



Release Notes for VSP 8600

Release 8.0
9036330-00 Rev AG
August 2022

© 2022, Extreme Networks
All Rights Reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see:
www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>

Contents

Chapter 1: About this Document	5
Purpose.....	5
Conventions.....	5
Text Conventions.....	5
Documentation and Training.....	7
Getting Help.....	8
Providing Feedback.....	9
Chapter 2: New in this Release	10
Chapter 3: Filenames for this Release	21
Chapter 4: Documentation Changes	22
Chapter 5: Upgrade Paths and Considerations	23
Supported Upgrade Paths.....	23
Upgrade Considerations.....	23
Downgrade Considerations.....	23
Real Time Clock.....	24
MACsec on 100 Gb Devices.....	25
Software Version Mismatch Generates Warning Messages when Installing a New IOC Module...	25
Chapter 6: VSP 8600 Series Hardware and Software Compatibility	26
Chapter 7: Software Scaling	27
Layer 2.....	27
IP Unicast.....	28
DvR.....	29
Layer 3 Route Table Size.....	29
IP Multicast.....	30
Filters, QoS, and Security.....	30
Fabric Scaling.....	30
OAM and Diagnostics.....	31
Chapter 8: Important Notices	32
Improved MAC Statistics Diagnostics.....	32
8624XS IOC Module Power Consideration.....	32
Feature Licensing.....	33
Limitations on license filename size.....	34
High Availability (HA).....	34
Network Load Balancing (NLB).....	35
System Name Prompt vs. IS-IS Host Name.....	35
VRRP IDs.....	35
Chapter 9: Known Issues and Restrictions	37
Known Issues and Restrictions.....	37

Contents

Filter Restrictions and Expected Behaviors.....	45
Chapter 10: Resolved Issues	47
Appendix A: Related Information	49
New MIBs.....	49
Modified MIBs.....	52

Chapter 1: About this Document

This section discusses the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

This document describes important information about this release for supported VSP Operating System Software (VOSS) platforms.

This document includes the following information:

- supported hardware and software
- scaling capabilities
- known issues, including workarounds where appropriate
- known restrictions

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that can be used throughout this document.

Table 1: Notes and warnings


Icon	Notice type	Alerts you to...
	Tip	Helpful tips and notices for using the product.

Table continues...





Icon	Notice type	Alerts you to...
	Note	Useful information or instructions.
	Important	Important features or instructions.
	Caution	Risk of personal injury, system damage, or loss of data.
	Warning	Risk of severe personal injury.

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	<p>Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command.</p> <p>If the command syntax is <code>cfm maintenance-domain maintenance-level <0-7></code> , you can enter <code>cfm maintenance-domain maintenance-level 4</code>.</p>
Bold text	<p>Bold text indicates the GUI object name you must act upon.</p> <p>Examples:</p> <ul style="list-style-type: none"> • Click OK. • On the Tools menu, choose Options.
Braces ({ })	<p>Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command.</p> <p>For example, if the command syntax is <code>ip address {A.B.C.D}</code>, you must enter the IP address in dotted, decimal notation.</p>
Brackets ([])	<p>Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command.</p> <p>For example, if the command syntax is <code>show clock [detail]</code>, you can enter either <code>show clock</code> or <code>show clock detail</code>.</p>

Table continues...

Convention	Description
Ellipses (...)	<p>An ellipsis (...) indicates that you repeat the last element of the command as needed.</p> <p>For example, if the command syntax is <code>ethernet/2/1 [<parameter> <value>]...</code>, you enter <code>ethernet/2/1</code> and as many parameter-value pairs as you need.</p>
<i>Italic Text</i>	<p>Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.</p>
Plain Courier Text	<p>Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages.</p> <p>Examples:</p> <ul style="list-style-type: none"> • <code>show ip route</code> • <code>Error: Invalid command syntax</code> <code>[Failed][2013-03-22 13:37:03.303</code> <code>-04:00]</code>
Separator (>)	<p>A greater than sign (>) shows separation in menu paths.</p> <p>For example, in the Navigation tree, expand the Configuration > Edit folders.</p>
Vertical Line ()	<p>A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command.</p> <p>For example, if the command syntax is <code>access-policy by-mac action { allow deny }</code>, you enter either <code>access-policy by-mac action allow</code> or <code>access-policy by-mac action deny</code>, but not both.</p>

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and software compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other resources](#) such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

[Extreme Portal](#)

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

[The Hub](#)

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

[Call GTAC](#)

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Service Notifications

You can subscribe to email notifications for product and software release announcements, Vulnerability Notices, and Service Notifications.

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form (all fields are required).
3. Select the products for which you would like to receive notifications.

*** Note:**

You can modify your product selections or unsubscribe at any time.

4. Select **Submit**.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.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.

Chapter 2: New in this Release

The following sections detail what is new in VSP 8600 Series Release 8.0.

New Transceivers and Components

This release introduces support for the following transceivers and cables. These transceivers and cables have been consolidated and qualified for use in Extreme Networks platforms, with enhanced diagnostics for transceivers. Enhanced diagnostic information includes power-on counters, comparison statistics for actual Tx and Rx dB values versus low alarm values, and the associated logging for these enhancements.

- 25 Gb transceiver modules
 - 25G SR SFP28 100m LC connector Multi-Mode MSA (PN: 25G-SR-SFP100M)

*** Note:**

25 Gb transceiver module use on VSP 8600 Series requires the 100G QSFP28 SFP28 adapter (PN: 10506) and channelization enabled on a 100 Gb port.

- 40 Gb transceiver modules
 - 40G SR4 QSFP+ 150m MPO connector Multi-Mode MSA (PN: 40G-SR4-QSFP150M)
 - 40G LM4 QSFP+ 160m 160m MMF. 1km SMF LC connector Multi-Mode MSA (PN: 40G-LM4-QSFP160M)
 - 40G LR4 QSFP+ 10km LC connector Single-Mode MSA (PN: 40G-LR4-QSFP10KM)
 - 40G ESR4 QSFP+ 400m 10G-SR interoperable MPO connector Multi-Mode MSA (PN: 40G-ESR4-QSFP400M-NT)
- 100 Gb transceiver modules
 - 100G CWDM4 QSFP28 2km LC connector Single-Mode MSA (PN: 100G-CWDM4-QSFP2KM)
 - 100G LR4 QSFP28 10km LC connector Single-Mode MSA (PN: 100G-LR4-QSFP10KM)
 - 100G SR4 QSFP28 100m MPO connector Multi-Mode MSA (PN: 100G-SR4-QSFP100M)
 - 100G SWDM4 QSFP28 100m LC connector Single-Mode MSA (PN: 100G-SWDM4-QSFP100M)
 - 100G 4WDM QSFP28 10km transceiver single-mode (PN: 100G-4WDM-QSFP10KM)
 - 100G 4WDM QSFP28 20km transceiver single-mode (PN: 100G-4WDM-QSFP20KM)
 - 100G 4WDM QSFP28 40km transceiver single-mode (PN: 100G-4WDM-QSFP40KM)
 - 100G ER4LT QSFP28 40km LC connector single-mode (PN: 100G-ER4LT-QSFP40KM)

