HITACHI
Inspire the Next

Readme for Network OS v.4.0.1

Appendix for Administrator Guide

FastFind Links

Hitachi Data Systems

MK-99CB2500044-01

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

Contents

Preface

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

This preface includes the following information:

- Intended Audience
- Product Version
- Release Notes
- Document Organization
- Document Conventions
- Getting Help
- Comments

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

Intended Audience

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

This document assumes the following:

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

Product Version

This document revision applies to Network OS for Brocade 10Gbps DCB switch module version v.4.0.1.

Release Notes

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

Document Organization

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

Chapter	Description
Chapter 1. Network OS 4.0.1 overview	Describes the overview of Network OS 4.0.1.
Chapter 2. Prerequisites of Network OS 4.0.1	Describes the prerequisites of Network OS 4.0.1.
Chapter 3. Feature Description	Provides the information of New/Enhanced feature for Network OS 4.0.1.

Document Conventions

This document uses the following typographic conventions:

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

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

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

Getting Help

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

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

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

Comments

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

1

Network OS 4.0.1 overview

In this chapter, the overview of Network OS 4.0.1 is described.

Brocade Network Operating System (NOS) v.4.0.1 is a first release to support the embedded DCB switch module of CB2500. The latest version for the switch module of CB2500 is NOS v.4.0.1_hit1.

Feature Outline

Network OS v4.0.1 includes the following features for the embedded DCB switch module against Brocade Network OS release.

- DiST (STPoVCS)
- Uni-directional Link Detection (UDLD)
- Flow-based QOS
- Flow-based sflow
- RSPAN
- Border Gateway Protocol (BGP)
- VRF-Lite
- QSFP breakout cable support for 40G port
- Inbuilt packet capture utility (PCAP)
- OSPF enhancements
- VRRP/VRRPe enhancements
- Multiple sflow collectors and IPv6 based sflow
- SPAN support on ISL port
- Management Services enhancements
- Port Security
- ICMP Rate Limiting

2

Prerequisites of Network OS 4.0.1

In this chapter, the prerequisites for Network OS 4.0.1 are described.

This chapter provides fundamental information about Network OS 4.0.1 as exemplified by the optional license, the standard compliance, scalability and compatibility.

Additionally the important notes described about the features of Hitachi Network OS 4.0.1 release.

OPTIONAL LICENSED SOFTWARE

Network OS v4.0.x supports the following licensed features:

- **Port upgrade license 18p** — Allows customers to instantly scale the fabric by provisioning additional ports via license key upgrade. (Applies to select models of switches).

- **40Gb activate license** — Allows customers to instantly scale the fabric by provisioning additional 40G ports via license key upgrade.

Software licenses are available in following formats.

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

Standards Compliance

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

The embedded DCB switch module conforms to the following Ethernet standards:

- IEEE 802.1D Spanning Tree Protocol

- IEEE 802.1s Multiple Spanning Tree

- IEEE 802.1w Rapid reconfiguration of Spanning Tree Protocol

- IEEE 802.3ad Link Aggregation with LACP

- IEEE 802.3ae 10G Ethernet

- IEEE 802.1Q VLAN Tagging

- IEEE 802.1p Class of Service Prioritization and Tagging

- IEEE 802.1v VLAN Classification by Protocol and Port

- IEEE 802.1AB Link Layer Discovery Protocol (LLDP)

- IEEE 802.3x Flow Control (Pause Frames)

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

- IEEE 802.1Qbb Priority-based Flow Control
- IEEE 802.1Qaz Enhanced Transmission Selection
- IEEE 802.1 DCB Capability Exchange Protocol (Proposed under the DCB Task Group of IEEE 802.1 Working Group)

The embedded DCB switch module conforms to the following Internet IETF RFCs:

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

Scalability

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

Network OS v4.0.1Scalability Numbers	VCS Fabric
Maximum # of VLANs	3,500
Maximum # of MAC addresses	120,000
Maximum # of port profiles(AMPP)	750
Maximum # of VLANS in port profiles	3,500
Maximum # of MAC Associations for AMPP	16,000
Maximum # of per priority pause levels	3
Maximum # of IGMP Snooping Interfaces supported	256
Learning rate for IGMP snooping (groups/second)	512
Maximum # of L2 (IGMP Snooping) multicast groups	6,000
# of L3 (S,G) forwarding Entries	2,000
# of L3 (*,G) forwarding Entries	256
# of L2/L3 Multicast Flows	10,000
PIM Interfaces Supported	32
IGMP interfaces supported	32
Learning Rate for PIM-SM (flows/second)	32
Maximum # VLAN per Edge Port in Trunk ModePort***	4,000
Maximum # of VF Ports (Per switch)	1,000
Maximum # of Enodes per Fabric	2,000
Number of active STP nodes (Per cluster)	8
No. of ports per node(phy) participating in xSTP	60
No. of ports per Cluster in xSTP	480
No.of vlags assuming 2-4 nodes in vlag	64(512* if bpdu guard enabled)
No. of vlags (participating xSTP)assuming 2-4 nodes in vlag	64
No. of vlags (participating xSTP)assuming 8 nodes in vlag	32
No. of PVST instances	128
No. of RPVST instances	128
No. of MSTP instances	32
Maximum # of VLAN in PVST	128
Maximum # of LAG groups	60
Maximum # of members in a standard LAG	16
Maximum # of members in a Brocade LAG	16
Maximum # of switches in a Fabric cluster mode	32
Maximum # of switches in Logical cluster	32
Maximum # of ECMP Paths	8
Maximum # of trunk members for fabric ports	16
Maximum # of VLAG groups	512
Maximum # of member ports in a VLAG	64
Maximum # of nodes in a VLAG	8
Maximum # of member ports per VLAG per Node	16

