID: 012100
Port Name: 10:00:00:05:1E:ED:95:38
Node Name: 20:00:00:05:1E:ED:95:38
(output truncated)

switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration zone zone1
switch(config-zone-zone1)# member-entry 20:00:00:05:1E:ED:95:38;alias2
switch(config-zone-zone1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

Adding a member to a zone

- In privileged EXEC mode, enter the **show name-server detail** command to list the WWNs of devices and targets available on the Brocade VCS Fabric cluster.
- Enter the **configure terminal** command to enter global configuration mode.
- Enter the **zoning defined-configuration zone** command and enter the name of an existing zone.

A subconfiguration mode prompt appears.

- Enter the subconfiguration mode **member-entry** command and specify the member you want to add.

The new member can be specified by a port WWN, a node WWN, or a zone alias.

Add multiple members in one operation by separating each member with a semicolon (;).

- Enter the **exit** command to return to the global configuration mode.
- Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

Example of adding three members to a zone, two node WWNs and an alias:

```
switch# show name-server detail
PID: 012100
Port Name: 50:05:07:61:00:1b:62:ed
Node Name: 50:05:07:61:00:1b:62:ed
(output truncated)
PID: 012200
Port Name: 50:05:07:61:00:09:20:b4
Node Name: 50:05:07:61:00:09:20:b4
(output truncated)

switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration zone zone1
switch(config-zone-zone1)# member-entry 50:05:07:61:00:1b:62:ed;50:05:07:61:00:09:20:b4;alias3
switch(config-zone-zone1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

Removing a member from a zone

- In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.

2. Enter the **zoning defined-configuration zone** command and enter the name of the zone from which you want to remove a member.
A subconfiguration mode prompt appears.
3. Enter the subconfiguration mode **no member-entry** parameter and specify the WWN or the alias of the member you want to remove.
You can remove only one member at a time. To remove more than one member, you must issue the **no member-entry** command for each member you want to remove.
4. Enter the **exit** command to return to global configuration mode.
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

NOTE

The parent configuration is removed when the last child object is removed, irrespective of the save operation.

Example of removing more than one member from a zone:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration zone zone1
switch(config-zone-zone1)# no member-entry 50:05:07:61:00:09:20:b4
switch(config-zone-zone1)# no member-entry alias3
switch(config-zone-zone1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

Deleting a zone

Before deleting a zone, ensure that the zone is not a member of any enabled zone configuration. Although the deletion will proceed in RAM, you will not be able to save the configuration to nonvolatile memory if an enabled zone configuration has the deleted zone as a member.

1. In privileged EXEC mode, enter the **show running-config zoning defined-configuration** command and verify that the zone you want to delete is not a member of an enabled zone configuration. If the zone is a member of an enabled zone configuration, remove it.
2. Enter the **configure terminal** command to enter the global configuration mode.
3. Enter the **no zoning defined-configuration zone** command and enter the name of the zone you want to delete.

4. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

NOTE

The parent configuration is removed when the last child object is removed, irrespective of the save operation.

Example of removing a zone from the defined configuration:

```
switch# show running-config zoning defined-configuration
zoning defined-configuration zone zone1
member-entry 10:00:00:00:00:00:01
!
zoning defined-configuration zone zone2
member-entry 10:00:00:00:00:00:02
!
switch# configure terminal
Entering configuration mode terminal
switch(config)# no zoning defined-configuration zone zone2
switch(config)# zoning enabled-configuration cfg-action cfg-save
Updating flash ...
switch(config)# exit
switch# show running-config zoning defined-configuration
zoning defined-configuration zone zone1
member-entry 10:00:00:00:00:00:01
```

Managing zones

Consider the following topics when managing zones.

Viewing the defined configuration

To view the defined configuration, in privileged EXEC mode enter the **show running-config zoning defined-configuration** command.

For each configuration, the command lists each member zone. For each zone, the command lists the WWN or alias name of each member. The following example illustrates this.

```
switch# show running-config zoning defined-configuration

zoning defined-configuration cfg cfg0
member-zone zone_0_1
member-zone zone_0_2
member-zone zone_0_3
member-zone zone_0_4
member-zone zone_same
!
zoning defined-configuration cfg cfg1
member-zone zone_1_1
member-zone zone_1_2
member-zone zone_1_3
member-zone zone_1_4
member-zone zone_same
!
zoning defined-configuration cfg cfg2
member-zone zone_2_1
member-zone zone_2_2
member-zone zone_2_3
member-zone zone_2_4
member-zone zone_same
!
zoning defined-configuration cfg cfg4
member-zone zone2
member-zone zone3
```

```

!
zoning defined-configuration zone zone0
 member-entry 11:22:33:44:55:66:77:80
 member-entry 11:22:33:44:55:66:77:81
 member-entry 11:22:33:44:55:66:77:82
 member-entry 11:22:33:44:55:66:77:83
 member-entry 11:22:33:44:55:66:77:84
!
zoning defined-configuration zone zone1
 member-entry 11:22:33:44:55:66:77:80
 member-entry 11:22:33:44:55:66:77:81
 member-entry 11:22:33:44:55:66:77:82
 member-entry 11:22:33:44:55:66:77:83
 member-entry 11:22:33:44:55:66:77:84
!
zoning defined-configuration zone zone2
 member-entry 11:22:33:44:55:66:77:80
 member-entry 11:22:33:44:55:66:77:81
 member-entry 11:22:33:44:55:66:77:82
 member-entry 11:22:33:44:55:66:77:83
 member-entry 11:22:33:44:55:66:77:84
!
      (output truncated)

```

Viewing the enabled configuration

To view the enabled configuration, in privileged EXEC mode enter the **show zoning enabled-configuration** command. The following information about the enabled configuration is displayed:

- The name of the configuration
- The configuration action
- The mode of the default zone — the mode that will be active if you disable the enabled configuration

NOTE

In Network OS 4.0.0 and later, the enabled-zone output is no longer available from the **show running-config zoning enabled-configuration enabled-zone** command. It is now available from the **show running-config zoning enabled-configuration** command.

The configuration name has CFG_MARKER asterisk (*) appended to it if an outstanding transaction exists; the asterisk is not present if no outstanding transaction exists. Similarly, the configuration action is flagged as "cfg-save" if no outstanding transaction exists; "cfg-none" indicates that an outstanding transaction exists. A CFG_MARKER flag is appended to the configuration if the enabled configuration does not exactly match the defined configuration. This scenario occurs when you have an enabled configuration and make changes to the defined-configuration, and then, instead of enabling the defined configuration, you issue the **cfg-save** command.



CAUTION

When edits are made to the defined configuration, and those edits affect a currently enabled zone configuration, issuing a "cfg-save" command makes the enabled configuration effectively stale. Until the enabled configuration is reenabled, the merging of new RBridges into the cluster is not recommended. This merging may cause unpredictable results, with the potential for mismatched enabled-zoning configurations among the RBridges in the cluster.

Example of viewing the zoning enabled configuration:

```

switch# show zoning enabled-configuration

zoning enabled-configuration cfg-name cfg1
zoning enabled-configuration enabled-zone zone1 member-entry 10:00:00:00:00:00:00:01
zoning enabled-configuration enabled-zone zone2 member-entry 10:00:00:00:00:00:00:02

```

Creating a zone configuration

A zone configuration cannot persist without any member zones. When creating a new zone configuration, the **zoning defined-configuration cfg** command places you in a command sub-configuration mode where you must add at least one member zone. While zone configurations without any member zones can exist in volatile memory, they are deleted when the transaction commits successfully.

The following procedure adds a new zone configuration to the defined configuration.

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning defined-configuration cfg** command and enter a new configuration name.
A subconfiguration mode prompt appears.
3. Enter the **member-zone** subconfiguration mode command and specify the name of at least one zone.
Add multiple zones in one operation by separating each zone name with a semicolon (;).
4. Enter the **exit** command to return to global configuration mode.
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

Example of creating a zone configuration with one member zone:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration cfg config1
switch(config-cfg-config1)# member-zone zone1
switch(config-cfg-config1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

NOTE

Zone aliases are not valid zone configuration members. Adding an alias to an existing zone configuration will not be blocked. However, the attempt to enable a zone configuration that contains aliases will fail with an appropriate error message.

Adding a zone to a zone configuration

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning defined-configuration cfg** command and enter the name of the configuration to which you want to add zones.
The command prompt changes to indicate a subconfiguration mode.
3. Enter the **member-zone** subconfiguration mode command and specify the name of at least one member zone.
Add multiple zones in one operation by separating each zone name with a semicolon (;).
4. Enter the **exit** command to return to global configuration mode.
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

Example of adding two zones to config1:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration cfg config1
switch(config-cfg-config1)# member-zone zone2;zone3
switch(config-cfg-config1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

Removing a zone from a zone configuration

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning defined-configuration cfg** command and enter the name of the configuration from which you want to remove a zone.

The command prompt changes to indicate a subconfiguration mode.

3. Enter the **no member-zone** subconfiguration mode command and specify the name of the zone you want to remove from the configuration.

You can remove only one member at a time. To remove more than one member, you must issue the **no member-zone** command for each member you want to remove.

4. Enter the **exit** command to return to global configuration mode.
5. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified configuration to nonvolatile memory.

NOTE

The parent configuration is removed when the last child object is removed, irrespective of the save operation.

Example of removing two zones from config1:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration cfg config1
switch(config-cfg-config1)# no member-zone zone2
switch(config-cfg-config1)# no member-zone zone3
switch(config-cfg-config1)# exit
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

Enabling a zone configuration

Only one zone configuration can be enabled in a VCS Fabric. The following procedure selects a configuration from the defined configuration and makes it the enabled configuration. If a zone configuration is currently enabled, the newly enabled configuration replaces the previously enabled configuration.

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.

2. Enter the **zoning enabled-configuration cfg-name** command with the name of the configuration you want to enable.

In addition to enabling the specified configuration, this command also saves any changes made to the zoning database in volatile memory to nonvolatile memory. The saved configuration is persistent.

If the configuration refers to a nonexistent zone or a zone with no members assigned to it, the operation fails and the command returns an error message. The following example enables config1.

Example of enabling a zone configuration:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration cfg-name config1
```

Example of a failed enable operation:

The enable operation fails because the configuration contains a zone without members.

```
switch(config)# do show running-config zoning
zoning defined-configuration cfg cfg1
member-zone-zone1
member-zone zone2
!
zoning defined-configuration zone zone1 <-----Zone with no member
!
zoning defined-configuration zone zone2
member-entry 20:03:00:11:0d:bc:76:09
!
zoning enabled-configuration cfg-name ""
zoning enabled-configuration default-zone-access allaccess
zoning enabled-configuration cfg-action cfg-none
switch(config)# zoning enabled-configuration cfg-name cfg1
% Error: Command Failed. Cfg contains empty zone object "zone1"
```

Disabling a zone configuration

Disabling the currently enabled configuration returns the fabric to no-zoning mode. All devices can then access one another or not at all, depending on the default zone access mode setting.

NOTE

For fabrics with many devices, Brocade recommends setting the default zone access mode to No Access before disabling a zone configuration to avoid RSCN storms.

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **no zoning enabled-configuration cfg-name** command.

In addition to disabling the currently enabled configuration, this command also saves any changes made to the zoning database in volatile memory to nonvolatile memory. The saved configuration is persistent.

Example of disabling a zone configuration:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no zoning enabled-configuration cfg-name
```

Deleting a zone configuration

The following procedure deletes a zone configuration from the defined configuration.

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.

2. Enter the **no zoning defined-configuration cfg** command and the name of the zone configuration you want to delete.
3. Enter the **zoning enabled-configuration cfg-action cfg-save** command to save the modified defined configuration to nonvolatile memory.

Example of deleting a zone configuration:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# no zoning defined-configuration cfg cfg2
switch(config)# zoning enabled-configuration cfg-action cfg-save
```

NOTE

If you try to delete the enabled configuration from the defined configuration, the **zoning enabled-configuration cfg-action cfg-save** command returns an error. However, if you commit the transaction with the **zoning enabled-configuration cfg-action cfg-disable** command, the operation proceeds without error.

Clearing changes to a zone configuration

The following procedure aborts all pending transactions and removes all uncommitted operations from the database. It returns the configuration in volatile memory to the state it was in when a **zoning enabled-configuration cfg-action cfg-save** or **zoning enabled-configuration cfg-name** command was last executed successfully.

1. In privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter the **zoning enabled-configuration cfg-action cfg-transaction-abort** command.

Example of aborting a transaction:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration cfg-action cfg-transaction-abort
```

Clearing all zone configurations

The following procedure clears all zone configurations from the defined configuration and enables the default zone.

NOTE

For fabrics with many devices, Brocade recommends setting the default access mode to No Access before clearing all zone configurations to avoid RSCN storms.

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **zoning enabled-configuration cfg-action cfg-clear** command.
3. Enter one of the following commands, depending on whether an enabled zone configuration exists:
 - If no enabled zone configuration exists, enter the **zoning enabled-configuration cfg-action cfg-save** command.
 - If an enabled zone configuration exists, enter the **no zoning enabled-configuration cfg-name** command to disable and clear the zone configuration in nonvolatile memory for all switches in the fabric.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration cfg-action cfg-clear
switch(config)# no zoning enabled-configuration cfg-name
```

Backing up the zone configuration

To back up your zoning configuration you copy it to a file and store it on a server or on an attached USB device. You can use the copy to restore the configuration if needed.

NOTE

Ensure that no transaction is pending before you perform the copy operation, otherwise failures will occur when you attempt to restore configuration entries from the saved file. Any transaction data would cause such a failure, including empty configuration definitions or empty zones.

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Empty the transaction buffer by either committing the transaction to nonvolatile memory or aborting the transaction.
 - To commit the transaction, enter the **zoning enabled-configuration cfg-action cfg-save** command, the **zoning enabled configuration cfg-name** command, or the **zoning enabled-configuration cfg-action cfg-disable** command.
 - To abort the transaction, enter the **zoning enabled-configuration cfg-action cfg-transaction-abort** command.
3. Enter the **exit** command to return to privileged EXEC mode.
4. Enter the **copy** command. For the source file, use **running-config** . For the destination file, use the file name you want the configuration copied to.

Example of making a backup copy on a USB device:

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning enabled-configuration cfg-action cfg-save
switch(config)# exit
switch# copy running-config usb://myconfig
```

Restoring a configuration from backup

When you restore a configuration from backup and add to the running configuration, the zone configuration identified in the backup copy as the enabled configuration becomes the new enabled configuration.

In privileged EXEC mode, enter the **copy** command. For the source file use the file where the saved configuration is stored. For the destination file, use **running-config**.

This operation updates the defined configuration in RAM.

NOTE

The **copy** command adds to the defined configuration. It does not replace the defined configuration.


The following example adds the configuration in the file named myconfig on the attached USB device to the defined configuration.

```
switch# copy usb://myconfig running-config
```

Zone configuration scenario example

This example creates the zone configuration shown below. The example assumes that two hosts need access to the same storage device, while each host needs private storage of its own. You create two zones: Zone A contains Host 1, its private storage device, and the shared storage device; Zone B contains Host 2, its private storage device, and the shared storage device. In addition, you create two zone configurations: cfg1 in which only Zone A is effective; cfg2, in which both zones are effective.

FIGURE 27 Zone configuration example



1. Log in to any switch in the Brocade VCS Fabric.
2. Enter the **show name-server detail** command to list the available WWNs.
3. Enter the **configure terminal** command to enter global configuration mode.
4. Enter the **zoning defined-configuration zone** command to create Zone A.
5. Enter the **zoning defined-configuration zone** command to create Zone B.
6. Enter the **zoning defined-configuration cfg** command to create the configuration **cfg1** with Zone A as its only member.
7. Enter the **zoning defined-configuration cfg** command to create the configuration **cfg2** with Zone A and Zone B as its members.
8. Enter the **zoning running-config defined-configuration** command to view the defined zone configuration.
9. Enter the **zoning enabled-configuration cfg-name** command to enable **cfg2**.

10. Verify the enabled zoning configuration, by means of the **show zoning enabled-configuration** command.