- 100Gb PSM4 2km SMF QSFP28 MPO (PN: 10405)
- 40 Gb cables
 - 40G Passive DAC QSFP+ 0.5m MSA (PN: 40G-DACP-QSFPZ5M), 5 meter
 - 40G Passive DAC QSFP+ 1m MSA (PN: 40G-DACP-QSFP1M), 1 meter
 - 40G Passive DAC QSFP+ 3m MSA (PN: 40G-DACP-QSFP3M), 3 meter
 - 40G Passive DAC QSFP+ to 4xSFP+ Breakout 2m MSA (PN: 40G-DACP-QSFP4SFP2M), 2 meter
 - 40G Passive DAC QSFP+ to 4xSFP+ Breakout 3m MSA (PN: 40G-DACP-QSFP4SFP3M), 3 meter
 - 40G Passive DAC QSFP+ to 4xSFP+ Breakout 5m MSA (PN: 40G-DACP-QSFP4SFP5M), 5 meter
 - 40G AOC QSFP+ 3m MSA (PN: 40G-AOC-QSFP3M), 3 meter
 - 40G AOC QSFP+ 5m MSA (PN: 40G-AOC-QSFP5M), 5 meter
 - 40G AOC QSFP+ 10m MSA (PN: 40G-AOC-QSFP10M), 10 meter
 - 40G AOC QSFP+ 20m MSA (PN: 40G-AOC-QSFP20M), 20 meter
 - 40G AOC QSFP+ 100m MSA (PN: 40G-AOC-QSFP100M), 100 meter
- 100 Gb cables
 - 100G Passive DAC QSFP28 3m MSA (PN: 100G-DACP-QSFP3M), 3 meter
 - 100G Passive DAC QSFP28 5m MSA (PN: 100G-DACP-QSFP5M), 5 meter
 - 100G Passive DAC QSFP28 to 4xSFP28 Breakout 3m MSA (PN: 100G-DACP-QSFP4SFP3M), 3 meter

Extreme Networks can announce support for additional optical components in the future. Check the following documents for a complete and current list of supported optical components:

Extreme Networks optical transceivers and components	Extreme Optics website
Compatibility for Extreme Networks SFP, SFP+, SFP28, QSFP+, and QSFP28 transceiver modules with the VSP Series switches	Extreme Hardware/Software Compatibility and Recommendation Matrices

Application Telemetry Host Monitoring

Application Telemetry is enhanced to support host monitoring using Extreme Management Center. You can capture and save the host-related traffic to a separate packet capture (PCAP) file, and then analyze the file.

For more information, see [Monitoring Performance for VOSS](#).

Auto-Complete for "sh" CLI Command

Prior to this release, when you typed `sh` at a command prompt and pressed **Tab** to auto-complete the command, the `sha512` command appeared. Now, when you type `sh` and press **Tab**, the command auto-completes to `show`.

The `sha512` command is now located under the `file-checksum` command.

Auto Forward Error Correction (FEC)

This release introduces the configuration option `auto`, which configures FEC based on port speed and pluggable module type. FEC is used for enhanced error correction when transmitting data over a noisy channel. FEC is not required on 100 Gb or 25 Gb long-range optics because these optics do error correction internally.

For more information, see [Administering VOSS](#).

Authentication for Privileged EXEC CLI Command Mode

For enhanced security, you can enable user authentication to enter Privileged EXEC command mode. Use the `sys priv-exec-password` command to enable password authentication.

After you enable password authentication for Privileged EXEC command mode, the system prompts you to enter a password to access Privileged EXEC command mode from User EXEC command mode. You must enter the same password that you used to log on to the switch.

For more information, see [Configuring Security for VOSS](#).

BGPv6

BGP peering over IPv6 transport uses a BGPv6 peer to exchange IPv6 routes over an IPv6 transport layer. This is different than BGP+, which enables exchange of IPv6 routes over a BGPv4 peer. Also with BGP+, you must use an IPv6 tunnel to install and configure IPv6 routes in an IPv6 Routing Table Manager (RTM). BGP+ uses an IPv4 mapped IPv6 address for the next hop address and requires you to configure IPv6 static routes and install IPv6 routes in an IPv6 RTM where the next hop for the static route is an IPv6 tunnel interface.

For more information, see [Configuring BGP Services for VOSS](#).

BGP Route Redistribution Enhancement

You can now configure a route metric-type for BGP route redistribution using the `ip bgp redistribute` command.

For more information, see [Command Line Interface Commands Reference for VOSS](#).

Distributed Virtual Routing (DvR) Controller

This release introduces support for the DvR Controller.

You must purchase the Layer 3 Virtualization or Layer 3 Virtualization with MACsec feature pack license to use this feature.

Note:

DvR Leaf is not supported.

DvR is only supported in warm standby High Availability mode on the VSP 8600 Series.

Distributed Virtual Routing (DvR) is a technology for router redundancy in a fabric deployment where IP subnets are stretched across multiple switches. DvR provides Default Gateway Redundancy and optimizes traffic flows to avoid traffic tromboning due to inefficient routing, thereby increasing the total routing throughput.

DvR can be deployed in Campus environments for stretching IP subnets between multiple aggregation layer switches and also simplifies data center deployments by introducing a Controller-Leaf architecture. In this architecture, Layer 3 configuration is required only on the Controller nodes,

whereas the Leaf nodes require only Layer 2 configuration. All Layer 3 configuration is automatically distributed to the Leaf nodes by the Controller nodes.

For typical Campus DvR deployments, configure aggregation layer switches as DvR Controllers. Wiring closet access switches are then typically dual-homed to a pair of DvR Controllers.

For more information, see [Configuring IPv4 Routing for VOSS](#).

DvR supports the redistribution of host routes into BGP so that non-Shortest Path Bridging (SPB) routers can benefit from the host accessibility information learned within a DvR domain.

For more information, see [Configuring VLANs, Spanning Tree, and NLB for VOSS](#)

DvR One IP Enhancement

You can now use a single IP address in the subnet for every Controller by configuring the DvR VLAN IP to be the same as the DvR gateway IP.

For more information, see [Configuring VLANs, Spanning Tree, and NLB for VOSS](#).

DHCP Relay scaling enhancement

DHCP Relay forwarding entries scaling increased from 1024 to 2048 on each switch.

For more information, see: [IP Unicast](#) on page 28.

Global Regular Expression Print (GREP) with CLI Show Commands

You can now use Global Regular Expression Print (GREP) with `show` commands to filter the output based on match criteria.

For more information, see [Configuring User Interfaces and Operating Systems for VOSS](#).

iBGP Virtualization

*** Note:**

DEMO FEATURE - iBGP Virtualization is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#)

This release supports internal Border Gateway Protocol (iBGP) peering, for both IPv4 and IPv6, over user-created Virtual Routing and Forwarding (VRF) instances.

For more information, see [Configuring BGP Services for VOSS](#)

IOC Module Pre-Configuration

Using IOC Module Pre-Configuration, you can configure a slot for an IOC Module before you insert the module in the chassis. By specifying the slot and module type, all configuration at the slot or port level become available for that slot. You can issue configuration commands for a specific slot before you insert an IOC Module in that slot.

For more information, see [Administering VOSS](#).