Maximum # of Management ACL	256
Maximum # of VMs supported in VM Aware Network Automation	8,000
Maximum # of ARP Entries	8,000
Maximum # of Unicast IPv4 routes in the hardware	4,000
Maximum # of OSPF areas	20
Maximum # of OSPF routers in a single area	64
Maximum # of OSPF adjacencies	100
Maximum # of OSPF routes	4,000
# of OSPF Interfaces	100
# of OSPF enabled subnets	100
# of local subnets in a single area	100
Maximum # of routes in SW	4,000
Maximum # of static routes	1,000
Maximum # of dynamic routes	4,000
Maximum # of VRRP instances per system	256
Maximum # of VRRP instances per interface	8
Maximum # of routers participating in a VRRP-E session	4
Maximum # of routes with ECMP supported	4,000
Maximum # of IP interfaces per system (VE Interfaces)	256
Maximum # of VRF per node	32
Maximum # of I-BGP peers	50
Maximum # of E-BGP peers	50
Maximum # of BGP routes in HW	4,000
Maximum # of RIB IN Routes	50,000
Maximum # of RIB OUT Routes	100,000
Maximum # BGP Peer Group	50
Maximum # of UDLD enabled interfaces	64
Maximum # of PVLAN domain supported	1,000
Maximum # of Secondary vlans per PVLAN supported	24
Maximum # of primary vlans per PVLAN supported in promiscuous mode	24
Sum total of all the rules across L2 Ingresss ACLs	508
Sum total of all the rules across L2 Egress ACLs	124
Sum total of all the rules across L3 Ingresss ACLs	508
Sum total of all the rules across L3 Egresss ACLs	511

*** Enabling 3500 VLANs in trunk mode on all the ports of the system might lead to system instability. Brocade recommends that “trunk mode VLAN all” should be restricted to a handful uplink ports or a set of selected ports where it is desirable to carry all the VLAN trunks.

Compatibility

In VCS Fabric mode, 4.0.1 has the connectivity to only Brocade NOS 4.0.1 into same fabric. If 4.0.1 connects to other than Brocade NOS 4.0.1, the VCS Fabric will be segmented. Regarding to the connectivity to standard Ethernet switch in VCS Fabric mode, please refer to the Network OS Administrator Guide.

Documentation Updates

When using the NOS 4.0.1 documentation, the embedded DCB switch module is equivalent to the VDX 6720 except where noted in the release note document. The most recent NOS documentation manuals are available on MyBrocade: http://my.brocade.com/

Following three NOS 4.0.0 documents are recommended references.

1. Network OS Administrator Guide: MK-99CB2500041

2. Network OS MIB Reference: MK-99COM058

3. Network OS Command Reference: MK-99CB2500042

The embedded DCB switch module uses additional commands to those listed in the Network OS Command Reference. Refer to Appendix A for additional commands.

IMPORTANT NOTES

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

Command Line Interface

- Some commands will not produce paginated output.
- Break command is not supported. Pl use ctrl-c as an alternative
- For certain commands (including no form with some commands), "?" will show unsupported additional options.
- Tab completion and <ctrl>-c (cancel) does not work for some commands.
- For some commands, "switchId" and "all" options are not applicable in this Brocade Network OS release but are still shown as options. These will be applicable and supported in future Brocade Network OS releases.