```
switch# show name-server detail
switch# configure terminal
Entering configuration mode terminal
switch(config)# zoning defined-configuration zone ZoneA
switch(config-zone-ZoneA)# member-entry 20:00:00:05:1e:ed;
95:38;50:05:07:61:00:09:20:b4;50:05:07:61:00:1b:62:ed
switch(config-zone-ZoneA)# exit
switch(config)# zoning defined-configuration zone ZoneB
switch(config-zone-ZoneB)# member-entry 20:00:00:05:1e:ed:18:c3;50:05:07:61:00:22:18:9b;
50:05:07:61:00:1b:62:ed
switch(config-zone-ZoneB)# exit
switch(config)# zoning defined-configuration cfg cfg1
switch(config-cfg-cfg1)# member-zone ZoneA
switch(config-cfg-cfg1)# exit
switch(config)# zoning defined-configuration cfg cfg2
switch(config-cfg-cfg2)# member-zone ZoneA;ZoneB
switch(config-cfg-cfg2)# exit
switch(config)# zoning enabled-configuration cfg-name cfg2
switch(config)# exit
switch# show zoning enabled-configuration
zoning enabled-configuration cfg cfg1
  member-zone ZoneA
!
zoning enabled-configuration cfg cfg2
  member-zone ZoneA
  member-zone ZoneB
!
zoning enabled-configuration zone ZoneA
  member-entry 20:00:00:05:1e:ed:95:38
  member-entry 50:05:07:61:00:09:20:b4
  member-entry 50:05:07:61:00:1b:62:ed
!
zoning enabled-configuration zone ZoneB
  member-entry 20:00:00:05:1e:ed:18:c3
  member-entry 50:05:07:61:00:22:18:9b
  member-entry 50:05:07:61:00:1b:62:ed
```

Merging zones

This section provides the background needed to merge zones successfully. The tables at the end of this section summarize scenarios involving Switch A and Switch B and the results to be expected following a merge.

Preconditions for zone merging

When a new switch is added to a VCS fabric, it automatically inherits the zone configuration information from the fabric. You can verify the zone configuration on any switch by using the procedure described in [Viewing the defined configuration](#) on page 140. Take care to avoid mismatched enabled-configuration scenarios.



CAUTION

When edits are made to the defined configuration, and those edits affect a currently enabled zone configuration, issuing a "cfg-save" command makes the enabled configuration effectively stale. Until the enabled configuration is reenabled, the merging of new R Bridges into the cluster is not recommended. This merging may cause unpredictable results, with the potential for mismatched enabled-zoning configurations among the R Bridges in the cluster.

If you are adding a switch that is already configured for zoning, you must clear the zone configuration on that switch before connecting it to the zoned fabric. Refer to [Clearing all zone configurations](#) on page 145 for instructions.

Adding a new fabric that has no zone configuration information to an existing zoned fabric is very similar to adding a new switch. All switches in the new fabric inherit the zone configuration data. If the existing fabric has an effective zone configuration, then the same configuration becomes the effective configuration for all switches in the added fabric.

NOTE

To prevent an unwanted zone merge, use the **no fabric isl enable** command on ISL interfaces instead of the **shutdown** command on tengigabitethernet ports.

Before the new fabric can merge successfully, it must satisfy the following criteria:

- Before merging
 - Ensure that all switches adhere to the default zone merge rules as described in [Zone merging scenarios](#) on page 150.
 - Ensure that the enabled and defined zone configurations match. If they do not match and you merge with another switch, the merge might be successful, but unpredictable zoning and routing behavior can occur. Refer to the Caution in this section and refer to [Viewing the defined configuration](#) on page 140.
- Merging and segmentation

The system checks each port as it comes online to determine whether the ports should be segmented. E_Ports come online on power up, enabling a switch, or adding a new switch, and the system checks the zone database to detect if the two database that can be merged safely. Refer to [Zone merging scenarios](#) on page 150.

Observe the following rules when merging zones:

- Merging rules
 - Local and adjacent configurations: If the local and adjacent zone database configurations are the same, they will remain unchanged after the merge.
 - Enabled configurations: If there is an enabled configuration between two switches, the enabled zone configurations must match.
 - Zone membership: If a zoning object has the same name in both the local and adjacent defined configurations, the content and order of the members are important.
 - Objects in adjacent configurations: If a zoning object appears in an adjacent defined configuration, but not in the local defined configuration, the zoning object is added to the local defined configuration. The modified zone database must fit in the nonvolatile memory area allotted for the zone database.
 - Local configuration modification: If a local defined configuration is modified because of a merge, the new zone database is propagated to the other switches within the merge request.
- Merging two fabrics

For best practices, the default-zone access modes should match, although this is not a requirement. Refer to [Zone merging scenarios](#) on page 150.

If the two fabrics have conflicting zone configurations, they will not merge. If the two fabrics cannot join, the ISLs between the switches will segment.

The transaction state after the merge depends on which switch is elected as the principal RBridge. The newly elected principal RBridge retains the same transaction information it had before the merge. Transaction data is discarded from any switch that lost its principal status during the merge.

- Merge conflicts

When a merge conflict is present, a merge does not take place and the ISLs will segment.

If the fabrics have different zone configuration data, the system attempts to merge the two sets of zone configuration data. If the zones cannot merge, the ISLs will be segmented.

- A merge is not possible under any of the following conditions:
 - Configuration mismatch: Zoning is enabled in both fabrics and the zone configurations that are enabled are different in each fabric.
 - Zone Database Size: The zone database size exceeds the maximum limit of another switch.

NOTE

If the zone members on two switches are not listed in the same order, the configuration is considered a mismatch, and the switches will segment from the fabric. For example: `cfg1 = z1; z2` is different from `cfg1 = z2; z1`, even though the members of the configuration are the same. If zone members on two switches have the same names defined in the configuration, make sure the zone members are listed in the same order.

Fabric segmentation and zoning

If the connections between two fabrics are no longer available, the fabric segments into two separate fabrics. Each new fabric retains the previous zone configuration.

If the connections between two fabrics are replaced and no changes have been made to the zone configuration in either of the two fabrics, the two fabrics can merge back into one single fabric. If any changes that cause a conflict have been made to either zone configuration, a fabric merge may fail.

Zone merging scenarios

The following tables provide information on merging zones and the expected results.

TABLE 25 Zone merging scenarios: Defined and enabled configurations

Description	Switch A	Switch B	Expected results
Switch A has a defined configuration. Switch B does not have a defined configuration.	defined:cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: none	defined: none enabled: none	Configuration from Switch A propagates throughout the fabric in an inactive state, because the configuration is not enabled.
Switch A has a defined and enabled configuration. Switch B has a defined configuration but no enabled configuration.	defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: cfg1	defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: none	Configuration from Switch A propagates throughout the fabric. The configuration is enabled after the merge in the fabric.
Switch A and Switch B have the same defined configuration. Neither have an enabled configuration.	defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: none	defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: none	No change (clean merge).
Switch A and Switch B have the same defined and enabled configuration.	defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: cfg1:	defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: cfg1:	No change (clean merge).
Switch A does not have a defined configuration. Switch B has a defined configuration.	defined: none enabled: none	defined:cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: none	Switch A absorbs the configuration from the fabric.

TABLE 25 Zone merging scenarios: Defined and enabled configurations (continued)

Description	Switch A	Switch B	Expected results
<p>Switch A does not have a defined configuration.</p> <p>Switch B has a defined and enabled configuration.</p>	<p>defined: none</p> <p>enabled: none</p>	<p>defined:cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>enabled: cfg1</p>	<p>Switch A absorbs the configuration from the fabric, with cfg1 as the enabled configuration.</p>
<p>Switch A and Switch B have the same defined configuration. Only Switch B has an enabled configuration.</p>	<p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>enabled: none</p>	<p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>enabled: cfg1</p>	<p>Clean merge, with cfg1 as the enabled configuration.</p>
<p>Switch A and Switch B have different defined configurations. Neither have an enabled configuration.</p>	<p>defined: cfg2 zone2: 10:00:00:90:69:00:00:8c; 10:00:00:90:69:00:00:8d</p> <p>enabled: none</p>	<p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>enabled: none</p>	<p>Clean merge. The new configuration will be a composite of the two.</p> <p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>defined: cfg2 zone2: 10:00:00:90:69:00:00:8c; 10:00:00:90:69:00:00:8d</p> <p>enabled: none</p>
<p>Switch A and Switch B have different defined configurations. Switch B has an enabled configuration.</p>	<p>defined: cfg2 zone2: 10:00:00:90:69:00:00:8c; 10:00:00:90:69:00:00:8d</p> <p>enabled: none</p>	<p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>enabled: cfg1</p>	<p>Clean merge. The new configuration is a composite of both, with cfg1 as the enabled configuration.</p>
<p>Switch A does not have a defined configuration.</p> <p>Switch B has a defined configuration and an enabled configuration, but the enabled configuration is different from the defined configuration.</p>	<p>defined: none</p> <p>enabled: none</p>	<p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>effective: cfg1</p> <p>zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>zone2: 10:00:00:90:69:00:00:8c, 10:00:00:90:69:00:00:8d</p>	<p>Clean merge. Switch A absorbs the defined configuration from the fabric, with cfg1 as the effective configuration.</p> <p>In this case, however, the effective configurations for Switch A and Switch B are different. You should issue a zoning enabled-configuration cfg-name command from the switch with the proper effective configuration.</p>

TABLE 26 Zone merging scenarios: Different content

Description	Switch A	Switch B	Expected results
Enabled configuration mismatch.	<p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>enabled: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p>	<p>defined: cfg2 zone2: 10:00:00:90:69:00:00:8c; 10:00:00:90:69:00:00:8d</p> <p>enabled: cfg2 zone2: 10:00:00:90:69:00:00:8c; 10:00:00:90:69:00:00:8d</p>	<p>Fabric segments due to mismatching zone configurations</p>
Configuration content mismatch.	<p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b</p> <p>enabled: irrelevant</p>	<p>defined: cfg1 zone1: 10:00:00:90:69:00:00:8c; 10:00:00:90:69:00:00:8d</p> <p>enabled: irrelevant</p>	<p>Fabric segments due to mismatching zone content</p>

TABLE 27 Zone merging scenarios: Different names

Description	Switch A	Switch B	Expected results
Same content, different enabled configuration name.	defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b	defined:cfg2 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: cfg2 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b	Fabric segments due to mismatching zone configurations
Same content, different zone name.	defined: cfg1 zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: irrelevant	defined: cfg1 zone2: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b enabled: irrelevant	Fabric segments due to mismatching zone content
Same name, same content, different order.	defined: cfg1zone1: 10:00:00:90:69:00:00:8a; 10:00:00:90:69:00:00:8b; 10:00:00:90:69:00:00:8c enabled: irrelevant	defined: cfg1zone1: 10:00:00:90:69:00:00:8b; 10:00:00:90:69:00:00:8c; 10:00:00:90:69:00:00:8a enabled: irrelevant	Fabric segments due to mismatching zone content
Same name, different types.	effective: zone1: MARKETING	enabled: cfg1: MARKETING	Fabric segments due to mismatching types

TABLE 28 Zone merging scenarios: Default access mode

Description	Switch A	Switch B	Expected results
Different default zone access mode settings.	default zone: All Access	default zone: No Access	Clean merge. No Access takes precedence and default zone configuration from Switch B propagates to fabric. default zone: No Access
Same default zone access mode settings.	default zone: All Access	default zone: All Access	Clean merge. Default zone configuration is All Access in the fabric.
Same default zone access mode settings.	default zone: No Access	default zone: No Access	Clean merge. Default zone configuration is No Access in the fabric.
Enabled zone configuration.	No enabled configuration. default zone = All Access	enabled: cfg2 default zone: All Access or No Access	Clean merge. Enabled zone configuration and default zone mode from Switch B propagates to fabric.
Enabled zone configuration.	No enabled configuration. default zone = No Access	enabled: cfg2 default zone: All Access	Fabric segments because Switch A has a hidden zone configuration (No Access) activated and Switch B has an explicit zone configuration activated.
Enable zone configuration.	enabled: cfg1 default zone: No Access	No enabled configuration. default zone: No Access	Clean merge. Enabled zone configuration from Switch A propagates to fabric.
Enable zone configuration.	enabled: cfg1 default zone: All Access	No enabled configuration. default zone: No Access	Fabric segments. You can resolve the zone conflict by changing the default zone to No Access on Switch A .

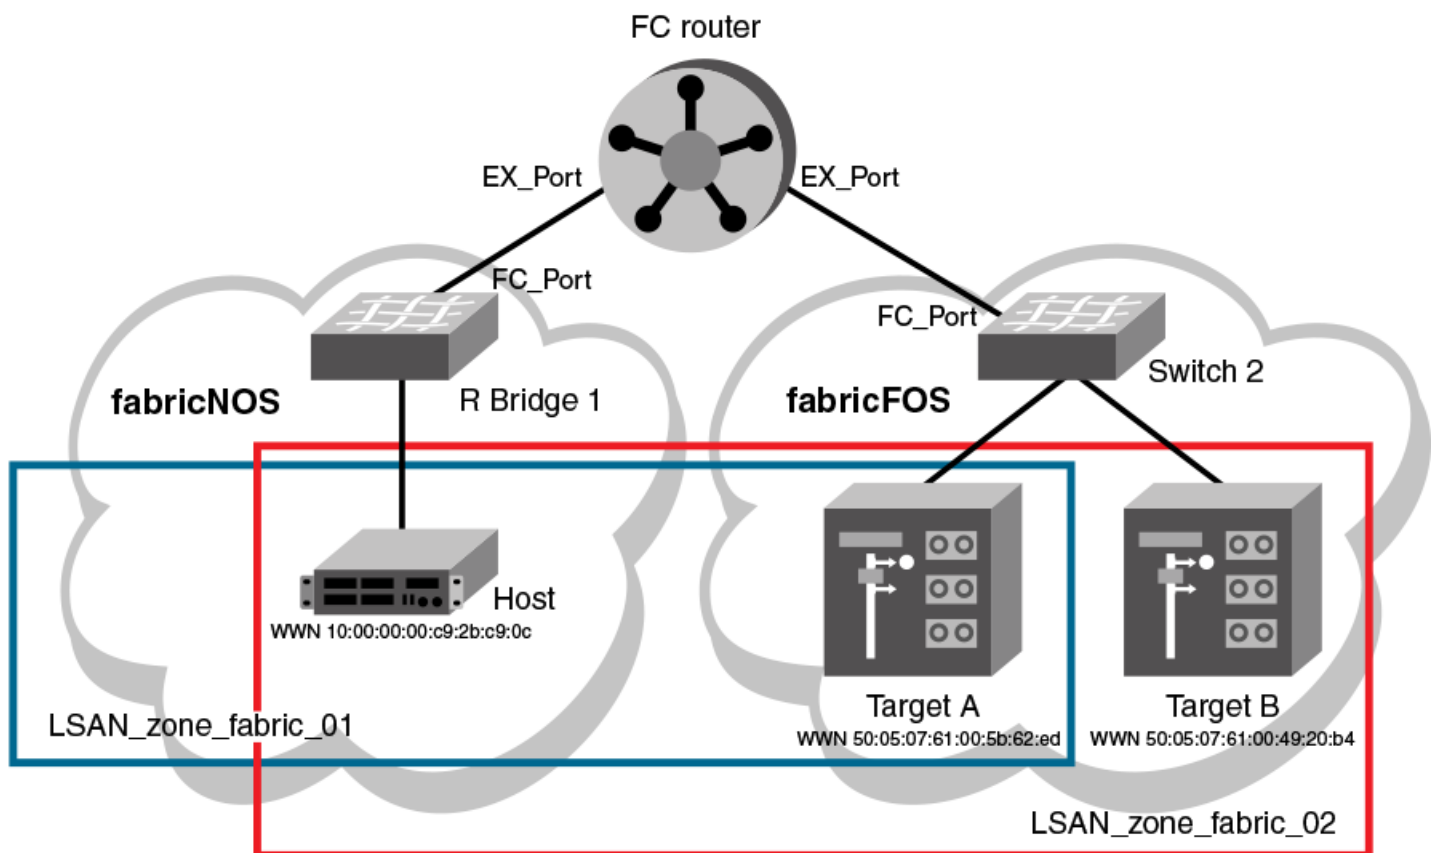
Configuring LSAN zones: Device-sharing example

The following example shows LSANs sharing devices in separate fabrics. The procedure illustrates the creation of two LSAN zones (called `lsan_zone_fabric_02` and `lsan_zone_fabric_01`), which involve the following devices and connections:

- RBridge1 and the host in a Network OS fabric named `fabric_01`.
- Switch2, Target A, and Target B in a Fabric OS fabric named `fabric_02`.
- RBridge1 is connected by one of its FC_Ports to an EX_Port on the FC router.
- Switch2 is connected to the FC router using another EX_Port or VEX_Port.
- Host has WWN 10:00:00:00:c9:2b:c9:0c (connected to RBridge1).
- Target A has WWN 50:05:07:61:00:5b:62:ed (connected to switch2).
- Target B has WWN 50:05:07:61:00:49:20:b4 (connected to switch2).

The following illustration shows the connectivity.

FIGURE 28 LSAN zones example



The following example steps create this set of LSAN zones.

1. Obtain the host WWN in fabric_01:
 - a) Log in to any switch in fabric_01.
 - b) On the fabric_01 switch, enter the **show name-server detail** command to list the WWN of the host (10:00:00:00:c9:2b:c9:0c).

NOTE

The **show name-server detail** output displays both the port WWN and node WWN; the port WWN must be used for LSANs.

```
switch# show name-server detail
PID: 012100
Port Name: 10:00:00:00:c9:2b:c9:0c
Node Name: 20:00:00:00:c9:2b:c9:0c
SCR: 3
FC4s: FCP
PortSymb: [27] "Brocade-1020|2.3.0.0|localhost.localdomain|Red Hat
Enterprise Linux Server release 5.5"
NodeSymb: NULL
Fabric Port Name: 20:21:00:05:1E:CD:79:7A
Permanent Port Name: 10:00:00:00:c9:2b:c9:0c
Device type: Physical Initiator
Interface: Fcoe 1/1/9
Physical Interface: Te 1/0/9
Share Area: No
Redirect: No
```

2. Obtain the target WWNS in fabric_02:
 - a) Log in as admin on switch2 in fabric_02.
 - b) On fabric_02, enter the **nsShow** command to list Target A (50:05:07:61:00:5b:62:ed) and Target B (50:05:07:61:00:49:20:b4).

```
switch:admin> nsshows
{
Type Pid    COS  PortName                               NodeName                               TTL(sec)
NL  0508e8; 3;  50:05:07:61:00:5b:62:ed; 50:05:07:61:00:1b:62:ed; na
FC4s: FCP [IBM  DNEF-309170  F90F]
Fabric Port Name: 20:08:00:05:1e:34:11:e5
Permanent Port Name: 50:05:07:61:00:5b:62:ed

NL  0508ef; 3;  50:05:07:61:00:49:20:b4; 50:05:07:61:00:09:20:b4; na
FC4s: FCP [IBM  DNEF-309170  F90F]
Fabric Port Name: 20:08:00:05:1e:34:11:e5
Permanent Port Name: 50:05:07:61:00:49:20:b4
The Local Name Server has 2 entries }
```

3. Create an LSAN zone in the Network OS fabric (fabric_01)
4. In fabric_01, enter the **zoning defined-configuration zone** command to create the LSAN `lsan_zone_fabric_01`, and include the host.

```
switch# config terminal
switch(config)# zoning defined-configuration zone lsan_zone_fabric_01
switch(config-zone-lsan_zone_fabric_01)# member-entry 10:00:00:00:c9:2b:c9:0c
```

5. In fabric_01, add Target A to the LSAN.

```
switch(config-zone-lsan_zone_fabric_01)# member-entry 50:05:07:61:00:5b:62:ed
switch(config-zone-lsan_zone_fabric_01)# exit
```

6. In fabric_01, enter the **zoning defined-configuration cfg** and **zoning enabled-configuration cfg-name** commands to add and enable the LSAN configuration.

```
switch(config)# zoning defined-configuration cfg zone_cfg
switch(config-cfg-zone_cfg)# member-zone lsan_zone_fabric_01
switch(config-cfg-zone_cfg)# exit
switch(config)# zoning enabled-configuration cfg_name zone_cfg
```

Create an LSAN zone in the Fabric OS fabric (fabric_02)

7. On switch2 (fabric_02), enter the **zoneCreate** command to create the LSAN lsan_zone_fabric2, which includes the host (10:00:00:00:c9:2b:c9:0c), Target A (50:05:07:61:00:5b:62:ed), and Target B (50:05:07:61:00:49:20:b4).

```
switch:admin> zonecreate "lsan_zone_fabric_02", "10:00:00:00:c9:2b:c9:0c;
                                                50:05:07:61:00:5b:62:ed;
                                                50:05:07:61:00:49:20:b4"
```

8. On switch2 (fabric_02), enter the **cfgShow** command to verify that the zones are correct.

```
switch:admin> cfgshow
Defined configuration:
  zone: lsan_zone_fabric_02
        10:00:00:00:c9:2b:c9:0c;
        50:05:07:61:00:5b:62:ed;
        50:05:07:61:00:49:20:b4
Effective configuration:
  no configuration in effect
```

9. On switch2 (fabric_02), enter the **cfgAdd** and **cfgEnable** commands to create and enable the LSAN configuration.

```
switch:admin> cfgadd "zone_cfg", "lsan_zone_fabric_02"
switch:admin> cfgenable "zone_cfg"
You are about to enable a new zoning configuration.
This action will replace the old zoning configuration with the
current configuration selected.
Do you want to enable 'zone_cfg' configuration (yes, y, no, n): [no] y
zone config "zone_cfg" is in effect
Updating flash ...
```

Display the configuration on the FC router:

10. Log in as an admin and connect to the FC router.

11. On the FC router, enter the following commands to display information about the LSANs.

The **lsanZoneShow -s** command shows the LSAN.