IPv6 Management Applications in Global Routing Table

In this release, the following IPv6 management applications in Global Routing Table (GRT) are supported:

- SSH client

- TFTP client
- Rlogin client

For more information, see [Administering VOSS](#).

IPv6 Loopback Statistics

You can now view statistics and icmp-statistics for IPv6 circuitless interfaces using the `show ipv6 interface statistics loopback` command.

For more information, see [Command Line Interface Commands Reference for VOSS](#).

IPv6 Shortcut Routing

This release supports IPv6 Shortcut Routing. IPv6 Shortcuts function in a very similar manner to IPv4 Shortcuts. Both IPv4 and IPv6 Shortcuts use IS-IS as the Interior Gateway Protocol (IGP) and the link state packet (LSP) for reachability information. However, IPv4 Shortcuts use TLV 135 and IPv6 Shortcuts use TLV 236. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as leaf nodes. IS-IS transports the IPv6 reachability information to remote BEBs and uses the shortest path, calculated by SPBM, for data forwarding.

IS-IS transports the IPv6 routes through TLV 236 in the LSP advertisements. These routes are installed in the Global Routing Table (GRT) with the node from which the LSPs carrying the IPv6 routes are received as the next hop.

For more information, see [Configuring Fabric Layer 3 Services for VOSS](#).

IPv6 Virtualization on VRF and Layer 3 VSNs

This release supports the following IPv6 features on Virtual Routing and Forwarding (VRF) and Layer 3 Virtual Services Networks (Layer 3 VSNs):

*** Note:**

You must purchase the Layer 3 Virtualization or Layer 3 Virtualization with MACsec feature pack license to use Layer 3 VSNs.

- IPv6 Interfaces and IPv6 Static Routes in VRFs
- IPv6 Inter-VSN Routing
- IPv6 ECMP and Alternative Route
- IPv6 VRRPv3
- IPv6 DHCP Relay
- Open Shortest Path First for IPv6 (OSPFv3)
- IPv6 Border Gateway Protocol (IPv6 BGP)
- IPv6 Route Policies and Route Redistribution
- IPv6 Layer 3VSN
- IPv6 IS-IS accept policies
- IPv6 ICMP Ping and Traceroute
- IPv6 Reverse Path Forwarding

*** Note:**

Because IPv6 RSMLT is not virtualized in this release, you cannot enable both RSMLT and an IPv6 interface on the same VRF.

For more information, see [Configuring IPv6 Routing for VOSS](#).

IPv6 scaling enhancements

This release adds IPv6 BGP peers and IPv6 OSPFv3 routes scaling.

For more information, see: [IP Unicast](#) on page 28.

Linux Kernel version 4.9

This release introduces a new Linux kernel on VSP 8600 Series. Introduction of the new Linux kernel has implications on software upgrades and downgrades. For more information, see [Supported Upgrade Paths](#) on page 23.

Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED)

Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) defined in ANSI/TIA-1057, is an extension to the LLDP standard protocol as defined in IEEE 802.1AB. LLDP-MED provides support to deploy Voice over Internet Protocol (VoIP) telephones into the LAN environment. LLDP-MED supports basic configuration, network policy configuration, location identification, and inventory management.

For more information, see [Administering VOSS](#).

Link-state tracking (LST)

*** Note:**

DEMO FEATURE - Link-state tracking (LST) is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

Link-state tracking (LST) binds the link state of multiple interfaces, creating LST groups with upstream (to-be-followed) and downstream (to-follow) interfaces. LST monitors the state of upstream interfaces and automatically transfers the upstream state to the downstream interfaces. If all the upstream interfaces in a LST group are down, the downstream interfaces are administratively configured as down after approximately five seconds. If any upstream interface in a LST group is up, the downstream interfaces are not affected. The role of the LST group is to keep the downstream interfaces in the same state as the upstream interface.

An interface can be an aggregation of ports, multi-link trunks (MLT) or link aggregation groups (LAG). Interfaces can only belong to one LST group. You can configure LST using CLI or EDM. LST receives updates from Port Manager, MLT, and VLACP regarding the upstream state of ports and trunks in the group.

For more information, see [Configuring Link Aggregation, MLT, SMLT and vIST for VOSS](#).

MLT and LACP Trunk Configuration Enhancement

You can now use `untag-port-default-vlan` with MLT and LACP trunk configurations on VSP 8600 Series.

For more information, see [Configuring VLANs, Spanning Tree, and NLB for VOSS](#).

Multiple CLI Users Per Role

* Note:

DEMO FEATURE - Multiple CLI Users Per Role is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

This release increases the number of CLI users per role (rwa, rw, ro) from 3 users (1 per role) to a maximum of 10 CLI users per switch, which includes:

- 3 default users (rwa, rw, ro)—User Type = default
- 7 user defined users—User Type = userDefined

User defined users can have ro or rw or rwa access rights.

For more information, see [Administering VOSS](#).

NTPv4 Client for IPv4 and IPv6

Network Time Protocol (NTP) is widely used to synchronize time between devices on networks. NTP version 4 (NTPv4) is an extension to the current NTPv3 where it supports IPv6 addresses, and is backward compatible with NTPv3. NTPv4 includes fundamental improvements that extend the potential accuracy to the tens of microseconds. It includes a dynamic server discovery scheme, so that in many cases, specific server configuration is not required.

For more information, see [Administering VOSS](#).

NTPv4 Master and Restrict

* Note:

DEMO FEATURE - NTPv4 Master Mode and Restrict is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information, see [VOSS Feature Support Matrix](#).

The switch can operate as both NTPv4 client and NTPv4 Server. You can configure the NTPv4 Server to operate in master mode. This release also introduces the NTPv4 Restrict capability on the switch that permits NTP traffic (with default restrictions) for all IP addresses or permits NTP traffic from the specified IP addresses or networks.

For more information, see [Administering VOSS](#).

Segmented Management Instance for use with NTPv4

This release introduces support for a Segmented Management Instance.

A Segmented Management Instance is required to provide access to specific management applications. After you create the Management Instance, you can add an IP address to it and configure route redistribution to advertise reachability of the Management Instance to the rest of the network. The first management application to use the Management Instance is NTPv4.

In this release, only the Management Instance CLIP is supported.

For more information, see [Administering VOSS](#).

sFlow Enhancement

The range for the sFlow sampling rate is changed from <8192-1000000> to <1024-1000000>. The change has no impact on upgrades.

For more information, see [Monitoring Performance for VOSS](#).

Policy Based Routing (Redirect Next Hop) per VRF

* Note:

DEMO FEATURE - Policy Based Routing (Redirect Next Hop) per VRF is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

There are two enhancements with the feature:

- Redirect next-hop for VRFs allows users to specify an optional VRF name in addition to the next hop address.
- Redirect next hop action when next hop is unreachable allows an optional “unreachable” action to be configured for redirect next hop ACEs when the specified next hop is unreachable.

For more information, see [Configuring QoS and ACL-Based Traffic Filtering for VOSS](#).

Read-Only User for EDM

* Note:

DEMO FEATURE - Read Only User for EDM is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information, see [VOSS Feature Support Matrix](#).

This release introduces support for read-only (RO) users in addition to the existing RWA user for EDM.

For more information, see [Configuring User Interfaces and Operating Systems for VOSS](#).