- Some CLI commands will generate an "Error:Access denied" message upon failure. This means the operation failed on the switch and may not be related to permissions.
- The "no" command always exists for all roles even if it is not required.
- Some no commands will execute without mandatory parameters that were originally used for configuration. Some needs mandatory parameters though help message does not suggest same
- Some no commands may produce an incorrect error message upon error.
- Incorrect range might be displayed in the help text for some of the show commands.
- Interface range command is not supported on breakout ports. Range command is not supported across multiple slots of the chassis
- System does not warn user on deleting the ip config when vrf is configured
- show interface stats brief does not distinguish loopback interfaces across rbridges
- Redistributed connected/static may be shown twice as part of config
- Some unsupported debug commands may be seen in NOS 4.1.0. Brocade recommends not to run them on switches:
  - Show confd-state -, for debugging purpose only.
  - Show parser dump -, for debugging purpose only
  - Show notification stream -, for debugging purpose only
  - Show features - no use
  - Show ssm -, for debugging purpose only.
  - Autoupgrade command in config mode
- 'snmp-server context CONTEXT_NAME vrf-name VRF-NAME command
- During "copy running-config startup-config" or "copy support" user might see occasional CPU spikes (to ~30-40%).
- While unconfiguring non-existent configs, for some features, "Error: Access Denied" may be displayed even though it is a no-op.

- Interface specific static arp entries are not shown when using show running command for an interface.
- show mac-address-table command on console with include option cannot be aborted with a break/ctl-C. Use a telnet session for the same.
- For ip access lists, display filtering based on sequence number alone does not work as expected.
- Security CLIs: In FC mode: the following are under rbridge-id context unlike earlier release
  - fcsp
  - secpolicy
  - system-monitor moves to rbridge context but system-monitor-mail is still in global mode
- DHCP/ipv6 autoconfig were moved from rbridge context in 3.x to mgmt. interface context in 4.x
- Though ICMPv6 RA guard CLI is available on all platforms , it is supported only for 6710/20/30
- TACACS/Radius local behavior is now changed and currently reflected using 'local backup'
- Do not use CLI 'no spanning-tree shutdown' from the vlan context from rspan-vlan
- Do not use lldp iscsi-priority' (and a couple of other similar CLIs from the same context) needs to be blocked on destination mirror port.
- "show chassis" output may show the PSU part number as "Unknown" after removal & re-insertion of the PSU
- Under certain scenarios, output of "show qos rcv-queue multicast ten <>" may not show accurate count of drops
- Certain oscmd commands may not work or give a different output under admin login
- Netconf commands 'debug internal rate-limit-delay' may fail
- debug ip bgp prefix-list <option> , debug ip bgp neighbor does not work
- 'no' command for 'qos map dscp-cos' does not work
- On rare scenario, configuration may not be applied to hardware on power-cycling the chassis

Breakout functionality on 40Gb ports

- Going to and from Breakout mode requires reload.
- The LED state for a breakout interface is deterministic. For all other supported platforms, the LED state for a breakout interface is not deterministic.
- In breakout mode, there is only SFP and no per-breakout media information. The show media command will displays the same media information for all breakout interfaces. The TX Power Field in the show media command is not supported by the 40G optics.

Restrictions for Ports in 1G Mode

- RMON stats are calculated incorrectly for packet sizes 64-127 bytes.
- Brocade Trunks cannot be formed with 1G, as all Brocade Trunks should be 10G.
- A LAG cannot be created between 1G and 10G ports.

Platform

- This platform does not support IP fragmentation. MTU errors are reported in “show interface” as “Errors” under the “Transmit Statistics”
- After "chassis disable" please wait for 60 seconds before doing the next "chassis enable".
- Chassis-name is limited to 15 characters
- 1G links must have auto-negotiation enabled. 1G links without auto-negotiation are not supported.
- Current 1G copper SFP's do not support exchanging flow-control settings during the auto-negotiation process. It is recommended to configure static mode of configuration of flow-control on both the ends of the desired link.
- System verification/diagnostics performed on a switch will require a reboot.
- Configuration of more than one In-band management port on a single switch is not recommended.
- Under certain stress conditions 'copy support' command might time out for some modules. In such cases it is recommended to retry 'copy support' with higher timeout multiplier value.

- It is highly recommended to copy configuration file to running-config and then save the running-config to startup-config, instead of directly copying the external configuration file to startup-config, especially when using fabric distributed features such as Zoning, VM Aware Network Automation and Virtual IP.

Virtual IP Address Support

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

Security, ACLs, Authentication, Authorization

- Netconf session may get closed for get-vlan-brief
- Login authentication service (aaa authentication login cli):
  - With "local" option specified as secondary authentication service, local authentication will be tried only when the primary authentication service (TACACS+/Radius/LDAP) is either unreachable or not available.
  - Behavior of "local" option in pre-4.1.0 releases is changed to the "local-auth-fallback" option.
  - When login authentication configuration is modified, the user sessions are not logged out as in pre-4.1.0 releases. All connected user sessions can be explicitly logged out using "clear sessions" CLI.
- ACLs are not supported for egress traffic flows
- Configuring TACACS+ or RADIUS without a key is not supported. If no key is configured, the switch uses a default key of "sharedsecret".
- There is a possibility that locked user accounts will get unlocked after a reboot if the running-config (before reboot) is different from startup-config of user accounts.
- Encrypted text (taken from running-config of any user account password with encryption turned on) should not be used as input for clear-text password for the same user. This may result in login failure of the user subsequently.