```
switch:admin> lsanzoneshow -s
Fabric ID: 2 Zone Name: lsan_zone_fabric_02
10:00:00:00:c9:2b:c9:0c Imported
50:05:07:61:00:5b:62:ed EXIST
50:05:07:61:00:49:20:b4 EXIST
Fabric ID: 75 Zone Name: lsan_zone_fabric_01
10:00:00:00:c9:2b:c9:0c EXIST
50:05:07:61:00:5b:62:ed Imported
```

The **fcRPhyDevShow** command shows the physical devices in the LSAN.

```
switch:admin> fcrphydevshow
Device      WWN                Physical
Exists
in Fabric   PID
-----
75          10:00:00:00:c9:2b:c9:0c  c70000
2           50:05:07:61:00:49:20:b4  0100ef
2           50:05:07:61:00:5b:62:ed  0100e8
Total devices displayed: 3
```

The **fcRProxyDevShow** command shows the proxy devices in the LSAN.

```
switch:admin> fcrproxydevshow
Proxy      WWN                Proxy Device  Physical State
Created
in Fabric  PID Exists        PID
-----
75          50:05:07:61:00:5b:62:ed  01f001  2           0100e8  Imported
2           10:00:00:00:c9:2b:c9:0c  02f000  75          c70000  Imported
Total devices displayed: 2
```

On the FC router, the host and Target A are imported, because both are defined by `lsan_zone_fabric_02` and `lsan_zone_fabric_01`. However, target B is defined by `lsan_zone_fabric_02` and is not imported because `lsan_zone_fabric_01` does not allow it.

Configuring Fibre Channel Ports

- [Fibre Channel ports overview.....](#) 157
- [Connecting to an FC Fabric through an FC Router.....](#) 158
- [Fibre Channel port configuration.....](#) 158

Fibre Channel ports overview

Fibre Channel (FC) ports provide the ability to connect a Brocade VCS Fabric cluster to a Fibre Channel switch in a Fabric OS network.

The FlexPort feature is the only method that a switch running Network OS 5.0.0 can connect FC ports. For instructions on how to use FlexPort, refer to the "Configuring FlexPort" chapter of the *Network OS Layer 2 Switching Configuration Guide*.

These connections can be regular but not long distance. The following Network OS Fibre Channel port types are supported:

- E_Port: Can be used to connect only to an EX_Port on a FC SAN with Fibre Channel Routing configured.
- F_Port:
 - Supports FC target connectivity (standards based F_Port).
 - Supports bidirectional traffic internally from VF_Port, or internal ISL port.

NOTE

You must enable **fcoepport default** for the interface for the Fibre Channel logins to be available to connect to F_Ports.

- Auto (G_Port) – This is the default.
- N_port:
 - Default port type in Access Gateway mode
 - Available in Access Gateway mode only
 - Supports bidirectional traffic internally from VF_Port.
 - External connection to F_Port on a FC SAN
- VF_port:
 - For FCoE initiator or target connectivity.
 - Supports bidirectional traffic internally to E_Port, F_Port, VF_Port (all in FCF mode), and N_Port (in AG mode).

FC ports can connect to a FC switch in a Fabric OS network through two methods:

- Using an Inter-Switch Link (ISL) connection from an FC E_port on a VDX 6740 switch to a FC router's EX_Port; this in turn connects to the FC switch in the Fabric OS network. The VDX FC ports are configured as E_Ports for this connection. Refer to [Connecting to an FC Fabric through an FC Router](#) on page 158.
- Using a direct connection from an FC port on a VDX 6740 switch's to a F_Port on a FC router in a Fabric OS network. The VDX FC ports are configured as N_Ports through the Access Gateway feature. Refer to [Using Access Gateway](#) on page 165.

Connecting to an FC Fabric through an FC Router

FC ports on VDX 6740 switches can provide a connection to a FC switch in the FC SAN.


The FC ports on the VDX 6740 switch can be configured as FC E_Ports to connect with EX_Ports on a FC router through Inter-Switch Links (ISL) links. In turn, EX_Ports on the FC router connect to F_Ports on the FC Fabric switch. These connections provide support for zoning across Network OS and Fabric OS fabric types, which can enable FCoE devices on the Brocade VCS Fabric cluster to access SAN storage and services. Refer to [Fibre Channel ports overview](#) on page 157 for information on how to create LSAN zones.

The following figure shows an FC connection between a Network OS fabric and Fibre Channel SAN.

NOTE

For details of Fibre Channel routing concepts, refer to the *Fabric OS Administrator's Guide*

FIGURE 29 FC connection between a Network OS fabric and a Fibre Channel SAN



Fibre Channel port configuration

Consider the topics discussed below when configuring Fibre Channel ports.

Using Fibre Channel commands

Network OS software provides the following high-level commands for managing Fibre Channel ports:

- **interface FibreChannel** - Global configuration mode command that allows you to enter the interface Fibre Channel configuration submode where you can enter commands to activate and deactivate a Fibre Channel port (**no shutdown** and **shutdown** commands) and to set port attributes (**desire-distance**, **fill-word**, **isl-r_rdy**, **speed**, **trunk-enable**, and **vc-link-init** commands).

- **show running-config interface FibreChannel** - A privileged EXEC mode command that displays Fibre Channel port configuration information.
- **show interface FibreChannel** - A privileged EXEC mode command that displays hardware counters that monitor activity and status of a Fibre Channel port.

Activating and deactivating Fibre Channel ports

When VCS mode is enabled and an FCoE license is installed, all FC ports are activated by default. When enabling a switch for Access Gateway Mode, all FC ports are re-enabled as N_Ports.

Prerequisites for enabling Fibre Channel ports

Follow these steps before enabling a Fibre Channel port.

1. An FCoE license must be installed on the Brocade VDX 6740 to allow Fibre Channel port activation. Refer to the *Network OS Software Licensing Guide* for details about installing the FCoE license. Once the FCoE license is installed, all Fibre Channel ports are activated by default.
2. The Flexports on the switch must be configured as fibre channel ports. For instructions on how to use Flexport, refer to the "Configuring Flexport" chapter of the *Network OS Layer 2 Switching Configuration Guide* and to the *Brocade VDX 6740 Hardware Reference Guide*.

NOTE

Access Gateway is a feature available only on VDX 6740 switches. When activated, all Fibre Channel ports are re-enabled as N_Ports.

Enabling a Fibre Channel port

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **interface FibreChannel rbridge-id/slot/port** command for the Fibre Channel port you want to enable.

A configuration submenu prompt appears.

3. Enter the **no shutdown** command.

The following example enables port 1 on RBridge 8.

```
switch# configure terminal
switch(config)# rbridge-id 8
switch(config-rbridge-id-8)# interface FibreChannel 8/0/1
switch(config-FibreChannel-8/0/1)# no shutdown
```

Disabling a Fibre Channel port

1. In privileged EXEC mode, enter the **configure terminal** command to enter the global configuration mode.
2. Enter the **interface FibreChannel rbridge-id/slot/port** command for the Fibre Channel port you want to disable.

A configuration submenu prompt appears.

3. Enter the **shutdown** command.

The following example disables port 1 on routing bridge 8.

```
switch# configure terminal
switch(config)# interface FibreChannel 8/0/1
switch(conf-FibreChannel-8/0/1)# shutdown
```

Configuring and viewing Fibre Channel port attributes

This section introduces the options for configuring a variety of Fibre Channel port attributes and confirming the status of those attributes.

Using Fibre Channel port commands

Network OS v2.1.1 and later allows you to configure and display the following attributes for a Fibre Channel port by using the commands shown below:

- Port speed — Enter the interface Fibre Channel configuration submode **speed** command to set the speed of a Fibre Channel port.
- Trunk port — Enter the interface Fibre Channel configuration submode **trunk-enable** command to configure the port for trunking.
- Buffer credit control — Enter the interface Fibre Channel configuration submode **isl-r_r-rdy** command to enable interswitch link receiver-ready (ISL R_RDY) mode on the port. Enter the interface Fibre Channel configuration submode **no isl-r_r-rdy** command to disable ISL R_RDY mode on the port. If ISL R_RDY is not set, then Inter-SwitchLink Virtual Channel ready (ISL VC_RDY) mode is set by default.
- Forward error correction — **fec-enable** is set by default for 16G FC port speeds; use **no fec-enable** to disable this capability.
- FC port type restriction — **config-mode** allows the following possible completions:
 - auto — Configures the port as a G-Port (locked).
 - eport — Configures the port as a E-Port (locked).
 - fport — Configures the port as a F-Port (locked).
 - nport — Configures the port as a N-Port (locked).
- Port group speed — The hardware connector group-speed configuration allows you to set the allowed speed ranges and protocol type for groups of Flexports.

ATTENTION

Setting ISL R_RDY is not recommended.

The following Fibre Channel port attributes are not supported by Network OS 5.0.0:

AL_PA offset 13	F_Port buffers	NPIV capability
Compression	Fault Delay	NPIV PP Limit
Credit Recovery	Frame shooter port	Persistent Disable
CSCTL mode	Locked L_Port	Port Auto Disable
D-Port mode	LOS TOV enable	QoS E_Port
Disabled E_Port	Mirror Port	Rate limit
Encryption		RSCN suppressed
EX_Port		

Viewing Fibre Channel port attributes

To view the Fibre Channel port attributes for a single port, in privileged EXEC mode, enter the **show running-config interface FibreChannel *rbridge-id/slot/port*** command for the port you want to view. To view the Fibre Channel port attributes for all Fibre Channel ports in the fabric, enter the **show running-config interface FibreChannel** command without any additional parameters.

Whether you view attributes for a single port or for all ports, the settings for the **isl-r_rdy**, **trunk-enable**, and **shutdown** attributes are always displayed. The **fec-enable** and **speed** attributes are displayed only if they are set to nondefault values.

The following example displays the Fibre Channel port attributes for a single port. In this case, the **speed** and **vc-link-init** attributes appear because they have been set to values other than their default values.

```
switch# show running-config interface FibreChannel 8/0/1
interface FibreChannel 8/0/1
  speed 8gbps
  no isl-r_rdy
  trunk-enable
  shutdown
```

The following example shows Fibre Channel attributes for all Fibre Channel ports. In this case, the **speed** and **fec-enable** attributes are set to their default values for all of the interfaces shown.

```
switch# show running-config interface FibreChannel
interface FibreChannel 3/0/1
  desire-distance 0
  no isl-r_rdy
  trunk-enable
  no shutdown
!
interface FibreChannel 3/0/2
  no isl-r_rdy
  trunk-enable
  no shutdown
!
interface FibreChannel 3/0/3
  no isl-r_rdy
  trunk-enable
  no shutdown
!
      (output truncated)
```

To view the setting of a single attribute on a specific port, regardless of whether the attribute is set to its default value, enter the **show running-config interface FibreChannel *rbridge-id/slot/port*** command.

The following example shows the setting of the speed attribute for port 66/0/1:

```
switch# show running-config interface FibreChannel 66/0/1 speed interface FibreChannel 66/0/1
  speed 16gbps
```

Setting Fibre Channel port speed

This procedure sets the ports speed to 4, 8, or 16 Gbps, or to autonegotiate (the default value).

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **interface FibreChannel *rbridge-id/slot/port*** command for the port on which you want to set the speed.

A configuration submenu prompt appears.

3. Enter the **speed** command followed by the desired speed in Gbps.

The following example sets the port speed to 4 Gbps.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# interface FibreChannel 8/0/1
switch(config-FibreChannel-8/0/1)# speed 4
```

Configuring a Fibre Channel port for trunking

A link can be configured to be part of a trunk group. Two or more links in a port group form a trunk group when they are configured for the same speed, the same distance level, and their link distances are nearly equal.

1. In privileged EXEC mode, enter the **configure terminal** command to enter global configuration mode.
2. Enter the **interface FibreChannel** *rbridge-id/slot/port* command for the desired port.

A configuration submenu prompt appears.

3. Enter the **trunk-enable** command.

The following example configures the link attached to port 4 on RBridge 8 to be part of a trunk group.

```
switch# configure terminal
Entering configuration mode terminal
switch(config)# rbridge-id 8
switch(config-rbridge-id-8)# interface FibreChannel 8/0/4
switch(config-FibreChannel-8/0/4)# trunk-enable
```

Monitoring Fibre Channel ports

To monitor a Fibre Channel port, in privileged EXEC mode, enter the **show interface FibreChannel** *rbridge-id/slot/port* command for the Fibre Channel port you want to monitor. The command output provides lots of information about the various hardware counters associated with the port.

This command has a basic version and a detail version. The basic version of the command provides general port information such as status, identification, and configuration information, along with interrupt statistics, link status counters, and so on, as shown in the following example:

```
switch# show interface FibreChannel 66/0/1

fibrechannel 66/0/1 is down (No_Light). Protocol state is down.
Pluggable media present
LineSpeed Actual:
PortSpeed:                N8Gbps
portDisableReason:        Persistently disabled port
PortId:                    010000
PortIfId:                  43010000
PortWwn:                   20:00:00:27:f8:d4:79:be
Distance:                  normal
FEC:                       Inactive

Last clearing of show interface counters: 00:00:00
Interrupts:                0          Link_failure: 0          Frjt:          0
Unknown:                   0          Loss_of_sync: 0          Fbsy:          0
Lli:                       0          Loss_of_sig:  0
Proc_rqrd:                 0          Protocol_err: 0
Timed_out:                 0          Invalid_word: 0
Rx_flushed:                0          Invalid_crc:  0
Tx_unavail:                0          Delim_err:   0
Free_buffer:                0          Address_err:  0
Overrun:                   0          Lr_in:       0
```

```

Suspended:      0          Lr_out:      0
Parity_err:     0          Ols_in:     0
2_parity_err:  0          Ols_out:   0
CMI_bus_err:    0

```

```

Rate info:
  Bandwidth:      8.00G
  Tx performance: 0 B/sec
  Rx performance: 0 B/sec

```

The detailed version of the command, illustrated below, tells you how much traffic has been transmitted or received, and how many times certain error conditions have occurred. Specifically, the **tim_txcrd_z** counters tell you how many times the port was unable to transmit frames because the transmit BB credit was zero. A number greater than zero indicates either that there is congestion on the port or that a device is affected by latency. A larger number indicates a greater problem. A sample is taken every 2.5 microseconds.

```

switch# show interface FibreChannel 66/0/1 detail

fibrechannel 66/0/1 is down (No_Light).  Protocol state is down.
Pluggable media present
LineSpeed Actual:
PortSpeed:                N8Gbps
portDisableReason:        Persistently disabled port
PortId:                    010000
PortIfId:                  43010000
PortWwn:                   20:00:00:27:f8:d4:79:be
Distance:                  normal
FEC:                       Inactive

Last clearing of show interface counters: 00:00:00
Rx Statistics:
  stat_wrx                 0          4-byte words received
  stat_frx                 0          Frames received
Tx Statistics:
  stat_wtx                 0          4-byte words transmitted
  stat_ftx                 0          Frames transmitted
Error Statistics:
  er_enc_in                0          Encoding errors inside of frames
  er_crc                   0          Frames with CRC errors
  er_trunc                  0          Frames shorter than minimum
  er_toolong                0          Frames longer than maximum
  er_bad_eof                0          Frames with bad end-of-frame
  er_enc_out                0          Encoding error outside of frames
  er_bad_os                 0          Invalid ordered set
  er_crc_good_eof          0          Crc error with good eof

Port Error Info:
  loss_of_sync:            0
  lossofsig:               0
  frjt:                    0
  fbsy:                    0

Buffer Information:
  Lx      Max/Resv   Buffer  Needed   Link   Remaining
  Mode    Buffers    Usage  Buffers  Distance Buffers
=====
-         8         0      0      -      0

Rate info:
  Bandwidth:      8.00G
  Tx performance: 0 B/sec
  Rx performance: 0 B/sec

```


Using Access Gateway

- Access Gateway basic concepts..... 165
- Enabling Access Gateway mode.....175
- Disabling Access Gateway mode..... 175
- Displaying Access Gateway configuration data..... 176
- VF_Port to N_Port mapping.....178
- Port Grouping policy..... 182
- Trunking in Access Gateway mode..... 188
- Access Gateway under FlexPort..... 189
- N_Port monitoring for unreliable links..... 190
- Displaying Access Gateway N_Port utilization data 191

Access Gateway basic concepts

On supported switches, the Access Gateway (AG) feature enables you to configure FC ports as N_Ports and to map specific VF ports to these N_Ports. This allows direct connection of hosts attached to the VF_Ports on the AG-supported switch with F_Ports on a Fibre Channel fabric edge switch instead of through ISL connections to a Fibre Channel Router (FCR). These connections can be regular or long distance.

NOTE

In this document, VCS native (or native) mode refers to a switch enabled in VCS mode. Access Gateway mode (or AG switch) refers to a switch in VCS mode enabled for the Access Gateway feature.

Through the use of N_Ports for direct connection to FC switches and VF_Port to N_Port mapping, Access Gateway provides the following benefits:

- As ISLs between switches and FCRs utilize possibly limited domain IDs to identify switches, direct connection from AG-switch N_Ports to Fibre Channel switch F_Ports can resolve scalability issues, as the number of Fibre Channel and VCS fabrics grow.
- Direct connection from AG-switch N_Ports to F_Ports allows greater interoperability with multivendor Fibre Channel fabrics, as connection to these fabrics through an FCR is limited.
- The use of N_Ports instead of ISL connections to FCRs increases the number of device ports available for FCoE hosts and devices behind LAG-supported FSBs connected to the AG-switch VF_Ports. In addition, through use of N_Port ID Virtualization (NPIV), multiple FCoE initiators can access the SAN through the same physical port.

After you configure a switch in AG mode, all FC Ports are enabled as N_Ports. These ports connect to F_Ports on the FC fabric. If the AG switch is connected to a FC switch, the connected N_Ports should come up automatically. Devices attached to VF_Ports come up when the **fcoeport default** command is executed on the individual switch interface port.

Switches supported for Access Gateway

For Network OS 5.0.0, the following Brocade switches support Access Gateway:


- VDX 6740 (but not VDX 6740T or VDX 6740T-1G)
- VDX 2740

Network diagrams

The following diagrams illustrate various connection configurations among switches and network elements, with and without Access Gateway.

The following figure illustrates the connection of eight hosts through an AG switch to a Fibre Channel fabric edge switch.

FIGURE 30 Hosts connecting to FC fabric through switch in AG mode




NOTE

An AG switch can connect to only one Fibre Channel SAN. Ports on this switch connecting to a second FC SAN are disabled. Multiple AG switches, each belonging to a different VCS cluster, can connect to the same SAN fabric.

The following figure illustrates an alternate connection of hosts (servers) to an FC fabric through a switch not in Access Gateway mode. A VDX FC port, configured as an E_Port, connects to the FC SAN through an ISL connection to an FC router. For more information on configuring FC ports for an ISL and FC router connection, refer to [Configuring Fibre Channel Ports](#) on page 157.

FIGURE 31 Connecting Network OS fabric to FC fabric without AG mode



Switches in AG mode are logically transparent to the host and the fabric. Therefore, you can increase the number of hosts that have access to the fabric without increasing the number of switch domains. This simplifies configuration and management in a large fabric by reducing the number of domain IDs and ports.


VCS mode must be enabled to enable Access Gateway on the switch. By default, both VCS mode and AG are enabled.