Reliability, Availability, Serviceability (RAS) Updates

KHI Enhancements

The output of the `show khi cpp port-statistics` command includes the receive and transmit packet-per-second rate per port, as well as the RxDiff and TxDiff delta fields.

For more information, see [Monitoring Performance for VOSS](#) and [Command Line Interface Commands Reference for VOSS](#).

show fulltech Command Enhancements

The description for the command `show fulltech` has been updated to indicate that the command output includes a recursive listing of filesystem contents.

For more information, see [Command Line Interface Commands Reference for VOSS](#)

Secure AAA Server Communication

* Note:

DEMO FEATURE - Secure AAA server communication is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

This switch software implements the Secure AAA server communication feature. AAA refers to Authentication, Authorization, and Accounting. This feature deploys Internet Protocol Security (IPsec) to provide per-packet confidentiality, authentication, integrity, and replay protection to the AAA server communication, including the security protocols, the Remote Access Dial-in User Services (RADIUS) and Terminal Access Controller Access Control System Plus (TACACS+).

The Internet Key Exchange (IKE) protocol is used for key management.

This feature provides the following enhancements to the security implementation:

- RADIUS secure communication using IPsec for IPv4
- RADIUS secure communication using IPsec for IPv6
- TACACS+ secure communication using IPsec for IPv4
- IPsec support for IPv4 protocol and configuring a Circuitless IP (CLIP) address on a loopback interface.
- Automatic configuration of shared key using IKE protocol for both IPv4 and IPv6.
- IKE support for two types of authentication methods for the IKE session establishment:
 - Pre-shared-key
 - Digital signature (digital certificate signed by trusted Certificate Authority (CA))

For more information, see [Configuring Security for VOSS](#).

SNMP Traps

The traps information is updated to include new features or functions in this release. This release now supports SNMP trap for unsuccessful logon attempts (rcnaSshUnauthorizedAccess) on VSP 8600 Series.

For more information, see [Troubleshooting VOSS](#).

SNMPv3 Remote Engine ID Discovery

Simple Network Management Protocol (SNMP) Inform packets must contain the management (remote) SNMP engine ID. In previous releases, manual configuration of the management SNMP engine ID was required. Remote engine ID discovery provides automatic discovery of manager SNMP engine IDs and removes the need for manual configuration. If the manager SNMP engine ID changes, the discovery process updates the engine table with the new manager SNMP engine ID.

For more information, see [Configuring Security for VOSS](#).

TCP Timestamp Control

TCP Timestamp Control (RFC 1323) provides protection against Wrapped Sequence numbers. However, it is possible to calculate the system uptime when the Timestamp option is enabled. The analysis of timestamp behavior can provide information on the system identity, which poses security threats and can cause a potential attack.

For more information, see: [Configuring Security for VOSS](#) and [Administering VOSS](#).

TLS Client for Secure Syslog

* Note:

DEMO FEATURE - TLS Client for Secure Syslog is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

The syslog server is installed on a host that serves as TLS server. The switch plays the role of a TLS client. A TLS handshake is initiated between the syslog server and the switch.

For more information, see [Troubleshooting VOSS](#).

Two-Factor Authentication - X.509v3 Certificates for SSH

* Note:

DEMO FEATURE - Two-Factor Authentication–X.509v3 Certificates for SSH is a demonstration feature on some products. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. For more information on feature support, see [VOSS Feature Support Matrix](#).

You can use Two-Factor Authentication with smart cards to authenticate users for SSH access to switches for device management. Two-Factor Authentication leverages a Public Key Infrastructure (PKI) security certificate to verify a cardholder's identity prior to allowing access to protected resources. You must enable Secure Shell (SSH) and X.509 V3 authorization on the switch, and provide the digital certificates to enable the identity management for the SSH client and server. Two-Factor Authentication requires: a VSP 8600 Series switch as the SSH server, a PC with Secure CRT 8.3.2 or 8.3.3 as the SSH client, a smart card reader, and Common Access Card (CAC) or Personal Identity Verification (PIV) cards. Optionally you can use a Windows Server 2008 or newer configured with RADIUS server and Active Directory. The switches use SSH and X.509 V3 certificates stored on the smart card. X.509 V3 digital certificates are documented in RFC5280.

For more information, see [Configuring Security for VOSS](#).

VLACP Flap Detect and Damping

Link instability or packet loss can cause the Virtual Link Aggregation Control Protocol (VLACP) state of a link to toggle (flap) rapidly, bringing services (such as IP multicast) up and down in rapid succession. This behavior can cause system-wide instability, including high CPU utilization. VLACP Flap Detect and Damping automatically shuts down selected VLACP links until a network administrator can resolve the root cause of the VLACP flapping. VLACP Flap Detect and Damping does not support auto-recovery. Therefore, a network administrator must re-enable the interface manually.

* Note:

Do not enable VLACP Flap Detect and Damping on Link Aggregation Control Protocol (LACP) enabled ports.

For more information, see [Configuring Link Aggregation, MLT, SMLT and vIST for VOSS](#).

VRRP Enhancement

A VRRP configuration consistency check was added to the `ip vrrp critical-ip-addr <A.B.C.D>` command to notify when a non-local IP address is specified.

Chapter 3: Filenames for this Release

Important:

Do not use Google Chrome or Safari to download software files. Google Chrome can change the file sizes. Safari changes the .tgz extension to .tar.

After you download the software, calculate and verify the md5 checksum. For more information, see [Administering VOSS](#).

The following table provides the filenames and sizes for this release.

Table 3: Software Filenames and Sizes

Description	VSP 8600 Series	File size
SHA512 Checksum files	VOSS8600.8.0.0.0.sha512	1565 bytes
MD5 Checksum files	VOSS8600.8.0.0.0.md5	605 bytes
MIB - supported object names	VOSS8600.8.0.0.0_mib_sup.txt	1,151,484 bytes
MIB - zip file of all MIBs	VOSS8600.8.0.0.0_mib.zip	1,152,056 bytes
MIB - objects in the OID compile order	VOSS8600.8.0.0.0_mib.txt	7,640,359 bytes
EDM Help files	VOSS86v800_HELP_EDM_gzip.zip	4,342,141 bytes
Logs reference	VOSS8600.8.0.0.0_edoc.tar	66,119,680 bytes
Software image	VOSS8600.8.0.0.0.tgz	167,933,807 bytes

The following table provides the open source software filenames and sizes for this release.

Table 4: Open Source Software Files

Master copyright file	Open source base software
VOSS8600.8.0.0.0_oss-notice.html 2,526,400 bytes	VOSS8600.8.0.0.0_OpenSource.zip 95,862,435 bytes

The Open Source license text for the switch is included on the product. You can access it by entering the following command in the CLI:

```
more release/w.x.y.z.GA /release/oss-notice.txt
```

where *w.x.y.z* represents a specific release number.

Chapter 4: Documentation Changes

The Features by Release table has been removed from this document. Product support information for features is now described in product support tables at the beginning of each feature description throughout the documentation suite, and in the [VOSS Feature Support Matrix](#).

Chapter 5: Upgrade Paths and Considerations

This section describes the upgrade path and any considerations that you should be aware of.

Supported Upgrade Paths

Validated upgrade paths are VSP 8600 Series 4.5.x, 6.1, 6.2 or 6.3 to VSP 8600 Series 8.0.