- There is no upper limit for the number of rules that can be added to a management access-list. But when the ACL is applied to a management interface, only the top 256 rules will be applied if the ACL contains more than 256 rules.

- Access to ONLY the following Active Directory (AD) servers is supported by Brocade LDAP client:

  - Windows 2000

  - Windows 2003

  - Windows 2008 AD

- The DNS configuration is primarily used for LDAP. It should be noted that DNS look-up will not be used by PING, Traceroute or any other services. These services will still require specifying the actual IP address.

- When more than 250 rules ACL's are configured (over supported scale), they may be partially installed & effective

- A hard-drop ACL rule on VDX-6740 may not drop UDLD packets

- Counter for hard-drop ACL may not count accurately

- Even though IGMP snooping feature is supported over VLAG, all the multicast data traffic will be forwarded only over the primary.

- When a MAC ACL with several clauses is applied to a port-channel which is a member of 750 or more VLANS, the MAC ACL counters may take several minutes to be enabled due to processing load associated with such configurations.

- Deny / Harddrop ACL on VE does not work when pkt ingresses from TRILL port

- There is very limited support of bulk calls are available in 4.0.1 release

  - Vlan- create/delete

  - Physical interfaces- Only selected set of configurations are supported (mode setting trunk/access)

  - SVI

- To configure radius authentication, it is required to open up the port for Radius Accounting too on Firewall.

Management Services

- SNMP is not aware of cluster. Hence if we query 1 node through SNMP, we will get the info related to that particular node only

SPAN & RSPAN

- CPU-originated packets cannot be output spanned.
- SPAN is supported only within a port-group on the VDX 6720.
- If SPAN has to be supported to multiple locations, please use RSPAN on vlan.
- On VDX 6720, only one port per port group can be configured as destination port for ingress spanning.
- On 6720, only one port per port group can be configured as destination port for egress spanning.
- On 6720,ISL port cannot be source or a destination SPAN port.
- On 6720, Inter-chip port spanning is not allowed.
- Spanning of LAG port is not supported. To span a LAG, user should individually enable spanning on all the member ports of the LAG.
- A profiled port cannot be a SPAN destination.

ICMPv6 RA Guard

- This feature is only supported by Brocade Fixed Port Switches VDX 6710, VDX 6720 and VDX 6730.

Trunking

- For the rest of the VDX platforms, Brocade trunk (BTRUNK) has a maximum throughput of 80G. Full link utilization of 8 ports in a trunk group is achievable with larger packet size (>128 Bytes).
- 40G BTRUNK is supported between embedded DCB switch and VDX 6740/VDX 6740-T (2-port trunk)

VCS

- Loopback connection is not supported in VCS mode. If a loopback connection is done, those interfaces become ISL interfaces.
- Fabric Cluster Mode:

- When a new switch is added to an existing VCS fabric and if the new switch takes the role of principal node, the other switches in the fabric will receive the configuration of the distributed features such as Virtual IP and VM-Aware Network Automation from the newly added switch. This will cause the existing distributed configuration to be overwritten by the newly added switch in the principal role. This can be avoided by following the new switch addition procedures in the Admin Guide.

- After a cluster reboot, Brocade recommends to do both "show fabric all" and "show vcs" to ensure that cluster is entirely formed without any issue. User might see that 'show vcs' takes an additional 2-3 minutes to show all participating switches. This is an existing behavior and doesn't affect data path functionality in most cases.

- "show fabric isl" & "show fabric trunk" may show the interfaces in random order without sorting

VLAG

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

- When configuring LACP LAG between VDX & non-Brocade switches it is highly recommended to enable the VLAG ignore-split on the VDX. Ignore split option is enabled by default in Brocade Network OS v4.1.0.

MAC Learning Considerations in VCS

- The CLI command "clear mac-address-table" has been enhanced to support clearing the mac-addresses associated with vLAG's. This command can be used to sync mac-address-tables of the VCS member switches.

- Post 3.x releases, FPMA mac addresses are not shown in "show mac-address-table dynamic". User can use ' show mac-address-table count' together to get this output

- Post 3.x releases, Internal Mac-addresses are shown in "show mac-address-table" output to support L3 use cases. The sync across the VCS has to be observed using "show mac-address-table dynamic".

- Under rare circumstances, end user might see mac address sync up issues on few nodes of a cluster (where 1 or more MAC addresses might be missing in some nodes). Brocade recommends to do "clear mac-address-table dynamic" in such cases.

- Static mac addresses will be displayed even when interfaces are down. This may cause blackholing of the traffic.

- There are 3 operational enhancements w.r.t VLAN Interfaces

- Removal of shutdown/ no shutdown at vlan interface level.
- Removal of vlans information entirely from 'show ip interface brief' cmd
- Output of 'show vlan brief' reflects the 'State' of VLAN as ACTIVE/INACTIVE (along with inactive reason . 'member port down') based on member ports' state.