While in AG mode, the switch functions both as a VCS and AG switch as follows:

- It supports N_Ports, VF_Ports, and Layer 2 interfaces.
- It can connect devices to a VCS fabric through Ethernet ports. VCS fabric services run on Ethernet ports which function under the "native" VCS switch configuration.
- It can connect hosts to a Fibre Channel switch through VF_Ports mapped to N_Ports.

The following figure illustrates the connection of FCoE hosts to two AG switches in a top-of-rack configuration. N_Ports on these switches connect to F_Ports on Brocade DCX Backbones in the FC SAN. In addition, Layer 2 interfaces on these switches connect to other VDX switches in the VCS fabric cluster.

FIGURE 32 Using Access Gateway to connect FC and VCS fabrics



Access Gateway and native VCS modes

In native VCS mode, the switch can function as part of a VCS Fabric cluster, but cannot connect to a FC fabric through N_Ports. When enabled in AG mode, the switch can still function as part of a VCS Fabric cluster, but can now connect directly with a FC "edge" fabric switch through F_Port to N_Port connections.

All VCS and Network OS features are available to a switch in AG mode, which is enabled by default. If you need to enable AG mode (for example, if upgrading from a legacy version), use the **ag enable** command, which initiates a system reboot.

If the switch is to be used as a FC/FCoE switch, you might need to disable Access Gateway mode and return to native mode, using the **no ag enable** command.

For more information enabling and disabling AG mode, refer to [Enabling Access Gateway mode](#) on page 175 and [Disabling Access Gateway mode](#) on page 175.

Access Gateway in a logical chassis cluster

Although operations of a VDX switch configured in Access Gateway mode are similar to a node configured in native VCS mode while in Logical Chassis Cluster mode, there are some unique considerations that you should be aware of.

Nodes in a logical chassis cluster are configured in Logical Chassis Cluster mode. In this mode, both the data and configuration paths are distributed to other nodes in the cluster. The entire cluster is configured from the principal node. For more information, refer to [Logical chassis cluster mode](#) on page 38.

Behavior and operations of an Access Gateway node configured in a logical chassis cluster, such as removing the node from or adding the node to the cluster, are similar to operations in other nodes in the cluster. Access Gateway does not have any global configuration, so the default configuration for the node in Access Gateway mode is similar to a cluster node in native VCS mode. Failover in the cluster, both controlled and uncontrolled, also occurs the same as for an AG switch.

There are two methods for adding an Access Gateway node to the cluster:

- You can enable Access Gateway on a standalone switch, then add the switch to the cluster. Before adding the switch, enable it for Access Gateway mode and configure Access Gateway features and policies through the Network OS AG commands. The node you add to the cluster must have same VCS ID as the cluster; otherwise the configuration on the node will revert to default configuration.
- You can enable Access Gateway on a switch that is already a node in the cluster by just enabling AG on the switch. The switch will reboot and join the cluster enabled in AG mode.

Access Gateway ports

Access Gateway supports VF_Ports that connect host systems to the switch, Layer 2 interface ports that connect the switch to the VCS Fabric cluster, and FC N_Ports that connect the switch to a FC fabric.

In order for hosts attached to VF_Ports on a VDX switch to connect with F_Ports on a FC switch, the VF_Ports must be mapped to the VDX switch N_Ports. Enabling AG mode configures all FC ports as N_Ports and maps VF_Ports to the N_Ports in a sequential, round-robin fashion. This allows for even distribution of logins in a typical configuration, where VF_Ports are sequentially allocated as ENodes log in. You can change this default mapping by mapping any VF_Port to an N_Port using Network OS commands, as described in [Configuring port mapping](#) on page 180.

The following types of ports are supported on a VDX switch in AG mode:

- N_Ports (node ports)—Connects a switch in AG mode to the F_Port on the Fibre Channel switch. By using FlexPort commands you can enable the following numbers of N_Ports on supported switches:
 - VDX 6740 — 32 N_Ports
 - VDX 2740 — 14 N_Ports

NOTE

In Network OS commands, N_Ports are designated by the format rbridge-id/slot/port_number. Therefore, 5/0/1 designates RBridge 5/slot 0/port 1.

- VF_Ports (Virtual Fabric ports)—For connection of FCoE hosts and devices behind LAG-supported FSB devices. You can map these ports to specific N_Ports for connection to a Fibre Channel switch. Consider the following specifications:
 - By default, each switch is assigned 64 VF_Ports.
 - There is no limit to the number of VF_Ports that you can map to an N_Port.
 - Up to 64 NPIV logins are allowed per VF_Port.

Default port numbers are specific to the switch platform:

- VDX 6740 — Valid VF_Ports are 0 to 1000.

- VDX 2740 – Valid VF_Ports are 0 to 1000.

NOTE

In Network OS commands, VF_Ports are designated by the format domain/rbridge-id/VF_Port. For example, 1/2/26 designates that VF_Port 26 resides in domain 1 and RBridge 2.

- L2 interface ports—Ethernet TRILL ports that connect with other VDX switches in the Brocade VCS Fabric cluster. Port numbers are specific to the switch platform:
 - VDX 6740 – Valid TRILL ports are (total number of physical ports) – (ports converted into Fibre Channel ports).
 - VDX 2740 – Valid TRILL ports are (total number of physical ports) – (ports converted into Fibre Channel ports).

Since all VDX switch FC ports are enabled as N_Ports in Access Gateway mode, FC hosts or targets cannot directly attach to the AG switch. When Access Gateway mode is enabled, you can configure additional FC port attributes for the N_Ports as you would on non-AG switches. If you are upgrading to Network OS 5.0.0 (rather than a clean install), refer to the *Network OS Upgrade Guide*.

Transitioning from native VCS to AG mode

Be aware of the following interface and port functions after enabling AG mode:

- The system reboots. Native VCS configuration is retained.
- CEE interfaces will be in a no-shutdown state.
- If FC ports are connected to FC switch F_Ports, connected N_Ports should come up automatically. Devices connected to mapped VF_Ports should come up after you enter the **fcoeport default** command on the interface port.
- VDX switch Ethernet ports are under the native VCS switch configuration.
- For default VF_Port to N_Port mapping, VF_Ports are mapped to N_Ports sequentially in a round-robin fashion as ENodes log in. Mapping that you implement through the **map fport interface fcoe port** command overrides the default mapping.
- VCS Fabric services run on VCS ports and not under the Access Gateway daemon.

Comparison of Access Gateway, ISL, and FC switch ports

A switch in Access Gateway (AG) mode uses VF_Ports and N_Ports to connect devices to the Fibre Channel (FC) switch. The connected FC switch connects to the AG switch N_Ports through F_Ports and presents a variety of ports for connection to FC fabric devices.

Access Gateway (AG) multiplexes host connections to the fabric. AG presents a VF_Port to a FCoE host and an N_Port to an edge Fibre Channel fabric switch. Multiple VF_Ports mapped to N_Ports provide multiple device ports for connection to the FC fabric.

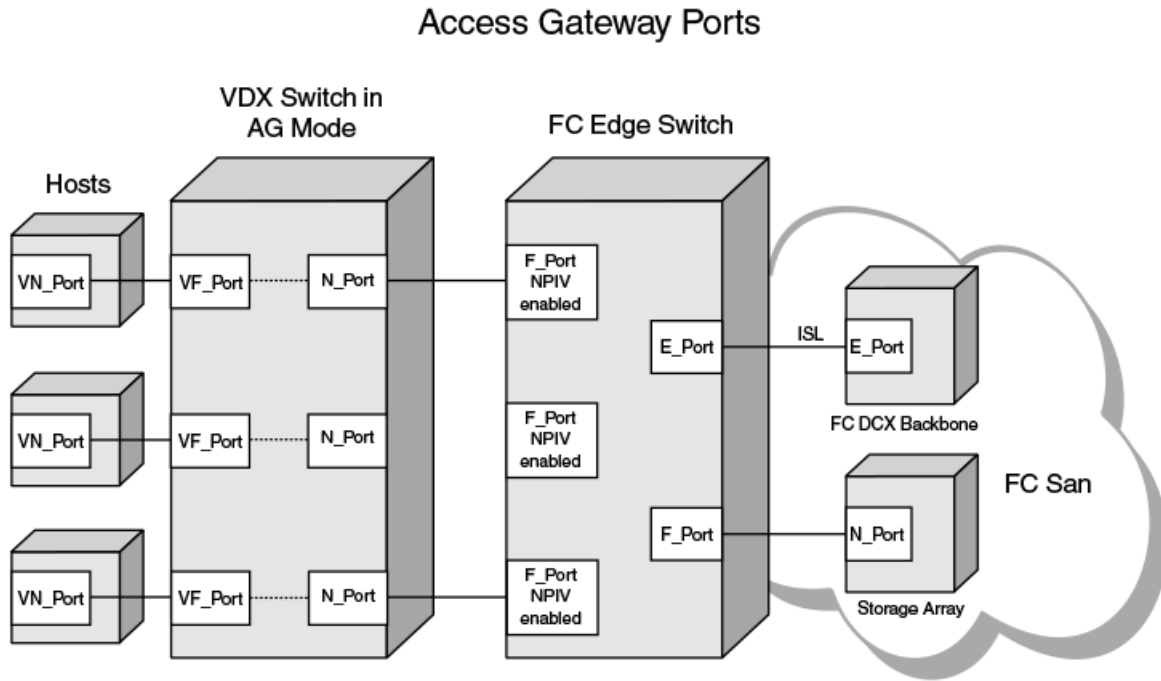
Using N_Port ID Virtualization (NPIV), AG allows multiple FCoE initiators to access the SAN on the same physical port. This reduces the hardware requirements and management overhead of hosts to the SAN connections.

In contrast to the AG switch, the connected FC switch presents F_Ports (or FL_Ports) to storage devices hosts and presents E_Ports, VE_Ports, or EX_Ports to other switches in the fabric.

A native switch using an ISL connection between its FC E_Port and an EX_Port on an FCR consumes domain ID resources that may impact scalability as VCS and FC fabrics grow. In addition, connection through a FCR may limit connection to multivendor FC fabrics. Finally, connection through an ISL provides limited device port connections to the FC fabric. For more information on configuring FC ports for connection to an FCR and FC fabric, refer to [Configuring Fibre Channel Ports](#) on page 157.

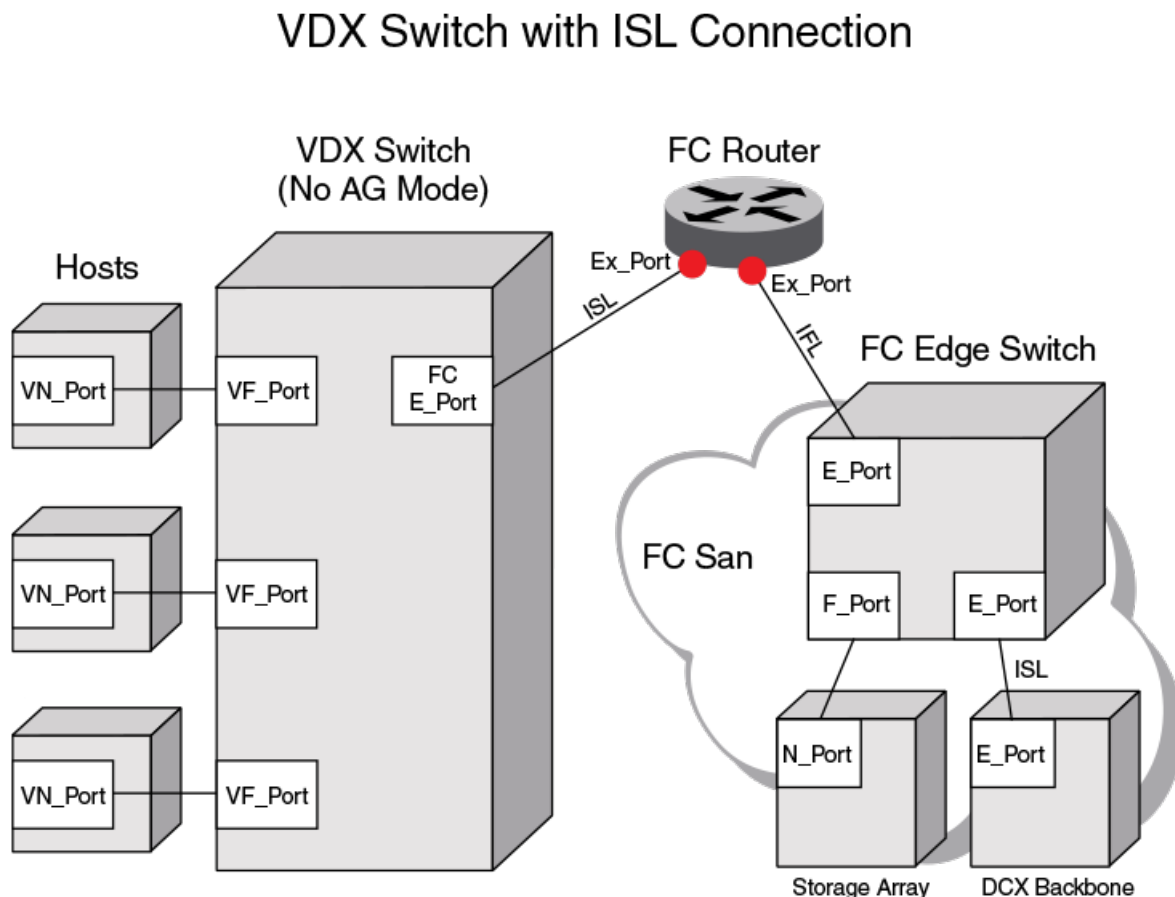
The following figure illustrates ports used for connecting hosts attached to an Access Gateway Switch to a Fibre Channel switch.

FIGURE 33 Access Gateway and FC switch ports



The following figure illustrates the ports used for connecting hosts attached to a non-AG (native) switch to a Fibre Channel switch, using an ISL connection.

FIGURE 34 Connection of native VCS and FC ports



Access Gateway features and requirements

Although Access Gateway provides standard features for connection to Fibre Channel SANs, you can configure a number of optional features as well. There are also requirements and limitations that you should be aware of when using this feature in a VCS cluster and FC fabric environment.

Port grouping

The Port Grouping (PG) policy is enabled by default when AG is enabled. This allows you to group N_Ports into a port group. By default, any VF_Ports mapped to these N_Ports are also members of that port group. Port Grouping allows you to isolate specific hosts to specific FC fabric ports for performance, security, or other reasons.

Automatic Login Balancing (LB) and Modified Managed Fabric Name Monitoring (M-MFNM) modes are enabled by default when the PG policy is enabled.

- When LB mode is enabled and an N_Port goes offline, existing logins from VF_Ports that are mapped to the offline N_Port are distributed to available N_Ports in the port group. If a new N_Port comes online, existing logins are not disturbed. LB mode can be disabled using Network OS commands.

- When LB mode is disabled, VF_Ports are not shared among N_Ports in the port group, as VF_Ports can only connect to N_Ports to which they are mapped. As a best practice to ensure device login, bind the ENode to a VF_Port and ensure that its mapped N_Port is online.
- M-MFNM mode ensures all N_Ports connect to the same FC fabric, preventing connections to multiple SANs. M-MFNM cannot be disabled as long as LB mode is enabled.

For more information on Port Grouping policy modes, refer to [Port Grouping policy modes](#) on page 186.

N_Port monitoring for unreliable links

The N_Port monitoring for unreliable links feature monitors links from all N_Ports on the VDX switch to F_Ports on the FC fabric. If online and offline static change notifications (SCNs) exceed a set threshold during a specific time period, the link is considered unreliable, and the N_Port is taken offline. The VF_Ports mapped to the N_Port also go offline. If the N_Port is in a port group and Automatic Login Balancing is enabled, the VF_Ports mapped to the N_Port are distributed among available N_Ports in the same port group.

Additional features and functions

Following are additional features and functions of Access Gateway:

- Access Gateway enables VDX FC ports as N_Ports. Hosts attached to VDX VF_Ports can connect directly with F_Ports on a Fibre Channel fabric edge switch through these N_Ports.
- Instead of using ISL connections and possibly limited domain resources, the use of N_Ports increases the number of available device ports on the switch. As the number of Fibre Channel and VCS Fabrics grow, scalability is less of an issue.
- Through the use of N_Port ID Virtualization (NPIV), multiple FCoE initiators can access the SAN through the same physical port.
- When enabling AG mode, VF_Ports are mapped to available N_Ports in a round-robin fashion as ENodes log in. However, you can re-map any VF_Port to switch N_Ports.
- Access Gateway can operate in both fabric cluster and logical chassis cluster modes.
- You can configure additional FC port attributes for the AG switch N_Ports as you would on non-AG switches. Refer to [Configuring Fibre Channel Ports](#) on page 157

Limitations

Following are limitations you should be aware of when using Access Gateway mode:

- Hosts connected to an Access Gateway switch cannot communicate with targets on the VCS Fabric.
- A VDX switch configured for Access Gateway can connect with only one FC Fabric. Ports connected to a second FC fabric are disabled.
- Access Gateway can operate in both fabric cluster mode and logical chassis cluster modes. The AG configuration is not distributed in fabric cluster mode and is distributed in logical chassis cluster mode.
- You can only configure VF_Port to N_Port mapping for devices directly attached to VF_Ports and F_Ports on the connected FC switch. These mappings control device logins through appropriate N_Ports.
- Since all switch FC ports are configured as N_Ports when AG mode is enabled:
 - FC hosts or targets cannot be directly attached to the AG switch.
 - The AG switch cannot be connected to a Fabric OS Access Gateway in a Cascaded configuration.
- Access Gateway does not "bridge" the VCS and FC fabrics:
 - Hosts connected to VF_Ports mapped to Access Gateway N_Ports appear on the FC fabric only.
 - Device FC IDs are assigned by the FC fabric F_Ports connected to the Access Gateway N_Ports.

- VF_Ports and N_Ports are under the Access Gateway daemon's configuration.
- Fibre Channel OS components, such as management server, name server, and zoning are restricted on Network OS Access Gateways just as they are on Fabric OS Access Gateways. Refer to the *Fabric OS Access Gateway Administrator's Guide* for a complete list.
- Although the **show fcoe login** command displays FCoE devices connected to the Access Gateway switch VF_Ports, these devices are in the FC fabric and cannot be detected by the VCS Fabric name server. Therefore, these devices cannot be zoned in a VCS Fabric.