At the time of publishing this document, there were no known restrictions on upgrades. Customers can upgrade directly from other releases to this release.

Upgrade Considerations

The [Administering VOSS](#) document includes detailed image management procedures that includes information about the following specific upgrade considerations:

- Pre-upgrade instructions for IS-IS metric type
- Upgrade considerations regarding MACsec replay-protect configuration
- Upgrade considerations for IS-IS enabled links with HMAC-MD5 authentication
- TACACS+ upgrade consideration

If your configuration includes one of the above scenarios, read the upgrade information in [Administering VOSS](#) before you begin an image upgrade.

 **Note:**

If your switch is configured in High Availability Hot Standby mode with SNMPv3 users using SHA, you must change to Warm Standby mode before upgrading to this release. For more information, see [Known Issues and Restrictions](#) on page 37.

Downgrade Considerations

Before you downgrade to an earlier software release, note the following downgrade considerations.

Real Time Clock

The latest VSP 8600 IOC modules have an updated real time clock (RTC) component, which is not compatible with some older software releases. The new modules should only be installed in a switch or chassis running the minimum supported software, which is 6.2.0.0.

Commissioning New RTC-updated Hardware

To determine if your hardware contains the updated RTC, use the `show sys-info card` command and check the H/W Revision field. If the IOC Module CardHWRevision is 14 or higher, then you have the updated RTC. With the updated RTC, you can only run 6.2.0.0 or higher software versions.

If you attempt to hot insert the latest IOC module (RTC updated) in a chassis running an older unsupported release, the IOC does not become operational. This card attempts to boot unsuccessfully and powers off after 5 boot attempts.

Downgrading New RTC-updated Hardware

If your chassis has any module with the new RTC component, you cannot downgrade the software to a version less than 6.2.0.0. During `software activate` execution, the switch prevents the downgrade and displays the following message:

```
ERROR: Hardware (revision 14) in slot <slot_number> is not supported in
this release. Cannot activate release <x.x.x.x>. Please refer to the
release notes or contact support.
```

If your chassis requires a software downgrade, you must remove all modules with the new RTC component from the chassis first.

*** Note:**

Removing these cards also results in a loss of configuration for the removed slots following a chassis boot.

IS-IS Authentication

If you already have IS-IS Authentication enabled and then downgrade to a previous release, the IS-IS adjacencies may not get established. This issue affects the 100 Gb 8606CQ links only, but it can result in traffic loss.

*** Note:**

This applies only when you downgrade the software from the current release to 6.1.x.

It does not apply when you downgrade the software to 6.2.x.

Use the following procedure as a workaround:

1. Disable IS-IS Authentication on 100 Gb ports on both peers.
2. Downgrade the software to the required release.
3. Re-enable IS-IS Authentication.

MACsec on 100 Gb Devices

When two VSP 8606CQ modules are connected back to back, the MACsec connection works only if the software version on both ends are the same. The modules must be running a supported release starting with Release 6.1.x. If, for example, one end is running 6.1.x and the other end is running a different release, MACsec will not work and traffic will drop.

With this new implementation of MACsec on the 100 Gb 8606CQ module, the MACsec statistics increment the `Unchecked Packets` counter on the receiving link and not the `Accepted` or `Validated` counter. This counter issue happens only when encryption is disabled on both the transmitting and receiving links.

Software Version Mismatch Generates Warning Messages when Installing a New IOC Module

When there is a mismatch between the software running on the switch and the software on the IOC module, the switch updates the IOC module to the version of software running on the switch. During this process you see errors that are similar to the following:

```
IO1 [12/06/17 11:50:43.513:UTC] 0x0024854b 00000000 GlobalRouter SF-APP WARNING
FabDrv_lc::mapLocalPortsToSysports dnxBcm_assignSysPortToModPort failed: unit=0 sysport=0
modId=40000 tmPort=1
IO1 [12/06/17 11:50:43.526:UTC] 0x0024854b 00000000 GlobalRouter SF-APP WARNING
FabDrv_lc::configIngressSideVOQs: UNKNOWN PORT TYPE OF 773 localPort = 1 modId = 6
IO7 [12/06/17 11:50:45.673:UTC] 0x0024854b 00000000 GlobalRouter SF-APP WARNING
FabDrv_lc::createVoqsForPort: dnxBcm_setPacketLengthAdjustForVoq failed: unit=0
voqBaseId=80000512 cos=4 PACKET_LENGTH_ADJUST=0
IO7 [12/06/17 11:50:45.688:UTC] 0x0024854b 00000000 GlobalRouter SF-APP WARNING
map_local_port_to_connectorPort: INVALID LOCAL PORT OF 10000000
```

The messages stop once the update of the IOC module software has completed. This has no impact on the switch operation.

*** Note:**

This issue applies only to a switch running a mix of releases. For example, there is a mismatch if the switch is running release 6.1.x or higher and it has an IOC running release 4.5.x.

Chapter 6: VSP 8600 Series Hardware and Software Compatibility

Part number	Model number	Initial software release	Supported new software release			
			6.1.0.0	6.2.0.0	6.3.0.0	8.0.0.0
EC8602001-E6	VSP 8608	4.5.0.0	Y	Y	Y	Y
EC8602002-E6	VSP 8608 with 3 SF modules and 4 AC PSUs	4.5.0.0	Y	Y	Y	Y
EC8602003-E6	VSP 8608 DC with 3 SF modules and 4 DC PSUs	4.5.0.0	Y	Y	Y	Y
EC8604001-E6	8600SF	4.5.0.0	Y	Y	Y	Y
EC8604002-E6	8624XS	4.5.0.0	Y	Y	Y	Y
EC8604003-E6	8624XT	4.5.0.0	Y	Y	Y	Y
EC8604004-E6	8616QQ	4.5.0.0	Y	Y	Y	Y
EC8604005-E6	8606CQ	4.5.0.1	Y	Y	Y	Y

Chapter 7: Software Scaling

This section lists software scaling capabilities for the VSP 8600 Series.

Layer 2

Table 5: Layer 2 Maximums

LACP aggregators	192 (up to 224 with channelization)
Layer 2 VSNs	4,000
MAC table size	256,000
MAC table size (with Switch Clustering)	128,000
Microsoft NLB cluster IP interfaces	200
MLT groups	192 (up to 224 with channelization)
MSTP instances	64
Port-based VLANs	4,059
Ports per LACP aggregator	8
Ports per MLT group	8
RSTP instances	1
SLPP VLANs	500
Switched UNI I-SIDs per switch (L2 only)	6000
Switched UNI endpoints per interface (same I-SID)	1
Switched UNI endpoints per interface (different I-SIDs)	4000
Transparent Port UNI services per switch (port mapped to I-SID)	192
VLACP interfaces	128