- Under certain conditions, MAC addresses may not be learnt even though ARP's may be learnt for those same MAC addresses

UDLD

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

STP/DiST

- VDX does not support tunneling non-standard BPDUs and thus IEEE BPDUs (0180:C200:0000) generated as tagged packets in STP/RSTP/MSTP modes may not be tunneled successfully across VCS fabric. However, VDX supports tunneling standards' based BPDUs such as untagged IEEE BPDUs and tagged or untagged PVST BPDUs (0100:0CCC:CCCD). Post 3.0.1, the tagged IEEE BPDU can be tunneled across VCS fabric using command: "tunnel tagged-ieee-bpdu" under interface configuration.
- In Fabric Cluster mode, global spanning-tree configurations (STP enable, STP Vlan configurations, STP over vLAG configurations) have to be performed in all the switches in VCS at the same time. For example, to run spanning-tree, it has to be enabled on all the switches including switches that don't have any edge ports. In case one want to enable the feature on a larger cluster size (> 8 nodes- Scale limits=8), to keep STP states/behavior consistent you need to enable the global configs (1) on all the nodes in the cluster and enable interface config (2) only up to 8 nodes in a cluster. For LC mode, 1 automatically happens and you do interface enabling on up to 8 different rbridges.
- By default global spanning-tree and interface level spanning-tree will be disabled, user has to explicitly enable on the desired ports. vlan spanning-tree state is default enabled
- BPDU tunnel configurations are permitted only when spanning-tree is disabled in VCS.

- For cisco proprietary Per Vlan Spanning Tree protocols (PVST and RPVST) user needs to configure Brocade switch to send BPDU on Cisco multicast destination mac address "0100.0ccc.cccd" for non-native vlans. By default, NOS 4.1.0 software use's brocade "0304.0800.0700" multicast mac to send BPDU's on non-native vlans.

- Since NI/FI/Cisco boxes use Cisco multicast mac address to send spanning tree BPDU on non-native vlans, this configuration is needed in VDX switches to interoperate. This is an interface specific configuration

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

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

Edge Loop Detection (ELD)

- ELD is supported on the edge ports that are connected either by end-hosts OR another switch OR another VCS.

- Maximum of 256 instances are supported in a fabric. Instance is counted per interface per vlan.

- To limit the number of instances utilized, it is recommended to enable ELD on only 1 vlan per interface.

- ELD is supported for edge interfaces connected to hosts too.

AMPP and Port-Profiles

- Port-profile status does not reflect the remote interface info in VCS fabric mode.

- Native VLAN support inside AMPP does not honor the global enable/disable flag.

- SPAN destination port cannot be a profiled port.

- All AMPP features that were supported both a physical interface and a VLAG.

- Brocade recommends deleting all manually created port-profiles when migrating from a legacy AMPP environment to VM Aware Network Automation.
- Vmkernel related port-profiles may unapply/reapply during HA resulting in vmotion failures.

vCenter

- VM-Aware Network Automation will work only with VMware vSphere version 4.0, 4.1, 5.0, 5.1 and 5.5.
- Receiving more than five vCenter events within a span of 30 seconds, results in asset discovery getting initiated. Post discovery cluster configuration will be in sync with vCenter.

QoS

- It is recommended to use the same CoS Tail drop threshold on all members of a port-channel to avoid unpredictable behavior.
- Asymmetric pause is supported on 1G port interfaces
- VDX 6740 supports 3 PFC queues
- Flow control is disabled by default on all interfaces.
- Trust- support is available only standalone mode, no VCS mode
- DSCP to CoS Mutation- all platforms
- DSCP to Traffic Class Mutation -all platforms
- DSCP to DSCP Mutation
- Random Early Discard (RED)
- Priority 7 is reserved for control traffic on VDX switches. User data traffic should use priorities 0 through 6.
- Brocade VDX architecture prioritizes Unicast traffic over Broadcast or Multicast traffic under port congestion.

VRRP

- VRRP and VRRP-E cannot be enabled together on VDX 6740 and 6740T platforms. Command “protocol vrrp-extended” is added to specifically enable VRRPE.
- VRRP-E global sessions may get disabled after firmware upgrade

- Large VRRP config may increase config download time

OSPF

- Graceful restart is not supported

BGP

- Following BGP features are not supported in this release:
  - Graceful Restart
  - AS Confederation
  - Outbound Route Filtering capability
  - Extended Community Filter support
- BGP Aggregate route is preferred over direct network
- Standard and Extended community may be allowed to be configured on same interface