FCoE and Layer 2 support and limitations

The following functions are supported:

- The following functionality is supported for a configuration consisting of a vLAG from a host CNA to two Access Gateway switches or to an AG switch and a VCS Switch (L2 vLAG for top of rack split):
 - The vLAG links can carry Layer 2 and Layer 3 traffic.
 - VLAG support is identical to support in native VCS mode.
 - A separate FCoE device login is supported through each AG switch.
- The following functionality is supported for a configuration with a LAG from an FSB to an AG switch:
 - LAG specifics, such as number of links and contiguous vs. discontinuous, is identical to native VCS support.
 - Multiple LAGs can connect to the AG switch (one per FSB).
 - LAG carries Layer 2 and Layer 3 traffic.
 - Devices connected via an FSB LAG cannot talk to a Cisco SAN.
 - LAGs and direct attached devices are supported on the same AG switch.
- The following functionality is supported for VF_Ports and CEE interfaces:
 - VF_Ports are dynamically bound to Ethernet interfaces as in native VCS mode.
 - As in native VCS mode, all CEE interfaces with connected devices must be configured for FCoE.
 - The CEE interface will come up as an ISL ET port if it is connected to a peer ET port on another switch in the VCS Fabric.
 - As in native VCS mode, 64 VF_Ports are allocated by default.
 - A VF_Port can accept up to 64 NPIV logins.
 - As in native VCS mode, VF_Ports are dynamically allocated as devices come up.
 - As in native VCS mode, VF_Ports can be statically bound to ENodes.
 - VCS Fabric services run on VCS ports and not under Access Gateway.
 - As in native VCS mode, the number of VF_Ports allocated can be changed dynamically:
 - > You can configure the maximum number of FCoE devices that can be logged into a switch by using the **fcoe_enodes** command.
 - > Newly allocated VF_Ports are mapped to existing N_Ports sequentially in a round-robin fashion, which assigns all VF_Ports sequentially and evenly to the N_Ports.
 - > Newly deallocated VF_Ports are removed from existing VF_Port to N_Port mappings.

Following are support limitations:

- If an interface is only handling L2 traffic, the corresponding VF_Port appears as disabled to AG.
- For vLAGs:
 - As in native VCS mode, a vLAG with FSB cannot support FCoE traffic. It can support L2 traffic only.
 - A vLAG from a host Converged Network Adapter (CNA) supports L2 and FCoE traffic.
- A single LAG/link is supported between one FSB and one AG switch. Subsequent LAG/links are treated as a TRILL loop and disabled.
- LAG member VF_Ports cannot be mapped to individual N_Ports or N_Port groups.

Enabling Access Gateway mode

Enabling Access Gateway (AG) mode on a VDX switch allows FCoE hosts and devices behind LAG-supported FSBs connected to VF_Ports to connect to a FC fabric.

NOTE

On supported devices, Access Gateway mode is enabled by default.

Enabling AG mode enables FC ports on the switch, configuring them as N_Ports. The N_Ports can connect directly to F_Ports on an edge FC switch. VF_Ports are mapped to N_Ports in a sequential, round-robin fashion as Enodes log in. You can change this default mapping using Network OS commands.

NOTE

Enabling AG mode is disruptive since the switch disables and reboots. If the switch is part of a logical cluster, you should back up the configurations before enabling AG mode.

Use the following procedure to enable Access Gateway mode.

1. Enter the the **ag enable** command.

```
switch# ag enable
```

The switch reboots and AG mode is enabled. Switch FC ports are automatically enabled as N_Ports and mapped to VF_Ports.

2. You can configure additional FC port attributes for the N_Ports as you would on switches in native mode. Refer to [Configuring Fibre Channel Ports](#) on page 157.

Disabling Access Gateway mode

Disabling Access Gateway (AG) mode returns the switch to native VCS mode. Disabling AG mode also removes all AG configuration, including port mapping and N_Port configuration.

Access Gateway mode must be enabled before you can disable it.

Use the following procedure to disable AG mode on the switch.

1. Display the current AG state and configuration by entering the **show ag rbridge-id rbridge-id** command in privileged EXEC mode.

```
switch# show ag rbridge-id 1
```

If Access Gateway configuration data displays, as shown in [Displaying Access Gateway configuration data](#) on page 176, AG is enabled.

If "AG mode not set" displays as shown in the following example, AG is not enabled.

```
switch# show ag rbridge-id 2
AG mode not set.
```

For more information on displaying the current AG configuration on a switch or all AG switches in a VCS cluster using this command, refer to [Displaying Access Gateway configuration data](#) on page 176.

2. Disable AG mode by entering the **no ag enable** command while in privileged EXEC mode:

```
switch# no ag enable
```

Displaying Access Gateway configuration data

Use the **show running-config rbridge-id *rbridge id* ag** and **show ag rbridge-id *rbridge-id*** commands to display Access Gateway configuration data.

To display AG configuration data, use the following methods:

- Use the **show running-config rbridge-id *rbridge-id* ag** command to display the configured N_Port to VF_Port mappings, port grouping information, and other parameters. This shows the factory-default configuration, unless parameters have been modified by the user.
- Use the **show ag rbridge-id *rbridge-id*** command to display the current and active status of AG configuration, such as the switch identification, number and type of ports, enabled policies, port grouping, and attached fabric details. This displays only ports that are currently online and current mappings. For example, this will show VF_Ports that have failed over to an N_Port if an N_Port that has gone offline.

NOTE

Display of current, active mapping, or configured mapping for a port group using the **show ag rbridge-id *rbridge-id*** and **show running-config rbridge-id *rbridge-id* ag** commands depends on the enabled or disabled state of Login Balancing mode. For more information, refer to [Automatic Login Balancing mode](#) on page 186.

1. Make sure you are in Privileged EXEC mode and have a switch prompt such as the following.

```
switch#
```


2. Perform one of the following steps:

- Enter the **show running-config rbridge-id *rbridge-id* ag** command as in the following example for RBridge 1.

```
switch# show running-config rbridge-id 1 ag
```

- Enter the **show ag rbridge-id *rbridge-id*** command as shown in the following example for RBridge 5.

```
switch# show ag rbridge-id 5
```

If Access Gateway is enabled, data such as the following displays for **show running-config rbridge-id *rbridge-id* ag**, as shown in the following example for RBridge 1:

```
switch# show running-config rbridge-id 1 ag
rbridge-id 1
ag
  nport 1/0/1
    map fport interface Fcoe 1/1/1 1/1/9 1/1/17 1/1/25 1/1/33 1/1/41 1/1/49 1/1/57
    !
  nport 1/0/2
    map fport interface Fcoe 1/1/2 1/1/10 1/1/18 1/1/26 1/1/34 1/1/42 1/1/50 1/1/58
    !
  nport 1/0/3
    map fport interface Fcoe 1/1/3 1/1/11 1/1/19 1/1/27 1/1/35 1/1/43 1/1/51 1/1/59
    !
  nport 1/0/4
    map fport interface Fcoe 1/1/4 1/1/12 1/1/20 1/1/28 1/1/36 1/1/44 1/1/52 1/1/60
    !
  nport 1/0/5
    map fport interface Fcoe 1/1/5 1/1/13 1/1/21 1/1/29 1/1/37 1/1/45 1/1/53 1/1/61
    !
  nport 1/0/6
    map fport interface Fcoe 1/1/6 1/1/14 1/1/22 1/1/30 1/1/38 1/1/46 1/1/54 1/1/62
    !
  nport 1/0/7
    map fport interface Fcoe 1/1/7 1/1/15 1/1/23 1/1/31 1/1/39 1/1/47 1/1/55 1/1/63
    !
  nport 1/0/8
    map fport interface Fcoe 1/1/8 1/1/16 1/1/24 1/1/32 1/1/40 1/1/48 1/1/56 1/1/64
    !
pg 0
  nport interface FibreChannel 1/0/1 1/0/2 1/0/3 1/0/4 1/0/5 1/0/6 1/0/7 1/0/8
  modes lb
  rename pg0
  !
  timeout fnm 120
  counter reliability 25
```

If Access Gateway is enabled, data such as the following displays for **show ag rbridge-id *rbridge-id***, as shown in the following example for RBridge 5:

```
switch# show ag rbridge-id 5
RBridge-ID 5:
-----
Name           : sw0
NodeName       : 10:00:00:05:33:f4:78:04
Number of Ports : 32
IP Address(es) : 10.37.209.80
Firmware Version : v4.1.0pgoel_pit02_nos4_1_10_10
Number of N_Ports (Fi) : 2
Number of VF Ports : 0
Policies Enabled : pg
Persistent ALPA : Disabled
Port Group information :
  PG_ID  PG_Name  PG_Mode  PG_Members
-----
  0      pg0    lb       5/0/1, 5/0/2, 5/0/3, 5/0/4,
                    5/0/5, 5/0/6, 5/0/7, 5/0/8
```

```

-----
Fabric Information :
Attached Fabric Name          N_Ports (Fi)
-----
10:00:00:05:33:72:f5:5a      5/0/1, 5/0/2

N_Port (Fi) information :
Port          PortID          Attached PWWN          IP_Addr          VF_Ports
-----
Fi 5/0/1      0x020200      20:02:00:05:33:72:f5:5a  10.37.209.86     None
Fi 5/0/2      0x020300      20:03:00:05:33:72:f5:5a  10.37.209.86     None
-----

VF_Port information :
VF_Port      Eth_Port      PortID          Attached PWWN          N_Port (Fi)
-----
None
-----

```

If AG is not enabled, "AG mode not set" displays, as shown in the following example for RBridge 2:

```

switch# show ag rbridge-id 2
AG mode not set.

```

NOTE

You can also enter **show ag rbridge-id all** to display AG configuration data for all switches in the VCS cluster.

VF_Port to N_Port mapping

To connect hosts attached to VDX Switch VF_Ports to Fibre Channel switch F_Ports, the appropriate VF_Ports must be mapped to VDX Switch N_Ports. Although ports have a factory-default mapping based on the VDX platform, you can change mapping using Network OS commands.

Consider the following when mapping ports:

- You can map multiple VF_Ports to an N_Port. There is no limit to the number of VF_Ports that you can map to an N_Port.
- You can only configure VF_Port to N_Port mapping for devices directly attached to VF_Ports on the VDX switch and F_Ports on the connected FC switch. These mappings control device logins through appropriate N_Ports.
- Consider the N_Port and VF_Port ranges allowed for a VDX platform. Refer to [Access Gateway ports](#) on page 169.
- If an N_Port is removed from a port group enabled for Automatic Login Balancing mode and moved to another port group, the VF_Ports mapped to that N_Port remain with the N_Port. If an N_Port is moved from a port group not enabled for Automatic Login Balancing mode, the VF_Ports that are mapped to the N_Port move to the default Port Group 0.

Displaying port mapping

You can display current and configured VF_Port to N_Port mapping on a specific switch or on all switches enabled for Access Gateway in the VCS cluster.

Display current, active VF_Port to N_Port mapping on a specific switch or on all switches enabled for Access Gateway in the VCS cluster using the **show ag map rbridge-id rbridge-id** command while in Privileged EXEC mode.

Display configured VF_Port to N_Port mapping on a switch using the **show running-config rbridge-id *rbridge id* ag** command.

Perform one of the following steps while in privileged EXEC mode:

- To display current, active, VF_Port mapping for a specific N_Port, enter **show ag map nport rbridge-id *rbridge-id***.

```
switch# show ag map nport interface fiberChannel 200/0/1 rbridge-id 200
```

- To display current, active VF_Port mapping to all N_Ports on a switch, enter **show ag map rbridge-id *rbridge-id***.

```
switch# show ag map rbridge-id 200
```

- To display VF_Port mapping to a N_Ports on all switches in the VCS cluster, enter **show ag map rbridge-id all**.

```
switch# show ag map rbridge-id all
```

- To display configured mapping on a switch, enter **show running-config rbridge-id *rbridge id* ag**.

```
switch# show running-config rbridge-id 1 ag
```

Current and configured mapping display

Display of current, active mapping, or configured mapping for a port group using the **show ag map** and **show running-config rbridge-id rbridge id ag** commands depend on the enabled or disabled state of Login Balancing mode. For more information, refer to [Automatic Login Balancing mode](#) on page 186.

The following example is sample output from the **show ag map rbridge-id rbridge-id** command, which shows current, active port mapping. The "Current_VF_Ports" column shows that there are no VF_Ports online.

```
switch# show ag map rbridge 5
RBridge-ID 5:
-----
N_Port(Fi)    PG_ID  PG_Name  Current_VF_Ports
-----
5/0/1         0      pg0      None
5/0/2         0      pg0      None
5/0/3         0      pg0      None
5/0/4         0      pg0      None
5/0/5         0      pg0      None
5/0/6         0      pg0      None
5/0/7         0      pg0      None
5/0/8         0      pg0      None
-----
```

The following example is sample output from the **show running-config rbridge-id rbridge id ag** command to show configured port mapping. The output lists N_Port numbers on the switch, and then mapped VF_Ports following "map fport interface fcoe."

```
switch# show running-config rbridge-id 1 ag
rbridge-id 1
ag
nport 1/0/1
  map fport interface Fcoe 1/1/1 1/1/9 1/1/17 1/1/25 1/1/33 1/1/41 1/1/49 1/1/57
  !
nport 1/0/2
  map fport interface Fcoe 1/1/2 1/1/10 1/1/18 1/1/26 1/1/34 1/1/42 1/1/50 1/1/58
  !
nport 1/0/3
  map fport interface Fcoe 1/1/3 1/1/11 1/1/19 1/1/27 1/1/35 1/1/43 1/1/51 1/1/59
  !
nport 1/0/4
  map fport interface Fcoe 1/1/4 1/1/12 1/1/20 1/1/28 1/1/36 1/1/44 1/1/52 1/1/60
  !
nport 1/0/5
  map fport interface Fcoe 1/1/5 1/1/13 1/1/21 1/1/29 1/1/37 1/1/45 1/1/53 1/1/61
  !
nport 1/0/6
  map fport interface Fcoe 1/1/6 1/1/14 1/1/22 1/1/30 1/1/38 1/1/46 1/1/54 1/1/62
  !
nport 1/0/7
  map fport interface Fcoe 1/1/7 1/1/15 1/1/23 1/1/31 1/1/39 1/1/47 1/1/55 1/1/63
  !
nport 1/0/8
  map fport interface Fcoe 1/1/8 1/1/16 1/1/24 1/1/32 1/1/40 1/1/48 1/1/56 1/1/64
  !
pg 0
  nport interface FibreChannel 1/0/1 1/0/2 1/0/3 1/0/4 1/0/5 1/0/6 1/0/7 1/0/8
  modes lb
  rename pg0
  !
  timeout fnm 120
  counter reliability 25
```

Configuring port mapping

When operating in Access Gateway mode, you can specify routes that AG will use to direct traffic from the devices (hosts or targets) on its VF_Ports to the ports connected to the fabric using its N_Ports. The process of specifying routes is called "mapping." When AG is

enabled on a switch, VF_Ports are assigned to available N_Ports in a round-robin fashion as ENodes log in. You can change this mapping using the following instructions.

Use the **map fport interface fcoe *port*** command to map specific VF_Ports to a an N_Port to ensure that all traffic from these VF_Ports always goes through the same N_Port. You must enter this command while in N_Port configuration mode for a specific N_Port. All VF_Ports mapped to an N_Port in an N_Port group will be part of that port group.

Remember the following points when mapping ports:

- The range of valid VF_Ports and N_Ports is specific to the VDX platform. Refer to [Access Gateway ports](#) on page 169 for valid port numbers.
- Newly allocated VF_Ports are mapped to existing N_Ports in a round-robin fashion.
- Newly deallocated VF_Ports are removed from existing mappings.
- If the AG switch is connected to a FC switch, the connected N_Port and devices on the mapped VF_Ports should come online automatically.

Use the following steps to configure VF_Port to N_Port mapping:

1. Perform steps under [Displaying port mapping](#) on page 178 to display current and configured port mapping.
2. Enter the **configure** command to access global configuration mode.

```
switch# configure
```

3. Enter the **rbridge-id *id*** command to enter RBridge ID mode for the specific switch.

```
switch(config)# rbridge-id 2
```

4. Enter the **ag** command to enter Access Gateway configuration mode.

```
switch(config-rbridge-id-2) # ag
```

5. Enter the **nport *port*** command for the N_Port where you want to change or set mapping to a VF_Port, where *nport* is the N_Port number in rbridge-id/slot/port format. This accesses the configuration mode for the N_Port.

```
switch(config-rbridge-id-2-ag) #
nport interface FiberChannel 2/0/4
```

6. Perform one of the following steps:

- To map a VF_Port to the N_Port, enter **map fport interface fcoe *port***, where *port* is the VF_Port in domain/rbridge-id/port format.

```
switch(config-rbridge-id-2-ag-nport-if-fi-
2/0/4) # map fport interface fcoe 1/2/26
```

- To remove a VF_Port mapped to the N_Port, enter the **no map fport interface fcoe *port*** command, where *port* is the VF_Port number in domain/rbridge-id/port format.

```
switch(config-rbridge-id-2-ag-nport-if-fi-
2/0/4) # no map fport interface fcoe 1/2/26
```

7. Return to privileged EXEC mode and enter the **show running-config rbridge-id *rbridge id* ag** command to verify configured VF_Port to N_Port mapping. Refer to [Displaying port mapping](#) on page 178 for more information.

```
switch# show running-config rbridge-id 2 ag
```

Port Grouping policy

The Port Grouping (PG) policy partitions the VF_Ports, host, target ports within an Access Gateway-enabled switch into independently operated groups. Port Grouping allows you to dedicate specific hosts to specific fabric ports for performance, security, or other reasons.

Port Grouping policy is enabled by default when you enable Access Gateway mode and cannot be disabled.

To create port groups, you group N_Ports under a specific port group ID. By default, any VF_Ports mapped to the N_Ports belonging to a port group are members of that port group. All N_Ports in the group are shared by all VF_Ports mapped to those N_Ports. ENodes can log in as long as an online N_Port exists in the group.

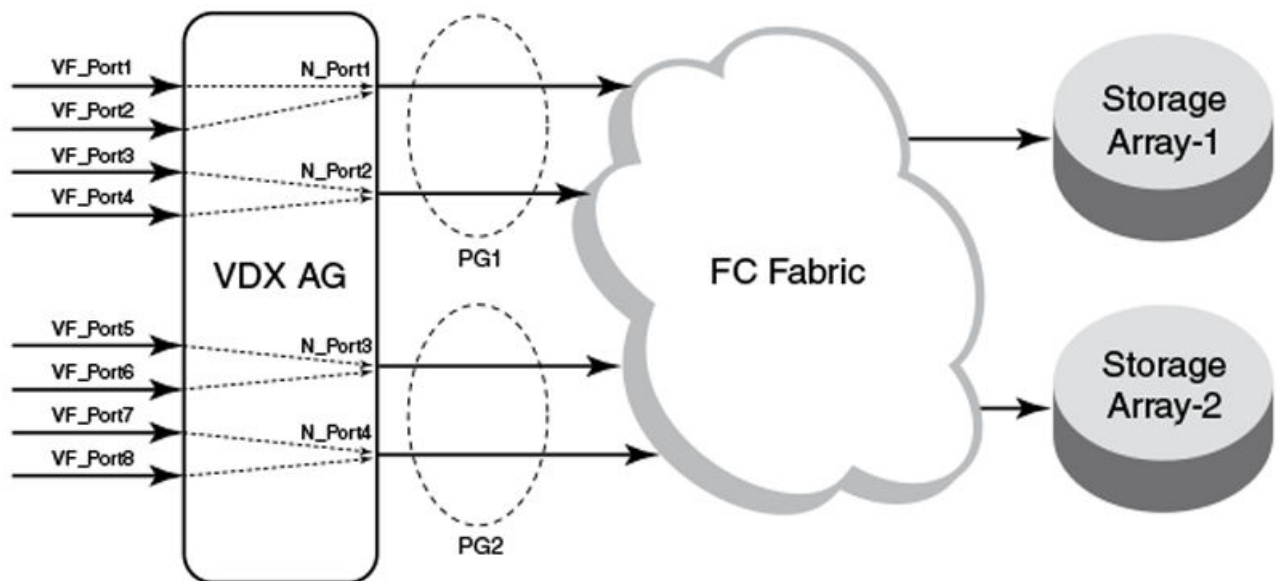
NOTE

In Network OS commands, N_Ports are designated by the format rbridge-id/port group/N_Port. For example, 5/0/1 designates RBridge 5/port group 1/N_Port 1.