IP Unicast

Table 6: IP Unicast Maximums

BGP+ peers	16
DHCP Relay forwarding entries (IPv4 or IPv6)	512 per VRF/2,048 per switch
ECMP groups/paths per group	1,000/8
IP interfaces (IPv4 or IPv6 or IPv4+IPv6)	4,059*
* NOTE : The maximum limit for IP interfaces is 3,584, if the limit of 512 VRRP interfaces is reached.	
IPv4 ARP table	64,000
IPv4 BGP peers	256
IPv4 CLIP interfaces	64
IPv4 RIP interfaces	200
IPv4 route policies (per VRF/per switch)	500/5,000
IPv4 static ARP entries (per VRF/per switch)	2,000/10,000
IPv4 static routes (per VRF/per switch)	2,000/10,000
IPv4 UDP forwarding entries	1,024
IPv4/IPv6 VRF instances	512*
<p>* Note: The maximum number of VRFs for inter-VRF redistribution is 256.</p>	
IPv6 BGP Peers (GRT and all VRFs combined)	256
IPv6 CLIP interfaces	64
IPv6 Ingress ACEs (Security and QoS)	2,000
IPv6 Neighbor table	16,000
IPv6 OSPFv3 routes (GRT and VRFs combined)	32,000
IPv6 RIPng peers	48
IPv6 RIPng routes	16,000
IPv6 Route Table size	32,000
IPv6 static neighbor records	128 per VRF or 1,000 per switch
IPv6 static routes	10,000
Layer 3 VSNs	512
Manually configured 6-in-4 tunnels	16
OSPF virtual instances	64
OSPF v2/v3 neighbors (GRT and all VRFs combined)	500
OSPFv2 areas	12 per VRF or GRT/80 per switch
OSPFv3 areas (GRT and VRFs combined)	64

Table continues...

OSPFv2/v3 interfaces (GRT and all VRFs, active/passive)	500/2,000
Routed Split Multi-LinkTrunking (RSMLT) interfaces	1,000
VRRP interfaces (IPv4 or IPv6)	512
VRRP interfaces with fast timers (200ms)	24
VRRP VRIDs	8 (combined across IPv4 and IPv6)

DvR

Table 7: DvR Maximums

DvR domains per SPB fabric.	16
Controller nodes per DvR domain with default route inject flag enabled. Total number of Controllers per domain cannot exceed 8. * Note: A DvR domain containing only Controller nodes and no Leaf nodes can have more than 8 Controllers per domain.	8
DvR host routes per DvR domain.	40,000

Layer 3 Route Table Size

Table 8: Layer 3 Route Table Size Maximums

IPv4 BGP routes (control plane only)	1.5 M
IPv4 OSPF routes	64,000
IPv4 RIP routes (per VRF/per switch)	2,000/16,000
IPv4 routes	252,000
IPv4 SPB Shortcut routes	16,000
IPv6 Shortcut routes (Incoming, processed ISIS routes):	32,000
IPv6 Shortcut routes (Outgoing, advertised ISIS routes):	
64-bit long prefix	22,000
128-bit long prefix	15,000
prefix length 64-128 (across all VRFs)	16,000

IP Multicast

Table 9: IP Multicast Maximums

IGMP interfaces	4,000
PIM interfaces (Active/Passive)	512/3,000
Multicast receivers/IGMP receiver entries (per switch)	6,000*
<p>* Note: 6000 is the the total number of unique SGVs for which there are receivers. The total number of receivers can be greater than 6000 if there are multiple receivers for the same group.</p>	
Multicast senders/IGMP sender entries (per switch)	6,000
PIM-SSM static channels	4,000
Total multicast routes (S,G,V) (per switch)	6,000

*** Note:**

IPv4 Routes, IPv4 SGV sender records, IPv6 Routes and IPv6 neighbor records reside in the same shared hardware table. If records of all 4 types are present together in this shared table, then the actual numbers that can be supported might be less than the scaling numbers indicated in the above tables.

Filters, QoS, and Security

Table 10: Filters, QoS, and Security Maximums

Total ACE - Ingress	3,500 (2,000 IPv4 ACEs and 1,500 IPv6 ACEs)
Total ACE - Egress	2,000
Total ACL - Ingress	2,000
Total ACL - Egress	1,000

Fabric Scaling

Table 11: Fabric Scaling Maximums

Number of SPB regions	1
Number of B-VIDs	2

Table continues...

Number of SPB adjacencies	192
SPBM enabled nodes per region (BEB + BCB)	2,000*
* NOTE : If there are VSP 4000 switches in the network, then the total number of SPBM enabled switches per region is reduced to 550.	
SPB multicast nodes per domain	1000
Multicast streams per BEB	6000
Multicast streams per BCB	15,000
Number of VLANs per FA enabled link	94
Number of BEBs this node can share services with (Layer 2 VSNs, Layer 3 VSNs, Multicast)	500**
** NOTE : vIST clusters are counted as 3 nodes.	
Maximum number of SPB Layer 2 multicast UNI I-SIDs	6,000
Maximum number of SPB Layer 3 multicast UNI I-SIDs	6,000
Maximum number of IP multicast S,Gs when operating as a BCB	50,000

OAM and Diagnostics

Table 12: OAM and Diagnostics Maximums

EDM sessions	5
FTP sessions	4
Mirrored destination ports	4
Mirroring ports	191
Rlogin sessions	8
sFlow sampling rate	5000 samples per second per IOC module
SSH sessions	8
Telnet sessions	8

Chapter 8: Important Notices

This section provides important information for this release.

Improved MAC Statistics Diagnostics

Use the command `show io control {slot[-slot] [, ...]}` to capture MAC address-movement statistics on a single slot, for improved diagnostics.

Example:

```
VSP8600:1#show io control
```

CP to SSIO MAC Statistics

Slot	Add	Delete	Refresh	Move
1	4168	24362	265886	0
2	4168	20207	0	0
4	4168	20184	0	0

HW to SSIO MAC statistics

Slot	Delete	Add	Report	Learn	Age	Move	Unsupported	Unknown
1	0	0	0	5724	4178	37	0	
0								
2	0	0	0	3637	2959	3110	0	
0								
4	0	0	0	174	2888	3633	0	0

8624XS IOC Module Power Consideration

1 Gbps or 10 Gbps copper transceivers in any port of the 8624XS IOC module continue to receive power even after you enter the `no sys power slot` command. This causes the remote end to declare the port UP and send traffic.

* Note:

This issue can cause a problem *only* if you use the `no sys power slot` command locally to power down and leave the module in the slot. Although all the ports are initially brought down gracefully as part of the execution of `no sys power slot`, the ports with 1 Gbps or 10 Gbps copper transceivers continue to receive power locally causing the PHY in the transceivers to renegotiate with the remote port. Eventually the port will be declared UP in the remote end.

However, the local end will still stay operationally down. Traffic loss results when the remote switch tries to send traffic to these ports.

To resolve this issue, use one of the following workarounds:

- Shut down the ports (`shutdown port`) in the remote switch before issuing the `no sys power slot` command locally.
- Configure VLACP on the links connected through the copper transceivers above if the far end switch supports VLACP. This provides a logical link down notification at the far end and prevents traffic loss.
- Remove the local IOC module that was powered down.

Feature Licensing

Licensing allows switch operators to select the features that best suits their needs.

The VSP 8600 Series supports a licensing model that has two main categories of licenses: Base License and Feature Pack Licenses. A Base License enables base software features and one is required per IOC in the chassis. You require a Feature Pack License to enable additional features that are grouped into Feature Packs. These licenses are optional.