L2/L3 Multicast

- The following PIM features are not supported in this release:
  - Non-Stop Routing (NSR)
  - IP version 6
  - Prefix list
  - Configuring the switch as the BSR (Bootstrap Router) candidate.
  - Configuring the switch as the Rendezvous Point or Rendezvous Point candidate
- The Rendezvous Point (RP) must be configured outside the VCS cluster.
- All PIM enabled routers should be directly connected to RP
- IGMP Snooping must be enabled in all the switches in the VCS cluster.
- IGMP timers configured on PIM enabled L3 interface are not considered over the timers on VLAN
- CLI incorrectly allows same interface to be selected as incoming and outgoing interface for PIM-DR

- IGMP leave from one receiver will affect other receivers if connected through a vLAG
- IGMP join does not get forwarded via vLAG on shutting the primary port until general query is received
- PIM OIF list may not be updated when static IGMP group from VE is removed

VRF

- Route leaks across VRF is not supported
- Management VRF is not supported
- VRF lite supports OSPF and static routing but not BGP
- On configure VRF on an interfaces, all previous IP config would be lost

Interoperability

- In a VPC environment where the Brocade VDX side has the active LACP settings and the Cisco side has the passive settings on the vLAG, the port-channel takes over 30 seconds to come up. Workaround: Reverse the settings and have the Brocade VDX LACP settings passive and the Cisco side set as active. The port channel will then restore after about 10 seconds.
- There is a compatibility issue between Brocade and Cisco chassis that can cause an LACP protocol timeout. If you have a Brocade VDX 6710 and a C24 VDX cluster and two Cisco Nexus 5k chassis configured in a VPC cluster using a combination of 1G fiber copper links, after shutting down links on the Cisco side, about 10 seconds of traffic loss can occur. The shutdown operation of the Nexus 1G port does not shut down the transmitter, so the Brocade VDX 6710 port is not able to detect link down. This leads to LACP protocol timeout.
- When interoperating with Brocade 8000, it is recommended to set the mac-aging time to 0 on the VDX switch to prevent any adverse impact caused by certain errors generated by the Brocade 8000.
- PVST-RPVST interop may not work between VDX and FCX/ICX

Miscellaneous

- When using STPoVCS, it is highly recommended to avoid "Peer-switch" configuration on the Cisco Nexus vPC configuration for the best performance.
- Brocade VDX switches load balance internal and external traffic based on hash functions using standard network headers as keys. Due to this implementation, users may experience traffic imbalance depending upon application flow definition.

- Packet drops will be seen for a short duration due to routing changes with link flaps and/or node failovers.
- On both ISL and Edge ports, sFlow sampling is supported only in inbound direction.
- SFlow collectors are not queried in snmp v1, v2 & v3 versions
- If multiple VLANs are configured on a switch, then in order to enable certain features such as IGMP or PVST it is recommended that specific features be enabled on a per-VLAN basis instead of enabling them globally.
- "clear ip route all" need to be issued once the maximum number of routes supported by a router is exceeded.
- SNMPset operation is not fully supported
- Under rare conditions, the switch may bootup with default configuration on power-cycling the switch
- Firmware downgrade is not blocked if the scale configured would not be supported in the downgraded release
- Please make sure to not have large no of unreachable TACACS+ accounting server configured, else it might cause unit to reboot. This issue is hit only with large config (4K vlan etc and 20K lines or config)
- The followings show unsupported features for the embedded DCB switch module in NOS 4.0.1 against Brocade Network OS release.
    - Logical Chassis mode
    - FCoE features
    - Private VLAN(PVLAN)

3

Feature Description

In this chapter, the New/Enhanced feature of the Network OS is described.

Network OS v4.0.1 includes support for new features such as DiST (STPoVCS), PVLAN, Uni-directional Link Detection (UDLD), Flow-based QOS, Flow-based sflow, RSPAN, Border Gateway Protocol (BGP), Inbuilt packet capture utility (PCAP), Management Services enhancements, DHCP IP Helper, DHCP-based Firmware download (DAD . DHCP Automatic Deployment), VRRP-E across VCS fabrics for embedded DCB switch module.

And, Network OS v4.0.1 deprecates 'ip gateway-address' CLI and replaces it by 'ip route' CLI

Feature Descriptions

Distributed Spanning Tree Protocol (DiST/STPoVCS)

Network OS v4.0.1 and later supports any version of STP to run in VCS mode and function correctly between interconnecting VCSs, or between VCS and other vendor's switches. This feature is called Distributed Spanning Tree Protocol (DiST).

The purpose of DiST is:

- To support VCS to VCS connectivity and automatic loop detection and prevention.
- To assist deployment plans for integrating with the legacy xSTP enabled switches in the network, for eventual replacement of such switches with fabrics.
- Support following flavors of spanning-tree protocol: STP, RSTP, MSTP, PVST+, and RPVST+

UDLD (UniDirectional Link Detection)

UniDirectional Link Detection (UDLD) protocol is a nonstandard Layer 2 protocol that detects when a physical link becomes unidirectional by means of the exchange of UDLD protocol data units (PDUs). A unidirectional loop can lead to the creation of a loop in a network, which the Spanning Tree Protocol (STP) could inadvertently allow to occur.