When Access Gateway mode is enabled, a default port group 0 (pg 0) is created that contains all N_Ports on the switch. The maximum number of port groups supported is 16, including pg 0.

The following figure illustrates two port groups connecting VF_Ports to an FC fabric. Ports in PG 1 are connecting to one storage array, while ports in PG2 are connecting to a different storage array.

FIGURE 35 Port groups connecting to FC fabric



Following are considerations and limitations for the Port Grouping policy.

- An ENode can log in
- A port cannot be a member of more than one port group.
- The PG policy is enabled by default in when you enable AG mode. A default port group "0" (PG0) is created, which contains all N_Ports and mapped VF_Ports on the switch.
- If an N_Port is added to a port group or deleted from a port group, it maintains its original mapping configuration. If an N_Port is deleted from a port group, it is automatically added to port group 0.

Displaying port grouping information

Display information for N_Port groups configured on the switch or all switches in the VCS cluster enabled for Access Gateway mode. Access Gateway must be enabled for this command to succeed.

Use the **show ag pg rbridge-id** *rbridge id* command while in Privileged EXEC mode to display information on N_Port groups configured on a switch. This information includes N_Ports and VF_Ports in the group and enabled PG modes.

1. Configure port groups using steps under [Creating and removing port groups](#) on page 183.
2. Perform one of the following steps while in Privileged EXEC mode:
 - To display the current information for port groups on a specific switch, enter **show ag pg rbridge-id** *rbridge-id*.

```
switch# show ag pg rbridge-id 5
```

- To display port grouping information for port groups on all Access Gateway switches in the VCS cluster, enter **show ag pg rbridge-id all**.

```
switch# show ag pg rbridge-id all
```

- To display current information on a specific port group (such as pg 11), enter **show ag pg pgid** *rbridge-id* *rbridge-id*.

```
switch# show ag pg pgid 11 rbridge-id 200
```

The following is an example of command output for RBridge 5:

```
switch# show ag pg rbridge-id 5
Rbridge-ID 5:
-----
PG_ID  PG_Name  PG_Mode  N_Ports (Fi)          VF_Ports
-----
0      pg0        lb       5/0/1, 5/0/2, 5/0/3, 5/0/4,  1/5/1, 1/5/2, 1/5/3, 1/5/4,
                    5/0/5, 5/0/6, 5/0/7, 5/0/8  1/5/5, 1/5/6, 1/5/7, 1/5/8,
                                      1/5/9, 1/5/10, 1/5/11, 1/5/12,
                                      1/5/13, 1/5/14, 1/5/15, 1/5/16,
                                      1/5/17, 1/5/18, 1/5/19, 1/5/20,
                                      1/5/21, 1/5/22, 1/5/23, 1/5/24,
                                      1/5/25, 1/5/26, 1/5/27, 1/5/28,
                                      1/5/29, 1/5/30, 1/5/31, 1/5/32,
                                      1/5/33, 1/5/34, 1/5/35, 1/5/36,
                                      1/5/37, 1/5/38, 1/5/39, 1/5/40,
                                      1/5/41, 1/5/42, 1/5/43, 1/5/44,
                                      1/5/45, 1/5/46, 1/5/47, 1/5/48,
                                      1/5/49, 1/5/50, 1/5/51, 1/5/52,
                                      1/5/53, 1/5/54, 1/5/55, 1/5/56,
                                      1/5/57, 1/5/58, 1/5/59, 1/5/60,
                                      1/5/61, 1/5/62, 1/5/63, 1/5/64
-----
```

Creating and removing port groups

You must create a port group with a unique ID before adding N_Ports to the group or enabling Port Grouping (PG) policy modes. Removing a port group removes all N_Ports, mapped VF_Ports, and associated PG modes.

Access Gateway must be enabled for the **pg** command to succeed.

When you enable Access Gateway mode, all ports belong to port group 0 (pg0). You can move ports to a separate port group by first creating a port group with a unique ID, and then adding N_Ports to that group. All VF_Ports mapped to added N_Ports also become members of the new group. Removing a port group removes all N_Ports, mapped VF_Ports, and associated PG modes.

Use the **pg *pgid*** command to configure a port group with a unique ID (*pgid*). The *pgid* is a number that cannot exceed the number of N_Ports allocated for the switch model, minus 1 for default pg 0. Therefore, for a VDX with 16 N_Ports, a valid *pgid* would be from 1 through 15. Once configured, you can access the port group for configuration tasks, such as adding and removing N_Ports, enabling port group modes, and renaming the group. Use the **no pg *pgid*** command to remove a port group. You enter these commands while in Access Gateway (ag) configuration mode.

NOTE

Port Grouping policy is enabled by default when you enable Access Gateway mode.

1. Enter the **configure** command to enter global configuration mode.

```
switch# configure
```

2. Enter the **rbridge-id *id*** command to enter RBridge ID mode for the specific switch.

```
switch(config)# rbridge-id 3
```

3. Enter the **ag** command to enter Access Gateway configuration mode.

```
switch(config-rbridge-id-3)# ag
```

4. Perform one of the following steps:

- To create a port group, enter the **pg *pgid*** command. The port group ID (*pgid*) must not exceed 64 characters.

```
switch(config-rbridge-id-3-ag)# pg 1
```

- To remove a port group, enter the **no pg *pgid*** command.

```
switch(config-rbridge-id-3-ag)# no pg 1
```

Creating a port group with the **pg *pgid*** command enters the PG configuration mode for the port group ID (*pgid*) so that you can add N_Ports and perform other PG policy configuration.

```
switch(config-rbridge-id-3-ag-pg-1)#
```

5. Verify that the port group was created using the **show ag pg *pgid* rbridge-id *rbridge-id*** command from privileged EXEC configuration mode. Refer to [Displaying port grouping information](#) on page 183 for details.

Naming a port group

You can name or rename a port group and use this name in place of the port group ID.

Access Gateway and the PG policy must be enabled for the **rename** command to succeed.

Use the **rename *pgid*** command while in the configuration mode for a port group to change the port group name. The name cannot exceed 64 characters.

1. Enter the **configure** command to enter global configuration mode.

```
switch# configure
```

2. Enter the **rbridge-id *rbridge-id*** command to enter RBridge ID configuration mode for the specific switch.

```
switch(config)# rbridge-id 3
```

3. Enter the **ag** command to enter Access Gateway configuration mode.

```
switch(config-rbridge-id-3)# ag
```


4. Enter the port group ID, such as pg 1, to enter configuration mode for the port group.

```
switch(config-rbridge-id-3-ag)# pg 1
```

5. Change the name of the port group using the **rename** *pgid* command. In the following example, port group is named pg-array24.

```
switch(config-rbridge-id-3-ag-pg-1)# rename pg-array24
```

The port group name must not exceed 64 characters.

Adding and removing N_Ports in a port group

After creating a port group, you must add N_Ports to the group. You must delete N_Ports from a group if you want to move them to another port group or not include them in a port group.

Access Gateway and the PG policy must be enabled for the **nport interface Fibrechannel** command to succeed.

Use the **nport interface Fibrechannel** *port* command while in command mode for a specific port group to add an N_Port to the group. Use the **no nport interface Fibrechannel** *port* command to remove an N_Port. Before you can add a port to a port group, you must remove it from the port group where it currently exists, unless the port is in port group 0 (pg 0). If you remove a port from a port group, it will default to port group 0. You cannot delete a port from port group 0.

NOTE

Under FlexPort, after you move a VF_Port from one port group to another, it is possible that an N_Port may not be available in the target port group.

1. Determine the port group on the switch where the port is currently a member by entering the **show ag pg rbridge-id** *rbridge-id* while in the privileged EXEC command mode.

```
switch# show ag pg rbridge-id 3
```

2. Enter the **configure terminal** command to enter global configuration mode.

```
switch# configure terminal
```

3. Enter the **ag** command to enter Access Gateway command mode.

```
switch(config-rbridge-id-3)# ag
```

4. Enter the port group ID, such as pg 1, to enter configuration mode for the port group where the N_Port currently resides.

```
switch(config-rbridge-id-3-ag)# pg 1
```

5. Delete the N_Port from the group using the **no nport interface Fibrechannel** *port* command. In the following example, N_Port 3 is removed from port group 1.

Before deleting the N_Port, all VF_Ports mapped to the N_Port should be remapped. If the port resides in port group 0 (pg 0), you do not need to remove it from the port group before adding it to a different port group and can skip to the next step.

```
switch(config-rbridge-id-3-ag-pg-1)# no nport interface Fibrechannel 3/0/3
```

You can delete multiple N_Ports by listing the ports separated by spaces as in the following example.

```
no nport interface Fibrechannel 3/0/3 3/0/5
```

- Enter the configuration mode for port group ID where you want to add the port, for example pg 2.

```
switch(config-rbridge-id-3-ag) # pg 2
```

- Add the N_Port to the group using the **nport interface Fibrechannel *port*** command where *port* is a supported N_Port number for the switch in rbridge-id/slot/port format.

```
switch(config-rbridge-id-3-ag-pg-2) # nport interface Fibrechannel 3/0/3
```

You can add multiple N_Ports by listing the ports separated by spaces. For example:

```
nport interface Fibrechannel 3/0/3 3/0/5
```

- Verify that the ports were added or removed from the port groups using the **show ag pg rbridge-id *rbridge-id*** while in the privileged EXEC command mode.

```
show ag pg rbridge-id 3
```

Refer to [Displaying port grouping information](#) on page 183 for details on this command.

Port Grouping policy modes

Port Grouping policy modes help manage VF_Port and N_Port operation when ports go offline or when all N_Ports in a group are not connected to the same FC fabric.

There are two Port Grouping policy modes that you can enable using Network OS commands:

- Login Balancing (LB) is enabled by default when you create a port group.
- Modified Managed Fabric Name Monitoring (M-MFNM) mode is enabled with LB mode. You cannot disable MFNM mode for a port group unless you disable LB mode.

Automatic Login Balancing mode

Automatic Login Balancing (LB) mode works to distribute logins across all available N_Ports in a port group. It is enabled by default when a port group is created.

Consider the following for LB mode:

- When LB mode is disabled for a port group, the same configured VF_Port to N_Port mapping displays for the **show running-config ag** or **show ag** commands. This is because configured and active mapping are the same.
- When LB mode is enabled for a port group, the **show ag** command displays the current, active mapping because VF_Port to N_Port mapping is based on the current distributed load across all N_Ports. The **show running-config ag** command displays the configured mapping only.
- If LB mode is enabled for a port group and a new N_Port comes online, existing logins are undisturbed. If an N_Port is disabled, its existing logins are distributed to available ports to maintain a balanced N_Port-to-VF_Port ratio.
- LB can be disabled using Network OS commands. When LB mode is disabled, VF_Ports are not shared among N_Ports in the port group, but can only connect to N_Ports to which they are mapped. If an N_Port is disabled, ENodes logged into mapped VF_Ports log out. As a best practice to ensure device login, bind the ENode to a VF_Port and ensure that its mapped N_Port is online.
- LB mode is disruptive.
- If an N_Port is removed from a port group enabled for LB mode and moved to another port group, the VF_Ports mapped to that N_Port remain with the N_Port.

This is not the case for port groups not enabled for LB mode. When you remove an N_Port from one of these port groups, the VF_Ports mapped to the N_Port move to the default Port Group 0 along with the N_Port. You can then move the N_Port to another group, but would need to re-map any VF_Ports to the N_Port.

- If an N_Port is in a port group and then Automatic Login Balancing is enabled, the VF_Ports mapped to the N_Port are distributed among online N_Ports in the same port group.
- You can disable or enable LB mode using the **no modes lb** or **modes lb** commands while in the port group configuration mode. Refer to [Enabling and disabling Login Balancing mode](#) on page 187.

Enabling and disabling Login Balancing mode

Although Login Balancing (LB) mode is enabled by default when you create a port group, you can disable and enable it using Network OS CLI commands.

Access Gateway and the PG policy must be enabled for the **no modes LB** or **modes lb** commands to succeed.

Enable or disable LB mode using the **no modes LB** or **modes lb** commands while in the port group's configuration mode.

1. Enter the **configure terminal** command to enter global configuration mode.

```
switch# configure terminal
```

2. Enter the **rbridge-id id** command to enter RBridge ID configuration mode for the specific switch.

```
switch(config)# rbridge-id 3
```

3. Enter the **ag** command to enter Access Gateway configuration mode.

```
switch(config-rbridge-id-3)# ag
```

4. Enter the port group ID, such as pg 8, to enter configuration mode for the port group.

```
switch(config-rbridge-id-3-ag)# pg 8
```

5. Perform one of the following steps:

- To enable LB mode, enter **modes mode_name**.

```
switch(config-rbridge-id-3-ag-pg-8)# modes lb
```

- To disable LB mode, enter **no modes mode_name**.

```
switch(config-rbridge-id-3-ag-pg-8)# no modes lb
```

Modified Managed Fabric Name Monitoring mode

Modified Managed Fabric Name Monitoring (M-MFNM) mode prevents connections from the AG VDX switch to multiple SANs to ensure that all N_Ports in a port group connect to the same FC fabric.

Modified Managed Fabric Name Monitoring (M-MFNM) mode is enabled with LB mode. It queries the FC fabric name for a default time out value of 120 seconds. If it detects an inconsistency, for example all the N_Ports within a port group are not physically connected to the same physical or virtual FC fabric, the following occurs:

- N_Ports are disabled to the fabric with the lower number of connected N_Ports.
- If more than one fabric has the same or the maximum number of ports connected, N_Ports are disabled to the fabric with the higher "fabric names" (WWN of the Principal Switch). Ports connected to the lowest "fabric name" stay online.

Consider the following about M-MFNM mode:

- M-MFNM mode is enabled by default when you enable LB mode. You cannot disable it unless you disable LB mode.
- You can change the default time out value (tov) for fabric name queries using the **timeout fnm value** when in ag configuration mode. Refer to [Setting and displaying the fabric name monitoring TOV](#) on page 188.

Setting and displaying the fabric name monitoring TOV

You can set the time out value (TOV) for M-MFNM queries of the fabric name to detect whether all N_Ports in a port group are physically connected to the same physical or virtual fabric.

Access Gateway and the PG policy must be enabled for the **timeout fnm** command to succeed.

Use the **timeout fnm value** command while in ag configuration mode to set time-out value (TOV) for M-MFNM queries of the fabric name. The valid range is 30 to 3600 seconds. The default value is 120 seconds.

1. Enter the **configure** command to access global configuration mode.

```
switch# configure
```

2. Enter the **rbridge-id id** command to enter RBridge ID configuration mode for the specific switch.

```
switch(config)# rbridge-id 3
```

3. Enter the **ag** command to enter Access Gateway configuration mode.

```
switch(config-rbridge-id-3)# ag
```

4. Enter the **timeout fnm value** command to change the time out value for fabric name queries.

```
switch(config-rbridge-id-3-ag)# timeout fnm 60
```

5. Enter timeout **timeout fnm** without a value to display the current M-MFNM timeout value.

```
switch(config-rbridge-id-3-ag)# timeout fnm
() (60)
```

Trunking in Access Gateway mode

The hardware-based Port Trunking feature enhances management, performance, and reliability of Access Gateway N_Ports connected to Brocade fabrics.

Trunking in Access Gateway (AG) mode creates trunk groups between N_Ports on the AG module and F_Ports on the Edge-switch module. AG trunking configuration is mostly on the Edge switch. Note the following for trunking under Access Gateway:

- On Access Gateway ports, trunking is enabled by default.
- You can create trunk groups of up to eight ports.
- All of the N_Ports in a trunk group should belong to the same AG-module port group.
- All of the F_Ports in a trunk group should belong to the same Edge-switch port group.
- The maximum number of trunks supported within a port group is four.
- On the AG module, ensure consistent speed settings and connector-group settings on all of the ports within each trunk.
- On the End-switch, ensure consistent speed settings and connector-group settings on all of the ports within each trunk.
- Round-robin assignment of VF_Ports to N_Ports affects the actual number of N_Ports in a trunk.

Setting up trunking for Access Gateway

Use the following steps to set up trunking under Access Gateway.

Make sure that all of the conditions specified under "Trunking in Access Gateway mode" are fulfilled. For Edge-switch implementation details, refer to the "Trunking in Access Gateway mode" chapter in the Fabric OS *Access Gateway Administrator's Guide*.

1. On the AG-mode switch, in privileged EXEC mode, enter the **configure** command to change to global configuration mode.
2. Enter the **interface FibreChannel** *rbridge-id/slot/port* command for each port that you are adding to the trunk.

A configuration submode prompt appears.

3. Enter the **trunk-enable** command.

The following example configures the link attached to port 4 on RBridge 8 to be part of a trunk group.

```
switch# configure
Entering configuration mode terminal
switch(config)# rbridge-id 8
switch(config-rbridge-id-8)# interface FibreChannel 8/0/4
switch(config-FibreChannel-8/0/4)# trunk-enable
```

4. Toggle the AG-mode port.

```
switch(config-FibreChannel-8/0/4)# shutdown
switch(config-FibreChannel-8/0/4)# no shutdown
```

5. Toggle the Edge-switch port.

Port trunking is now in effect between the specified ports.

NOTE

If you move trunk-group ports out of their common port group, the trunk slave ports will be disabled. If this happens, recover by entering **shutdown** and then **no shutdown** for each disabled port.

Access Gateway under FlexPort

FlexPort functionality enables specific ports to be dynamically reconfigured as either Fibre Channel (FC) or Ethernet ports, in several modes and speeds. This section deals with Access Gateway (AG) under FlexPort.

NOTE

For details of FlexPort implementation, refer to "Configuring FlexPort," in the *Network OS Layer 2 Switching Configuration Guide*.

NOTE

In Network OS 5.0.0, Forward Error Correction (FEC) is not supported in Access Gateway under FlexPort.

Configuring Access Gateway under FlexPort

To use Access Gateway under FlexPort, the best practice is to convert needed Ethernet ports to Fiber Channel ports before enabling Access Gateway.

For ports that you plan to configure as Fibre Channel (FC) ports under Access Gateway (AG) mode, note the following default settings:

- All N_Ports are in Port Group 0.
- All ports are configured as Ethernet ports, and none are configured as Fibre Channel (FC) ports.

Do not make the mistake of first configuring Fibre Channel over Ethernet (FCoE) and then using FlexPort to reconfigure Ethernet ports as FC ports. This mistaken order may lead to suboptimal login balance.

NOTE

If login balance among online N_Ports becomes uneven, refer to "Restoring_Port login balance." You can also turn off login balancing for a Port Group (as described in "Enabling and disabling Login Balancing mode") and manually map VF_Ports to N_Ports.