Licenses are tied to the switch Base MAC address. After you generate the license through Extreme Networks Support Portal at <https://extremeportal.force.com/ExtrLicenseLanding>, you can install the license on the switch.


Offer Level	Period	Support
Factory Default	30-days	Can configure all features, excluding MACsec.
Trial	60 days	<p>Can test licensed features. The following types of Trial Licenses are available:</p> <ul style="list-style-type: none"> • allows the use of all features excluding MACsec • allows the use of all features including MACsec <p> Note: You can activate a Trial License once per switch.</p>
IOC Base License		<p>Can use Base software features on the switch.</p> <p>IOC Base license is required for each IOC module that you plan to install in the chassis. If the number of IOCs exceeds the licensed IOC quantity, the ports on the excess IOCs are license-locked and appear administratively down.</p>

Table continues...

Offer Level	Period	Support
Feature Pack		<p>Features that are not available in the Base License are grouped into Feature Packs based on use case. A license is required to use a Feature Pack. A Feature Pack License applies to the entire chassis; you do not need to purchase this license type for each installed IOC module.</p> <p>Feature Pack licenses that the VSP 8600 supports:</p> <p>Layer 3 Virtualization:</p> <ul style="list-style-type: none"> • Layer 3 Virtual Services Networks (VSNs) • DvR Controller • Greater than 25 VRFs • Greater than 17 BGP Peers <p>Layer 3 Virtualization with MACsec:</p> <ul style="list-style-type: none"> • Layer 3 Virtual Services Networks (VSNs) • DvR Controller • Greater than 25 VRFs • Greater than 17 BGP Peers • MACsec

For more information about licenses, see [Administering VOSS](#).

Limitations on license filename size

When you dynamically load a named license file, ensure that the file name has a maximum of 42 characters *including* the .xml extension. In other words, the length of the file name must be less than or equal to 42 characters, including the extension.

Otherwise, the license file does not load successfully on system reboot.

High Availability (HA)

VSP 8600 supports controller redundancy, thus enabling High Availability (HA). Each IOC module supports both I/O and supervisor/controller functionality. An IOC inserted in Slot 1/2 acts as the Master/Standby Controller in an HA configuration.

VSP 8600 supports two HA modes: Warm Standby and Hot Standby.

- In Warm Standby mode, the configurations are synchronized between Master and Standby IOCs. In Warm Standby mode, if there is a software failure on the Master IOC, the Standby

IOC immediately takes over and reboots all the other IOCs. If Fabric or vIST is provisioned, non-stop forwarding can be achieved by network-based resiliency enabled by these technologies.

- In Hot Standby mode, both configuration and protocol states are synchronized between Master and Standby, thus ensuring a hitless switchover upon Master IOC failure.

! Important:

Hot Standby does not support configurations with SPBM.

Network Load Balancing (NLB)

VSP 8600 supports Network Load Balancing (NLB) in Unicast mode only.

System Name Prompt vs. IS-IS Host Name

Starting with Release 6.1, the software no longer allows spaces in the system name prompt, but it still allows spaces in the IS-IS host name. When you upgrade, the software replaces spaces in the system name with underscores while leaving the IS-IS host name unchanged.

VRRP IDs

Because there is a hardware limitation of using only eight MAC addresses for VRRP, the number of VRIDs is also limited to eight. You can use any eight values for VRIDs between 1 and 255. However, once you choose the eight VRID values, you must reuse the same eight values across all VLANs on the device.

As VRRP virtual MAC for IPv4 and IPv6 for a same VRID is different, IPv4 and IPv6 VRRP instance with same VRID will consume 2 VRRP MAC entries. For example: if VRID 1 is used for IPv4 and IPv6 is used, virtual MAC for IPv4 and IPv6 are 00:00:5e:00:01:01 and 00:00:5e:00:02:01 respectively. These virtual MAC addresses use 2 VRRP MAC addresses in hardware.

Using the syntax for establishing the VRID and Virtual IP Address (`ip vrrp address [VRRP ID] [VRRP Virtual IP Address]`), the following example uses VRIDs from 2 through 9. This example shows only the relevant commands to illustrate this issue.

```
VSP8600:1(config-if)#ip vrrp address 2 2.1.1.10
VSP8600:1(config-if)#ip vrrp address 3 3.1.1.10
VSP8600:1(config-if)#ip vrrp address 4 4.1.1.10
VSP8600:1(config-if)#ip vrrp address 5 5.1.1.10
VSP8600:1(config-if)#ip vrrp address 6 6.1.1.10
VSP8600:1(config-if)#ip vrrp address 7 7.1.1.10
```

Important Notices

```
VSP8600:1(config-if)#ip vrrp address 8 8.1.1.10  
VSP8600:1(config-if)#ip vrrp address 9 9.1.1.10
```

At this point you have used all the VRIDs in the selected range (2–9). Now you must start reusing the VRIDs from 2 to 9 for all other VRRP enabled VLANs. The following example shows what happens when you do not reuse a VRID from the selected range.

```
VSP8600:1(config-if)#ip vrrp address 10 10.1.1.10  
Error: maximum number of VRRP entries exceeded
```

The following example shows the correct reuse of one of the VRIDs from the selected range.

```
VSP8600:1(config-if)#ip vrrp address 2 10.1.1.10
```

Chapter 9: Known Issues and Restrictions

This section details the known issues and restrictions found in this release.

Known Issues and Restrictions

This chapter details the known issues and restrictions found in this release. Where appropriate, use the workarounds provided.

General Issues and Restrictions

Issue number	Description	Workaround
VOSS-4712	When there are broadcast packets in the VLAN, these packets are sent to all ports in the VLAN. The packets get dropped because the port is operationally down. However, outPkts stats increment and the unicast packets are not sent to that port because the port is down.	Ignore the stats counter when port is down.
VOSS-5191	The OSPF MD5 related functionality cannot be enabled from EDM.	Use CLI to configure OSPF MD5 related functionality.
VOSS-5511	Half duplex option is not supported, but it can be configured on VSP 8600 port.	Do not configure half-duplex.
VOSS-5702	Multicast traffic will not have DSCP marked (when enabled on incoming port), when IGMP snooping is enabled on the VLAN.	No workaround.
VOSS-5990	Path MTU discovery feature is not supported for IPv6. Due to this, packets larger than IPv6 interface MTU size are dropped but no ICMP error message is sent to the source host indicating the reason for this drop.	No workaround.
VOSS-6102	<code>sys action reset counters</code> command does not reset ISIS control packets.	Use <code>clear isis</code> command to reset stats.
VOSS-6103	<code>sys action reset counters</code> command does not reset ISIS int-counters.	Use <code>clear isis</code> command to reset stats.
VOSS-6104	<code>sys action reset counters</code> command does not reset any ISIS system stats.	Use <code>clear isis</code> command to reset stats.

Table continues...