This proprietary UDLD protocol is compatible only with the Brocade IP product line UDLD protocol. It can be configured on all physical ports in Standalone mode and on all physical edge ports in a Virtual Cluster Switching (VCS) environment. When a physical link is detected as unidirectional, traffic is blocked on the physical link. When a unidirectional link is detected as bidirectional, traffic is automatically unblocked on the physical link.

Flow based features (sFlow/QoS)

Flow-based sFlow is used to analyze a specific type of traffic (flow based on access control lists, or ACLs). This involves configuring an sFlow policy map and binding it to an interface.

RSPAN

RSPAN extends SPAN by enabling remote monitoring of multiple switches across your network.

The traffic for each RSPAN session is carried over a user-specified RSPAN VLAN that is dedicated for that RSPAN session in all participating switches. The SPAN traffic from the sources is copied onto the RSPAN VLAN and then forwarded over trunk ports that are carrying the RSPAN VLAN to any RSPAN destination sessions monitoring the RSPAN VLAN.

Border Gateway Protocol (BGP)

Border Gateway Protocol (BGP) is an exterior gateway protocol that performs inter-autonomous system (AS) or inter-domain routing. It peers to other BGP-speaking systems over TCP to exchange network reachability and routing information.

Support for BGP on NOS platforms is for BGP4 (compliant with RFC 1771 and 4271), and provides the following:

- Connectivity from the VCS to a core/external network or cloud

Administrative distance for BGP routes cannot be changed using route-map configuration

Virtual Routing & Forwarding Lite (VRF-Lite)

VRF is a technology that controls information flow within a network by isolating the traffic by partitioning the network into different logical VRF-domains. Every VRF capable router supports one routing table for each VRF instance. Typical full-blown implementations of VRFs are designed to support BGP/MPLS VPNs, whereas VRF-lite implementations typically are much simpler with moderate scalability (as compared to BGP/MPLS VPNs). Brocade NOS v4.0.x will support VRF-Lite.

Inbuilt packet capture utility (PCAP)

The packet capture utility, executed by means of the "capture packet interface" command, enables capturing packets from an interface that are to / from CPU, as well as transit packets if a trap is enabled by means of ACL logging. This command can provide significant help in debugging, especially for capturing & viewing Layer 2 TRILL and Layer 3 packets using "show capture packet". Moreover, the packets captured can be saved & exported in the PCAP format for enabling to be viewed offline using commonly used tools like WireShark.

Management Services

Brocade Network OS v4.0.1 supports various enhancements to existing management services, including TACACS+ and SNMP. SNMP supports v1/v2c/v3. SNMP is not cluster-aware. New supports added for 4.0.1 are:

- UDLD Traps

- Community MIB
- Password encryption for SNMPv3

Netconf Bulk config support is available for limited yang calls.

IPv4 and IPv6 Management Services

Network OS v4.0.x supports various IPv4 and IPv6 Management Services, including In band Management using IPv4 addressing, Secure Syslog and IP Static Routes in VCS. The services are:

- Inband Management is available. In band Management can be used for switch management and SNMP functions such as sourcing SNMP traps.
- IPv6 addressing is available on all VDX switches for the management interfaces. VDX switches can be managed through IPv6 management network and IPv6 services such as ping, traceroute, and DNS DHCPv6 are supported
- DNS Server configuration allows either ipv4 or ipv6 DNS server addresses to be configured on the switch.
- Static Routing is available on the VDX devices in the VCS fabric. Static routes can be used to avoid overhead of running dynamic protocols in simple networks, overwrite the routes calculated by routing protocols, inject networks that don't have routing protocol enabled and keep default back up paths when routing protocol instabilities affect the network.
- Network OS v4.0.x supports ICMPv6 RA Guard only on VDX6710/20/30 platforms to mitigate any attack vectors based on illicit ICMPv6 Router Advertisement messages.

Secure Syslog

Secure Syslog facility using Transport Layer Security (TLS) Protocol is available now for both IPv4 and IPv6 server addressing. User can choose both secure syslog server and non-secure syslog server.

LLDP Protocol

Network OS v4.0.x supports two LLDP neighbors.

ACLs

In Brocade Network OS v4.0.x, both Ingress Layer 2 MAC access control lists (ACLs) and Layer 3 IP access control lists are supported. Brocade Network OS v4.0.x also supports standard and extended ACLs.

MAC ACLs are supported on the following interface types:

- Physical interfaces
- Logical interfaces (LAGs)
- VLANs

IP ACLs are supported on the following interface types:

- Logical interfaces (LAGs)
- Management interfaces
- VLANs

The following QoS features are supported in Brocade Network OS v4.0.x.