1. Disable Access Gateway.

```
switch# no ag enable
```

2. Using FlexPort commands, convert needed Ethernet ports to FC ports.

For details, refer to "Configuring FlexPort," in the *Network OS Layer 2 Switching Configuration Guide*.

3. Configure FCoE.

For details, refer to the "FCoE interface configuration" section of the *Network OS Layer 2 Switching Configuration Guide*.

4. Enable Access Gateway.

```
switch# ag enable
```

The VF_Ports are distributed evenly across all available N_Ports.

Restoring N_Port login balance

If login distribution among online N_Ports becomes uneven, use the **clear fcoe login** command to redistribute the logins.

NOTE

The following configurations may cause uneven login distribution:

- Access Gateway under FlexPort (following the initial setup)—Conversion of additional Ethernet ports to FC ports
- Toggle of N_Ports—If associated VF_Ports log in to another N_Port and remain there

1. To view the list of logged-in devices, enter the **show fcoe login** command.

```
switch# show fcoe login
```

2. To log out the current device and log in to the least-loaded N_Port, enter the **clear fcoe login device** command.

```
switch# clear fcoe login device 10:00:00:05:1e:8e:be:40
```

3. To log out of all devices in the Port Group, redistribute the VF_Ports to the available N_Ports, and automatically log back in to all the devices, enter the **clear fcoe login rbridge-id** command.

N_Port monitoring for unreliable links

N_Port monitoring monitors links between N_Ports on the switch configured in Access Gateway mode and F_Ports on the connected FC fabric. When links are considered unreliable, the N_Port is disabled.

Links from all N_Ports are monitored for the number of online and offline static change notifications (SCNs) that occur during a five-minute period. If the number of SCNs on a link exceeds a set threshold, the link is considered unreliable, and the port is taken offline. VF_Ports mapped to the N_Port also go offline. Once the number of SCNs drops below the set threshold, the port is deemed reliable again and the N_Port and the mapped VF_Ports go back online.

The default threshold is 25 SCNs per 5 minutes. You can set from 10 to 100 SCNs per 5 minutes. While in **ag** command mode, you can use the **counter reliability value** command to modify the default threshold.

Setting and displaying the reliability counter for N_Port monitoring

You can set the reliability count of static change notifications (SCNs) counted during a five-minute period before the link between a N_Port on a Switch in Access Gateway mode and an F_Port on a FC fabric is considered unreliable.

Access Gateway mode must be enabled for this procedure to succeed.

To set the reliability count, use the **counter reliability value** command while in ag configuration mode for the switch. The default value is 25 SCNs per 5 minutes. You can set from 10 to 100 SCNs per 5 minutes.

1. Enter the **configure terminal** command to enter global configuration mode.

```
switch# configure terminal
```

2. Enter the **rbridge-id id** command to enter RBridge ID mode for the specific switch.

```
switch(config)# rbridge-id 2
```

3. Enter the **ag** command to enter Access Gateway configuration mode.

```
switch(config-rbridge-id-2)# ag
```

4. Enter the **counter reliability value** command to change the counter value.

```
switch(config-rbridge-id-2-ag)# counter reliability 50
```

5. Enter the **counter reliability** command without a value to display the current reliability counter. In the following example, a counter value of 50 is returned.

```
switch(config-rbridge-id-2-ag)# counter reliability
() (50)
```

Displaying Access Gateway N_Port utilization data

Under Access Gateway, the **show ag nport-utilization** command displays N_Port utilization information. You can display this information either for a specific RBridge or for all the RBridges.

The information displayed indicates the highest bandwidth utilization and associated time stamp. Two actions clear such data:

- Enabling the port
- Running **clear ag nport-utilization**

NOTE

For trunk slave ports, no utilization information is printed. Instead, the bandwidth of such ports is included in the bandwidth of the trunk master port.

In Privileged EXEC mode, enter **show ag nport-utilization**.

```
switch# show ag nport-utilization
```

Data such as the following displays:

```
N_Port(Fi) information :
Port          PortID      Attached PWWN      IP_Addr      VF_Ports
-----
Fi 1/0/7      0xa90900   2f:00:00:05:1e:80:31:4f   10.17.31.169   1/1/1, 1/1/2
highest bandwidth utilization of 11 % recorded at Wed Apr 30 14:07:42 2014

Fi 1/0/8      0xa90900   2f:00:00:05:1e:80:31:4f   10.17.31.169   None
trunk slave. bandwidth/traffic added to trunk master
```


Using System Monitor and Threshold Monitor

- System Monitor overview.....193
- Configuring System Monitor.....194
- Threshold Monitor overview.....197
- Configuring Threshold Monitor.....200

System Monitor overview

System Monitor provides customizable monitoring thresholds, which allow you to monitor the health of each component of a switch. Whenever a switch component exceeds a configured threshold, System Monitor automatically provides notification by means of e-mail or RASLog messages, depending on the configuration.

Because of platform-specific values that vary from platform to platform, it was previously not possible to configure platform-specific thresholds through a global CLI command. In Network OS 4.0.0 and later, it is possible to monitor individual switches in a logical chassis cluster or fabric cluster. This is done in RBridge ID configuration mode, by addressing the RBridge ID of the selected switch.

Threshold and notification configuration procedures are described in the following sections.

Monitored components

The following FRUs and temperature sensors are monitored on supported switches:

- **LineCard** —Displays the threshold for the line card.
- **MM** —Displays the threshold for the management module.
- **SFM** —Displays the threshold for the switch fabric module device.
- **cid-card** —Displays the threshold for the chassis ID card component.
- **compact-flash** —Displays the threshold for the compact flash device.
- **fan** —Configures fan settings.
- **power** —Configures power supply settings.
- **sfp** —Displays the threshold for the small form-factor pluggable (SFP) device.
- **temp**—Displays the threshold for the temperature sensor component.

NOTE

CID cards can be faulted and removed. The system continues to operate normally as long as one CID card is installed. If both CID cards are missing or faulted, the switch will not operate.

Monitored FRUs

System Monitor monitors the absolute state of the following FRUs:

- Fan
- Power supply
- CID card

- SFP
- Line card

Possible states for all monitored FRUs are removed, inserted, on, off, and faulty. A state of none indicates the switch is not configured. If the FRU is removed, inserted, or goes into a faulty state, System Monitor sends a RASLog message or an e-mail alert, depending on the configuration.

Based on the configured threshold, each component can be in a marginal state or a down state. If a component is in a marginal state or a down state, System Monitor generates a RASLog message to alert the user. It also generates a separate RASLog message for the overall health of the switch.

NOTE

For details about each RASLog message, refer to the "RAS System Messages" chapter of the *Network OS Message Reference*.

The following table lists the marginal and down thresholds for components monitored by System Monitor on supported switches.

TABLE 29 Hardware platform marginal and threshold settings for supported switches

Platform	Hardware component	Marginal threshold	Down threshold
Brocade VDX 6740	Power supply	1	2
	Temperature sensor	1	2
	Compact flash	1	0
	Fan	1	2
Brocade VDX 8770-4	Power supply	1	2
	Temperature sensor	1	2
	Compact flash	1	0
	Fan	1	2
Brocade VDX 8770-8	Power supply	6	7
	Temperature sensor	1	2
	Compact flash	1	0
	Fan	1	2

Configuring System Monitor

This section contains example basic configurations that illustrate various functions of the **system-monitor** command and related commands.

NOTE

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

Setting system thresholds

Each component can be in one of two states, down or marginal, based on factory-defined or user-configured thresholds. (The default thresholds are listed in [Configuring System Monitor](#) on page 194.)

1. Issue the **configure terminal** command to enter global configuration mode.

2. Enter RBridge ID configuration mode, as in the following example.

```
switch(config)# rbridge-id 154
```

3. Change **down-threshold** and **marginal-threshold** values for the SFM.

```
switch(config-rbridge-id-154)# system-monitor sfm threshold down-threshold 3 marginal-threshold 2
```

NOTE

You can disable the monitoring of each component by setting **down-threshold** and **marginal-threshold** values to 0 (zero).

Setting state alerts and actions

System Monitor generates an alert when there is a change in the state from the default or defined threshold.

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge ID configuration mode (for RBridge ID 154 in this case).

```
switch(config)# rbridge-id 154
```

To enable a RASLog alert when the power supply is removed, enter the following command:

```
switch(config-rbridge-id-154)# system-monitor power alert state removed action raslog
```

NOTE

There are no alerts for MM, compact-flash, or temp. There are no alert actions for SFPs.

Configuring e-mail alerts

Use the **system-monitor-mail fru** command to configure e-mail threshold alerts for FRU, SFP, interface, and security monitoring. For an e-mail alert to function correctly, you must add the IP addresses and host names to the domain name server (DNS) in addition to configuring the domain name and name servers. A single email configuration is applicable for all switches in a logical chassis cluster. For complete information on the **system-monitor-mail relay host** command, refer to the *Network OS Command Reference*.

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter the following command to enable e-mail alerts and to configure the e-mail address.

```
switch(config)# system-monitor-mail fru enable email-id
```

Sendmail agent configuration

The following **system-monitor-mail relay host** commands allow the sendmail agent on the switch to resolve the domain name and forward all e-mail messages to a relay server.

- To create a mapping:

```
switch(config)# system-monitor-mail relay ip-address 1.2.3.4 domain-name domain_name1.brocade.com
```

- To delete the mapping:

```
switch(config)# no system-monitor-mail relay ip-address 1.2.3.4 domain-name domain_name1.brocade.com
```

- To change the domain name:

```
switch(config)# system-monitor-mail relay ip-address 1.2.3.4 domain-name domain_name2.brocade.com
```

NOTE

You must delete the first domain name before you can change it to a new domain name.

- To delete the domain name and return to the default:

```
switch(config)# no system-monitor-mail relay ip-address 1.2.3.4 domain-name domain_name2.brocade.com
```

Viewing system SFP optical monitoring defaults

You can view the optical monitoring default values by entering **show defaults threshold** followed by the SFP type.

The following example command will display the defaults for type 1GLR SFPs:

```
switch# show defaults threshold sfp type 1GLR
```

Displaying the switch health status

To display the health status of a switch, enter **show system monitor**.

```
switch# show system monitor
** System Monitor Switch Health Report **
RBridge 154      switch status      : MARGINAL
                 Time of Report      : 2013-03-24 20:51:53
                 Power supplies monitor : MARGINAL
                 Temperatures monitor  : HEALTHY
                 Fans monitor          : HEALTHY
                 Flash monitor         : HEALTHY
```

Threshold Monitor overview

The **threshold-monitor** commands allow you to monitor CPU and memory usage of the system, interface and SFP environmental status, and security status and be alerted when configured thresholds are exceeded. These commands are configured in RBridge ID configuration mode to support fabric cluster and logical chassis cluster topologies.

In addition to the policy keywords (available for **interface**, **SFP**, and **security monitoring**), you can use the **custom** keyword create your own custom policies that have non-default thresholds, and apply them by means of the **apply** operand. This allows you to toggle between default settings and saved custom configuration settings and to apply actions and thresholds separately. For example, you can choose to use default threshold settings together with a customized subset of available actions, or you can modify some of the threshold settings and use the default action settings. You can also pause monitoring and actions by means of the **pause** keyword.

For detailed information on the variables and keywords (operands) of the **threshold-monitor** series of commands, refer to the *Network OS Command Reference*.

CPU and memory monitoring

When configuring CPU monitoring, specify a value in the 1-100 range. When the CPU usage exceeds the limit, a threshold monitor alert is triggered. The default CPU limit is 75 percent. With respect to memory, the limit specifies a usage limit as a percentage of available resources.

When used to configure memory or CPU threshold monitoring, the limit value must be greater than the low limit and smaller than the high limit. The alert provided is a RASLog message, with the following options configurable under the **raslog** option of the **threshold-monitor cpu** or the **threshold-monitor memory** commands:

high-limit	Specifies an upper limit for memory usage as a percentage of available memory. This value must be greater than the value set by limit . When memory usage exceeds this limit, a RASLog CRITICAL message is sent. Valid values range from 0 through 80 percent.
limit	Specifies the baseline memory usage limit as a percentage of available resources. When this value is exceeded, a RASLog WARNING message is sent. When the usage returns below the value set by limit , a RASLog INFO message is sent. Valid values range from 0 through 80 percent.
low-limit	Specifies a lower limit for memory usage as percentage of available memory. This value must be smaller than the value set by limit . When memory usage exceeds or falls below this limit, a RASLog INFO message is sent.
poll	Specifies the polling interval in seconds. Valid values range from 0 through 3600.
retry	Specifies the number of polling retries before desired action is taken. Valid values range from 1 through 100.

NOTE

For CPU and memory thresholds, the low limit must be the lowest value and the high limit must be the highest value.

The table below lists the factory defaults for CPU and memory thresholds.

TABLE 31 Default values for CPU and memory threshold monitoring

Operand	Memory	CPU
low-limit	40%	N/A
limit	60%	75%
high-limit	70%	N/A
poll	120 seconds	120 seconds
retry	3	3

SFP monitoring

The SFP parameters that can be monitored are listed and described below.

TABLE 32 SFP parameter descriptions

SFP parameter	Description	Suggested SFP impact
Temperature	Measures the temperature of the SFP, in degrees Celsius.	High temperature suggests the SFP might be damaged.
Receive power (RXP)	Measures the amount of incoming laser, in μ Watts.	Describes the condition of the SFP. If this parameter exceeds the threshold, the SFP is deteriorating.
Transmit power (TXP)	Measures the amount of outgoing laser power, in μ Watts.	Describes the condition of the SFP. If this parameter exceeds the threshold, the SFP is deteriorating.
Current	Measures the amount of current supplied to the SFP transceiver.	Indicates hardware failures.
Voltage	Measures the amount of voltage supplied to the SFP.	A value higher than the threshold indicates the SFP is deteriorating.

SFP thresholds

You can customize SFP thresholds or actions by using the **threshold-monitor sfp** command, which enables you to perform the following tasks.

- Customize SFP configurations or accept SFP defaults.
- Manage the actions and thresholds for the Current, Voltage, RXP, TXP, and Temperature areas of the SFP.
- Suspend SFP monitoring.

If you do not provide the SFP type parameters, the default thresholds and actions are used. SFP types, monitoring areas, and default threshold values for the 16-Gbps and QSFP SFPs are detailed below.

TABLE 33 Factory thresholds for SFP types and monitoring areas

SfpType	Area	Default Value	
1 GSR	Temperature (C)	100	-40
	Voltage (mV)	3600	3000
	RXP (μ W)	1122	8
	TXP (μ W)	1000	60
	Current (mA)	12	2
1 GLR	Temperature (C)	90	-45
	Voltage (mV)	3700	2900
	RXP (μ W)	501	6
	TXP (μ W)	794	71
	Current (m)	45	1
10 GSR	Temperature (C)	90	-5
	Voltage (mVolt)	3600	3000
	RXP (μ W)	1000	32
	TXP (μ W)	794	251
	Current (mA)	11	4
10 GLR	Temperature (C)	88	-5

TABLE 33 Factory thresholds for SFP types and monitoring areas (continued)

SfpType	Area	Default Value	
	Voltage (mV)	3600	2970
	RXP (μ W)	1995	16
	TXP (μ W)	1585	158
	Current (mA)	85	15
10 GUSR	Temperature (C)	100	-5
	Voltage (mV)	3600	2970
	RXP (μ W)	2000	32
	TXP (μ W)	2000	126
	Current (mA)	11	3
QSFP	Temperature (C)	75	-5
	Voltage (mV)	3600	2970
	RXP (μ W)	1995	40
	TXP (μ W)	0	0
	Current (mA)	10	1

Threshold values

High and low threshold values are the values at which potential problems might occur. For example, in configuring a temperature threshold for SFPs, you can select the temperatures at which a potential problem can occur because of overheating or overcooling.

A combination of high and low threshold settings can cause the following actions to occur:

- Above high threshold — A default or user-configurable action is taken when the current value is above the high threshold.
- Below high threshold — A default or user-configurable action is taken when the current value is between the high and low threshold.
- Below low threshold — A default or user-configurable action is taken when the current value is below the low threshold.
- Above low threshold — monitoring is not supported for this value.

Security monitoring

You can monitor all attempts to breach your SAN security, helping you fine-tune your security measures. If there is a security breach, you can configure an email or RASLog alert to be sent. The following security areas are monitored:

- Telnet Violation, which occurs when a Telnet connection request reaches a secure switch from an unauthorized IP address.
- Login Violation, which occurs when a secure fabric detects a login failure.

The following table lists the factory defaults for security area settings.

TABLE 34 Security area default settings

Area	High threshold	Low threshold	Buffer	Timebase
Telnet Violation	2	1	0	Minute
Login Violation	2	1	0	Minute

Interface monitoring

You can set thresholds for error statistics on all external Gigabit Ethernet interfaces. When any monitored error crosses the configured high or low threshold, an alert can be generated or a problem interface can be isolated (refer to [Port Fencing](#) on page 200).

Interface error types

The following table describes the interface counters that can be monitored on external interfaces.

TABLE 35 Interface errors that can be monitored on external interfaces

Interface area	Description	Port Fencing support	Threshold defaults
MissingTerminationCharacter	Number of frames terminated by anything other than the Terminate character; this includes termination due to the Error character.	No	Low 12 Buffer 0 High 300
CRCAAlignErrors	Total number of frames received that had a length (excluding framing bits but including Frame Check Sequence (FCS) octets) of between 64 and 1518 octets. The error indicates either a bad FCS with an integral number of octets (an FCS error) or a bad FCS with a non-integral number of octets (an alignment error).	No	Low 12 Buffer 0 High 300
IFG	Minimum-length interframe gap (IFG) between successive frames is violated. A typical IFG is 12 bytes.	Yes	Low 5 Buffer 0 High 100
SymbolErrors	An undefined (invalid) symbol received on the interface. Large symbol errors indicate a bad device, cable, or hardware.	No	Low 0 Buffer 0 High 5

NOTE

The default setting for above high threshold, above low threshold, below high threshold, and below low threshold actions is "[none]."

Port Fencing

A port that is consistently unstable can harm the responsiveness and stability of the entire fabric and diminish the ability of the management platform to control and monitor the switches within the fabric. *Port Fencing* is not enabled by default; it disables the interface if a user-defined high threshold is exceeded. When a port that has exceeded its user-defined high threshold is fenced by software, the port is placed in the "Disabled" state and held offline. After a port is disabled, user intervention is required for frame traffic to resume on the port.

NOTE

Port Fencing is supported for the "RX IFG Violated" error only.

Configuring Threshold Monitor

The following basic configurations illustrate various functions of the **threshold-monitor** commands.

NOTE

For CLI details, refer to the *Network OS Command Reference*

Viewing threshold status

To view the status of currently configured thresholds, enter the **show running-config threshold-monitor** command with the RBridge ID, as follows:

```
switch# show running-config rbridge-id rbridge_id threshold-monitor
```

NOTE

Default values are not displayed under the **show running-config threshold-monitor** command. Only custom values are displayed when a user applies a policy.