- Layer2 and Layer3 QoS
- Traffic Policing
- Control Traffic Prioritization
- BUM Traffic Control
- Enhanced Transmission Selection (ETS) and Priority Flow Control (PFC)
- DSCP Trust
- DSCP to CoS Mutation
- DSCP to Traffic Class Mutation
- DSCP to DSCP Mutation
- Random Early Discard (RED)

The NOS v4.0.x release will support ACL-based QoS.

Open Shortest Path First (OSPF) in VCS

Open Shortest Path First (OSPF) is a link-state routing protocol supported in NOS v4.0.x. OSPF routing is supported on all VDX Switches.

OSPF can be configured only in a Virtual Cluster Switching (VCS) environment. OSPF can be configured on either a point-to-point or broadcast network. OSPF can be enabled on the following interfaces: GigabitEthernet, TenGigabitEthernet, FortyGigabitEthernet, and VE.

In NOS v4.0.x OSPF interfaces can be configured as passive-interfaces. A passive interface does not send or process received “hello’s” thus not forming adjacencies or advertising routes.

VRRP and Brocade VRRP-E in VCS

In Network OS v4.0.x, VRRP is available on all VDX Switches and supports two versions of VRRP protocol for IPv4:

- Standard VRRP—the standard router redundancy protocol, VRRP v2 supports the IPv4 environment. Also, the Brocade version of standard VRRP is compliant with RFC 3768.
- VRRP-E (Extended)—A Brocade proprietary protocol similar to standard VRRP. It does not interoperate with VRRP.

The VRRP and VRRP-E protocol supports:

- VRRP—FortyGigabitEthernet, TenGigabitEthernet, GigabitEthernet, and VE.
- For VRRP-E—VE ports only.

IP Multicast in VCS

Network OS v4.0.x supports PIM-SM, which is the most commonly deployed flavor of PIM with IGMP. PIM-SM is more effective in large networks that are sparsely populated with hosts interested in multicast traffic.

IP multicast is supported only on the VDX6740/T and VDX 8770. PIM-SM can be enabled on the following interfaces

- GigabitEthernet
- TenGigabitEthernet
- FortyGigabitEthernet
- VE Interface

Port Security

Following mac port security features can be used to enhance the security at layer 2:

Mac-limiting - This feature can be used to configure a group of MAC addresses that are allowed to access a given interface. When secure MAC addresses are assigned to a secure port, the port does not forward packets with source addresses outside the group of defined addresses. Another method is to limit the number the MAC addresses that are allowed on a given interface. This feature is available on all VDX platforms.

OUI based port security – This security feature can be used to limit the MACs allowed based on a vendor OUI. When user configures an OUI on a secure port, only the traffic that is coming from the devices which are part of configured OUI will be forwarded. This feature is available only on the VDX 8770 and VDX 6740/6740T platforms.

Port Security with Sticky MAC - Port security with sticky MAC is similar to static secure MAC in functionality, but sticky macs are dynamically learnt MACs. In this security mechanism dynamically learnt MAC will be retained even after link goes down. This feature is available on all VDX platforms.

ICMP Rate Limiting

The following features are supported:

Enabling/Disabling ICMPv6 error messages for an unreachable address

This feature allows the ability to enable or disable generation of a Destination Unreachable message in response to a packet that cannot be delivered to its destination for reasons other than congestion

Enabling ICMPv6 Echo Reply messages

This feature allows the ability to enable or disable sending of an Echo Reply message in response to an Echo Request message sent to an IPv6 multicast or anycast address. NOTE: The number of responses may be traffic conditioned to limit the effect of a denial of service attack.

ICMPv6 rate limiting

This feature allows for rate limiting the transmission of ICMP responses. User can limit the rate at which ICMP/ICMPv6 messages are sent out on a network.

Protecting Against TCP SYN Attacks

This feature protects against TCP SYN attacks by allowing the configuration of the Brocade VDX devices to drop TCP SYN packets when excessive numbers are encountered. Threshold values can be configured for TCP SYN packets that are targeted at the router itself or passing through an interface, and drop them when the thresholds are exceeded.

A. Software Fixes

In this chapter, Software Fixes information for this release is described.

Software Fixes

Closed with Code Change in Network OS v4.0.1_hit1

This section lists the defects closed with a code change as of February 4, 2015 in Network OS v4.0.1_hit1.

Item	Descriptions	
#1	Summary	Switch Module: Power on retry over event happens intermittently at low temperature..
	Symptom	The fail message of "Switch Module: Power on retry over" would be logged on system event log (SEL) at low temperature (about 5 degree C). When this issue happens, user can not log-in to the switch CLI.
	Technical Severity	Medium
	Reported in release	NOS 4.0.1_hit

Hitachi Data Systems

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

Regional Contact Information

Americas
+1 408 970 1000
info@hds.com

Europe, Middle East, and Africa
+44 (0) 1753 618000
info.emea@hds.com

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

MK-99CB2500044-01