To display the default values of thresholds and alert options, enter the **show defaults threshold** command, as in the following example for interfaces.

```
switch# show defaults threshold interface type Ethernet
```

```
Type: GigE-Port
```

Area	High Threshold			Low Threshold			Buffer Value	Time Base
	Value	Above Action	Below Action	Value	Above Action	Below Action		
MTC	300	none	none	12	none	none	0	minute
CRCAAlign	300	none	none	12	none	none	0	minute
Symbol	5	none	none	0	none	none	0	minute
IFG	100	none	none	5	none	none	0	minute

```
MTC - Missing Termination Character
```

CPU and memory threshold monitoring

NOTE

Support for the custom **policy** operand is not provided for CPU and memory threshold monitoring.

Configuring CPU monitoring thresholds and alerts

CPU monitoring allows you to set alerts for CPU usage.

1. Enter **configure terminal** to enter global configuration mode.

```
switch# configure terminal
switch(config)#
```

2. Enter **rbridge-id rbridge-id#** to change to RBridge ID configuration mode.

```
switch(config)# rbridge-id 154
switch(config-rbridge-id-154)#
```

3. Enter **threshold-monitor cpu ?** to view the available options:

```
switch(config-rbridge-id-154)# threshold-monitor cpu ?
```

The following example changes the thresholds from the default, adjusts polling and retry attempts, and causes a RASLog message to be sent when thresholds are exceeded.

```
switch(config-rbridge-id-154)# threshold-monitor cpu actions raslog limit 65 poll 60 retry 10
```

NOTE

This command does not support **low-limit** or **high-limit** under the **raslog** alert option.

Configuring memory monitoring thresholds and alerts

CPU monitoring allows you to set alerts for memory usage.

1. Enter **configure terminal** to enter global configuration mode.

```
switch# configure terminal
switch(config)#
```

2. Enter **rbridge-id rbridge-id#** to change to RBridge ID configuration mode.

```
switch(config)# rbridge-id 154
switch(config-rbridge-id-154)#
```

3. Enter **threshold-monitor memory ?** to view the available options.

```
switch(config-rbridge-id-154)# threshold-monitor memory ?
```

The following example changes the thresholds from the default and causes no message to be sent when thresholds are exceeded.

```
switch(config-rbridge-id-1)# threshold-monitor memory actions none high-limit 60 low-limit 40
```

Configuring SFP monitoring thresholds and alerts

The following is an example of configuring SFP monitoring.

1. Enter **configure terminal** to enter global configuration mode.

```
switch# configure terminal
switch(config)#
```

2. Enter **rbridge-id rbridge-id#** to change to RBridge ID configuration mode.

```
switch(config)# rbridge-id 154
switch(config-rbridge-id-154)#
```

3. Enter **threshold-monitor sfp** and create a custom policy.

```
switch(config-rbridge-id-154)# threshold-monitor sfp policy mypolicy type lglr area temperature
alert above highthresh-action raslog email
```

NOTE

Refer also to [Security monitoring](#) on page 199 for more information.

4. Apply the policy.

```
switch(config-rbridge-id-154)# threshold-monitor sfp apply mypolicy
```

Security monitoring

Security monitoring allows you to set security threshold and alert options, including login-violation or telnet-violation alerts.

Viewing security defaults

To display the default values of security threshold and alert options, enter the **show defaults security area** command with the **login-violation** or **telnet-violation** options.

```
switch# show defaults security area login-violation
```

Configuring security monitoring

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge ID configuration mode, as in the following example.

```
switch(config)# rbridge-id 154
```

3. Enter the **threshold-monitor security** command to configure custom login-violation monitoring, as in the following example.

```
switch(config-rbridge-id-154)# threshold-monitor security
policy mypolicy area login-violation alert above highthresh-action
raslog below highthresh-action email lowthresh-action none
```

4. Apply the policy.

```
switch(config-rbridge-id-154)# threshold-monitor security apply mypolicy
```

Configuring Interface monitoring

The following sections discuss how to view interface threshold defaults and configure interface monitoring.

Viewing interface threshold defaults

Use the following command to view interface threshold defaults.

```
switch# show defaults threshold interface type Ethernet
```

Refer to [Viewing threshold status](#) on page 201 for the results of this command.

Configuring interface monitoring

1. Issue the **configure terminal** command to enter global configuration mode.
2. Enter RBridge ID configuration mode (in this case, for RBridge ID 154).

```
switch(config)# rbridge-id 154
```

3. Enter the **threshold-monitor interface** command to configure custom interface monitoring, as in the following example.

```
switch(config-rbridge-id-154)# threshold-monitor interface policy mypolicy type ethernet area
missingterminationcharacter alert above lowthresh-action email
```

4. Apply the policy.

```
switch(config-rbridge-id-154)# threshold-monitor interface apply mypolicy
```

Pausing and continuing threshold monitoring

By default, threshold monitoring is enabled.

To disable monitoring of a particular type, enter the **threshold-monitor [cpu | interface | memory | security | sfp] pause** command.

To re-enable monitoring, enter the **no** of the **threshold-monitor** command.

NOTE

Not all functions of the **threshold-monitor** command can be disabled. Continue to enter **?** at each level of the command synopsis to confirm which functions can be disabled.

Using VMware vCenter

- [vCenter and Network OS integration overview.....](#) 205
- [vCenter discovery.....](#) 206
- [vCenter configuration.....](#) 206

vCenter and Network OS integration overview

The VMware vCenter Server allows for the management of multiple ESX /ESXi servers and virtual machines (VMs) from different ESX servers through a single graphical user interface (GUI). It provides unified management of all the hosts and VMs in the data center, from a single console with an aggregate performance monitoring of clusters, hosts and VMs.

The VMware vCenter and Brocade Network OS integration supported in Brocade VCS Fabric mode enables you to discover VMware ESX servers managed by a vCenter server. VMware's server hosts (ESX servers) are connected directly to the physical switches through the switch ports (edge ports in Brocade VCS Fabric mode). The server hosts implement a virtual switch (vSwitch), which is used to provide connections to the VMs. The fundamental requirement for the vCenter and Network OS integration is the IP-level management connectivity of the vCenter Server 4.0 version and later with the Brocade VDX switches.

NOTE

The Network OS integration with vCenter requires vCenter versions 4.0, 4.1, 5.1 or 5.5.

You can view virtual switches and virtual machines, their associated MAC addresses, and network policies using the Network OS command line interface (CLI). Refer to the *Network OS Command Reference* for details about the **vcenter** and **vnetwork** commands.

vCenter properties

The vCenter manages the VMware ESX/ESXi hosts. The vCenter user interface is provided through a vSphere client on the same management network as the vCenter, and virtual machines (VMs) are created using the vSphere client user interface. In addition to creating the VMs, the server administrator associates the VMs with distributed virtual switches, distributed virtual port groups, standard virtual switches (vSwitches) and standard port groups.

The vCenter automatically generates some of the VM properties (such as the MAC address), and some properties must be configured (such as the VLAN properties). Most of the VM configuration, including network policies, is done using the vCenter's vSphere user interface and is beyond the scope of this document.

For VMWare configuration information, visit the VMware documentation site.

vCenter guidelines and restrictions

Follow these guidelines and restrictions when configuring vCenter:

- Special characters in the port group names are replaced with the URL-encoded values.
- Standard port groups with the same name that reside in different ESX/ESXi hosts must have identical VLAN settings across all hosts.
- For all vCenter port groups, Network OS automatically creates a port profile with the following format: `auto-vcenter_name-datacenter_ID-port-group-name`. User editing of these auto port groups is not supported.
- Network OS supports vCenter discovery that is based on events.
- Network OS supports LLDP and QoS (IEEE 8021.p) for distributed virtual switches (dvSwitches).

- Network OS supports up to 750 port groups in the vCenter.
- Using port-profile names fewer than 63 characters has been shown to conserve CPU resources.
- CDP/LLDP-receiving interface ports must not have any conflicting configurations (such as switch port and FCoE port configurations) on the interface that prevent them from being in a port-profiled mode.
- Before configuring a vCenter in the fabric, remove all the manually created port profiles that have vCenter inventory MAC associations.
- Up to four multiple data centers are supported.
- Duplicate vCenter asset values are not supported, such as duplicate MAC addresses and duplicate Host names.

vCenter discovery

A Brocade VDX switch connected to VMware ESX/ESXi hosts and virtual machines must be aware of network policies in order to allow or disallow traffic; this requires a discovery process by the VDX switch. During VDX switch configuration, relevant vCenters that exist in its environment and the discovery of virtual assets from the vCenter occurs in the following circumstances:

- When a switch boots up
- When a new vCenter is configured on the VDX switch and activated (activation turns on the timer processing, set to 30-minute intervals)
- When the discovery is explicitly initiated with the CLI

The following assets are discovered from the vCenter:

- Hosts and data centers associated with the vCenter
- Virtual machines (VMs) that have been created on the hosts
- VMware distributed virtual port groups (dvPortGroups)
- Standard port groups, with QoS priority associated with a dvPortGroup
- Standard virtual switches
- Distributed virtual switches

vCenter configuration

Configuring vCenter consists of three basic steps performed in this order:

1. Enabling VMware vSphere QoS.
2. Enabling CDP/LLDP on switches.
3. Adding and activating the vCenter.

These steps and postconfiguration steps are discussed in this section.

Step 1: Enabling QoS

You must edit the network resource pool settings and set QoS priorities. Refer to the latest VMware vSphere Networking documentation.

Step 2: Enabling CDP/LLDP

In order for an Ethernet Fabric to detect the ESX/ESXi hosts, you must first enable Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) on all the virtual switches (vSwitches) and distributed vSwitches (dvSwitches) in the vCenter Inventory.

For more information, refer to the VMware KB article 1003885.

Enabling CDP/LLDP on vSwitches

Complete the following steps to enable CDP/LLDP on virtual switches (vSwitches).

1. Login as root to the ESX/ESXi Host.
2. Use the following command to verify the current CDP/LLDP settings.

```
[root@server root]# esxcfg-vswitch -b vSwitch1
```

3. Use the following command to enable CDP/LLDP for a given virtual switch. Possible values here are **advertise** or **both**.

```
[root@server root]# esxcfg-vswitch -B both vSwitch1
```

Enabling CDP/LLDP on dvSwitches

Complete the following steps to enable CDP on distributed virtual switches (dvSwitches).

1. Connect to the vCenter server by using the vSphere Client.
2. On the vCenter Server home page, click **Networking**.
3. Right-click the distributed virtual switches (dvSwitches) and click **Edit Settings**.
4. Select **Advanced** under **Properties**.
5. Use the check box and the drop-down list to change the CDP/LLDP settings.

Step 3: Adding and activating the vCenter

After CDP is enabled on all the vSwitches and dvSwitches in the vCenter, configuration on the Network OS side is a two-step process, consisting of adding the vCenter and activating the vCenter.

Adding the vCenter

You must add the vCenter before initiating any discovery transactions. To authenticate with a specific vCenter, you must first configure the URL, login, and password properties on the VDX switch.

NOTE

By default, the vCenter server accepts only HTTPS connection requests.

1. Enter the **vcenter** command with the name, URL, user name, and password of the vCenter.

```
switch(config)# vcenter myvcenter url https://10.2.2.2 username user password pass
```

2. An invalid state or condition of a vCenter can cause the deletion of all auto-port-profiles in a system. To prevent this from happening, configure the **ignore-delete-all-response** operand of the **vcenter** command to ignore the "delete-all" responses from the vCenter.

```
switch# vcenter MYVC discover ignore-delete-all-response 5
```

Activating the vCenter

After adding the vCenter, you must activate the configured vCenter instance.

NOTE

In VCS mode, you can configure the vCenter by using any node. Discovery is initiated by the primary node.

1. Enter the **configure terminal** command.
2. Enter the **vcenter** command to activate the vCenter.

```
switch(config)# vcenter myvcenter activate
```

Immediately following first-time vCenter activation, the Network OS starts the virtual asset discovery process. Use the **show vnetwork vcenter status** command to display the vnetwork status, as in the following example.

```
switch# show vnetwork vcenter status
vCenter          Start                Elapsed (sec)      Status
=====          =====
myvcenter        2011-09-07 14:08:42  10                 In progress
```

When the discovery process completes, the status displays as "Success." Network OS has performed all the necessary configurations needed for the vCenter Server, and is now ready for CDP transmissions from the virtual switches to identify which ESX/ESXi host is connected to which physical interface in the Ethernet Fabric.

Discovery timer interval

By default, Network OS queries the vCenter updates every thirty minutes. If any virtual assets are modified (for example, adding or deleting virtual machines (VMs), or changing VLANs), Network OS detects those changes and automatically reconfigures the Ethernet Fabric during the next periodic rediscovery attempt.

Use the **vcenter interval** command to manually change the default timer interval value to suit the individual environment needs.

```
switch(config)# vcenter myvcenter interval ?
Possible completions:
<NUMBER:0-1440> Timer Interval in Minutes (default = 30)
```

NOTE

Best practice is to keep the discovery timer interval value at the default (30). A value of 0 disables the periodic vCenter discovery.

User-triggered vCenter discovery

The discovery of virtual assets from the vCenter occurs during one of the following circumstances:

- When a switch boots up.
- When a new vCenter is configured on the VDX switch and activated (activation turns on the timer processing, set to 180-second intervals.)
- When the discovery is explicitly initiated with the CLI.

To explicitly initiate vCenter discovery, perform the following task in global configuration mode.

1. An invalid state or condition of a vCenter can cause the deletion of all auto-port-profiles in a system. To prevent this from happening, configure the **ignore-delete-all-response** operand of the **vcenter** command to ignore the "delete-all" responses from the vCenter.

```
switch(config)# vcenter MYVC discover ignore-delete-all-response 5
```

2. Return to privileged EXEC mode with the **exit** command.

```
switch(config)# exit
switch#
```

3. Use the **vnetwork vcenter** command to trigger a vCenter discovery manually.

```
switch# vnetwork vcenter myvcenter discover
```

Viewing the discovered virtual assets

Enter one of the following **show vnetwork** asset commands:

```
switch# show vnetwork dvpgs datacenter
  datacenter_name
  vcenter
  vcenter_name
switch# show vnetwork dvs datacenter
  datacenter_name
  vcenter
  vcenter_name
switch# show vnetwork hosts datacenter
  datacenter_name
  vcenter
  vcenter_name
switch# show vnetwork pgs datacenter
  datacenter_name
  vcenter
  vcenter_name
switch# show vnetwork vcenter status datacenter
  datacenter_name
  vcenter
  vcenter_name
switch# show vnetwork vmpolicy
  datacenter
  datacenter_name
  vcenter
  vcenter_name
switch# show vnetwork vms datacenter
  datacenter_name
  vcenter
  vcenter_name
switch# show vnetwork vss datacenter
  datacenter_name
  vcenter
  vcenter_name
```

where:

- **dvpgs** — Displays discovered distributed virtual port groups.
- **dvs** — Displays discovered distributed virtual switches.
- **hosts** — Displays discovered hosts.
- **pgs** — Displays discovered standard port groups.
- **vcenter status** — Displays configured vCenter status.

- **vmpolicy** — Displays the following network policies on the Brocade VDX switch: associated media access control (MAC) address, virtual machine, (dv) port group, and the associated port profile.
- **vms** — Displays discovered virtual machines (VMs).
- **vss** — Displays discovered standard virtual switches.

Refer to the *Network OS Command Reference* for detailed information about the **show vnetwork** commands.

Configuring Remote Monitoring

- [RMON overview.....211](#)
- [Configuring and managing RMON.....211](#)

RMON overview

Remote monitoring (RMON) is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. The RMON specification defines a set of statistics and functions that can be exchanged between RMON-compliant console managers and network probes. As such, RMON provides you with comprehensive network-fault diagnosis, planning, and performance-tuning information.

Configuring and managing RMON

Both alarms and events are configurable RMON parameters.

- Alarms allow you to monitor a specific management information base (MIB) object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold). Alarms are paired with events; the alarm triggers an event, which can generate a log entry or an SNMP trap.
- Events determine the action to take when an event is triggered by an alarm. The action can be to generate a log entry, an SNMP trap, or both. You must define the events before an alarm can be configured. If you do not configure the RMON event first, you will receive an error when you configure the alarm settings.

By default, no RMON alarms and events are configured and RMON collection statistics are not enabled.

Configuring RMON events

You can add or remove an event in the RMON event table that is associated with an RMON alarm number.

To configure RMON events, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
switch# configure terminal
```

2. Configure the RMON event.

```
switch(config)# rmon event 27 description Rising_Threshold log owner john_smith trap syslog
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Save the *running-config* file to the *startup-config* file.

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

Configuring RMON Ethernet group statistics collection

You can collect RMON Ethernet group statistics on an interface. RMON alarms and events must be configured for you to display collection statistics. By default, RMON Ethernet group statistics are not enabled.

Ethernet group statistics collection is not supported on ISL links.

To collect RMON Ethernet group statistics on an interface, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
switch# configure terminal
```

2. Enter the **interface** command to specify the interface type and RBridge-id/slot/port number.

```
switch(config)# interface tengigabitethernet 1/0/1
```

3. Enable the DCB interface.

```
switch(conf-if-te-1/0/1)# no shutdown
```

4. Configure RMON Ethernet group statistics on the interface.

```
switch(conf-if-te-1/0/1)# rmon collection stats 200 owner john_smith
```

5. Return to privileged EXEC mode.

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

6. Enter the **copy** command to save the running-config file to the startup-config file.

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

Configuring RMON alarm settings

To configure RMON alarms and events, perform the following steps from privileged EXEC mode.

1. Enter the **configure terminal** command to access global configuration mode.

```
switch# configure terminal
```

2. Configure the RMON alarms.

Example of an alarm that tests every sample for a rising threshold

```
switch(config)# rmon alarm 5 1.3.6.1.2.1.16.1.1.1.5.65535 interval 30
                    absolute rising-threshold 95 event 27 owner john_smith
```

Example of an alarm that tests the delta between samples for a falling threshold

```
switch(config)# rmon alarm 5 1.3.6.1.2.1.16.1.1.1.5.65535 interval 10 delta
                    falling-threshold 65 event 42 owner john_smith
```

3. Return to privileged EXEC mode.

```
switch(config)# end
```

4. Save the *running-config* file to the *startup-config* file.

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

5. To view configured alarms, use the **show running-config rmon alarm** command.

Monitoring CRC errors

Certain interface counters, such as those for CRC errors, may not be available by means of SNMP OIDs. In this case it is recommended that either RMON or CLI be used to monitor those statistics.

The following synchronizes the statistics maintained for the interface and RMON, as well as ensures proper reporting from an operational standpoint.

1. First use the **clear counters all** command in global configuration mode.

```
device# clear counters all
```

2. Then use the **clear counters rmon** command.

```
device# clear counters rmon
```

3. Finally, execute the **rmon collection stats** command on each interface, as in the following example.

```
device(config)# interface tengigabitethernet 170/0/1  
device(conf-if-te-170/0/1)# rmon collection stats 2 owner admin
```

4. Use an appropriate RMON MIB for additional monitoring.

For example, to obtain CRC statistics on a Brocade VDX platform, the following RMON MIB could be used: Object-etherStatsCRCAAlignErrors, OID- .1.3.6.1.2.1.16.1.1.1.8