Issue number	Description	Workaround
VOSS-7148	EDM: In the Virtual IF tab, the options SHA-1 and SHA-2 are not available to configure virtual link authorization.	Use CLI to configure virtual link authorization.
VOSS-7500	COM+ does not display correct number of IP OSPF ECMP routes.	No workaround. COM+ is no longer supported.
VOSS-7709	On the 8608CQ IOC module, the output of the show interface gigabitEthernet statistics command does not display a value in IN PACKET for packets that have ethertype/length field of 0.	No workaround.
VOSS-7941	<p>When there is a mismatch between the software running on the switch and the software on the IOC module, the switch updates the IOC module to the version of software running on the switch. During this process you see errors that are similar to the following:</p> <pre> IO1 [12/06/17 11:50:43.513:UTC] 0x0024854b 00000000 GlobalRouter SF- APP WARNING FabDrv_lc::mapLocalPortsToSysports dnxBcm_assignSysPortToModPort failed: unit=0 sysport=0 modId=40000 tmPort=1 IO1 [12/06/17 11:50:43.526:UTC] 0x0024854b 00000000 GlobalRouter SF- APP WARNING FabDrv_lc::configIngressSideVOQs: UNKNOWN PORT TYPE OF 773 localPort = 1 modId = 6 IO7 [12/06/17 11:50:45.673:UTC] 0x0024854b 00000000 GlobalRouter SF-APP WARNING FabDrv_lc::createVoqsForPort:dnxBcm_s etPacketLengthAdjustForVoq failed: unit=0 voqBaseId=80000512 cos=4 PACKET_LENGTH_ADJUST=0 IO7 [12/06/17 11:50:45.688:UTC] 0x0024854b 00000000 GlobalRouter SF-APP WARNING map_local_port_to_connectorPort: INVALID LOCAL PORT OF 10000000 </pre> <p>The messages stop when the update of the IOC module software has completed. This has no impact on the switch operation.</p> <p> Note:</p> <p>This issue applies only to a switch running a mix of releases. For example, there is a</p>	No workaround, but there is no operational impact.

Table continues...

Issue number	Description	Workaround
	mismatch if the switch is running release 6.1.x or higher and it has an IOC running release 4.5.x.	
VOSS-8017	SNMPv3 privacy option supports DES and AES128 only. There is no support for higher AES options like AES192, AES256, and AES512.	No workaround.
VOSS-8278	EDM does not have a field to configure the RSA user key.	Use the CLI to configure the RSA user key.
VOSS-8444 VOSS-8758	Disabling IS-IS incorrectly can cause unforeseen problems including traffic loss.	Use the following procedure to disable IS-IS on a switch: <ol style="list-style-type: none"> 1. Shut down NNI ports first. 2. Disable IS-IS globally.
VOSS-8469	For Windows Server Certificate Authorities, the IPsec tunnel cannot use digital certificates as the authentication method.	Use EJBCA as the CA.
VOSS-8516	Secure Copy (SCP) cannot use 2048-bit public DSA keys from Windows.	Use 1024/2048-bit RSA keys or 1024-bit DSA keys.
VOSS-8549	Configuring inter-VRF redistribution on more than 256 VRFs can deplete virtual memory and cause the following warning: VmSize of proc cbcp-main.x(4429) is 1867272KB, above 90% of available 1782579KB(index 0).	Configure inter-VRF redistribution on a maximum of 256 VRFs.
VOSS-8831	When ingress mirroring is configured on an NNI port, two mirrored copies will be made for an incoming mac-in-mac packet that contains a multicast BMAC DA, and also if the ISID carried in the packet is terminated on that fabric connect node.	No workaround.
VOSS-9977	Filter statistics do not increment if the incoming packet is marked for drop AND the filter has an action of mirror. For example: Packets might be marked for drop because the port is not a member of the VLAN specified in the packet. The mirror action does take place (along with other actions, if any, such as internalQos). Filter statistics increment normally if the packet is not marked for drop or if the packet does not contain a mirroring action (even if the packet is marked for drop).	If traffic is getting dropped because the port is not a member of the VLAN then make sure the port is part of the VLAN present in the packet.

Table continues...

Known Issues and Restrictions

Issue number	Description	Workaround
VOSS-9985	<p>If an IGMPv3 interface has both static and dynamic receivers on the same port, the switch clears the static port from the outgoing port list when the dynamic receiver disappears.</p> <p>To avoid this potential traffic loss, avoid having both static and dynamic receivers on an IGMPv3 interface.</p>	No workaround.
VOSS-10091	<p>After deleting an IPVPN, you might see the following error message:</p> <pre>ercdDeleteIpmcRecord:1734 ercdIpmcLookupAvlTree() failed SrcIp: 0x1b000093, DstIp: 0xe6290000 vlan_id 0xffff</pre> <p>This issue has no impact on the switch operation.</p>	No workaround.
VOSS-10362	<p>There is no consistency check to prevent a user from assigning a new I-SID value to a VLAN that already has an I-SID assigned to it. This is currently the existing behavior for I-SID Assignment and users should be aware of this to prevent unintended consequences.</p>	No workaround.
VOSS-10557	<p>SNMP Get tools do not translate the port number to a name.</p>	To get the port name, use the CLI or EDM.
VOSS-10681	<p>After deleting an L3VSN VLAN running IPMC traffic and then recreating it in the GRT (VRF 0), you might see OSAL backtrace messages such as the following:</p> <pre>1 2018-06-26T06:33:45.090-05:00 VSP8608-1(120.169) CP1 - 0x0022c590 - 00000000 GlobalRouter OSAL INFO [bt] Execution path: 1 2018-06-26T06:33:45.090-05:00 VSP8608-1(120.169) CP1 - 0x0022c590 - 00000000 GlobalRouter OSAL INFO [bt] /opt/ appfs/lib/cp/libndutl.so.1(nd_utl_backtrace+0x4c) [0xfc8ff20] 1 2018-06-26T06:33:45.090-05:00 VSP8608-1(120.169) CP1 - 0x0022c590 - 00000000 GlobalRouter OSAL INFO [bt] cbc- main.x(show_stackframe+0x1c) [0x1141c4f0]</pre> <p>This issue has no impact on the switch operation.</p>	No workaround.

Table continues...

Issue number	Description	Workaround
VOSS-10839	The <code>no mvpn enable</code> and <code>no ipvpn</code> commands could cause IS-IS adjacency flapping in setups with a large number of multicast streams and receivers. SPBM traffic cannot pass through the switch until the adjacencies are up again.	Use one of the following workarounds: <ul style="list-style-type: none"> • Increase the IS-IS hold down timer. • Remove the multicast streams or the multicast receivers in that VRF and then execute <code>no mvpn enable</code> or <code>no ipvpn</code>.
VOSS-10852	<p>In an IP Multicast over Fabric Connect scenario with a local SMLT sender and A and B vIST peers, the multicast traffic is hashed to A on VLAN xxx. VLAN xxx is not yet configured on A and B.</p> <ol style="list-style-type: none"> 1. Configure VLAN xxx with <code>ip spbmultipcast enable</code> on A. The sender is created on A and tries to sync to B. However, B ignores the message since VLAN xxx is not yet configured on B. 2. Configure VLAN xxx with <code>ip spbmultipcast enable</code> on B. The local senders on A are not sent to B until the periodic resync that occurs every 15 minutes. During this 15 minutes if an SMLT outage occurs and traffic is hashed to B, there will be minimal traffic outage until B creates the distribution tree on the SPBM core. <p>! Important:</p> <p>Configure VLAN xxx with <code>ip spbmultipcast enable</code> on B. The local senders on A are not sent to B until the periodic resync that occurs every 15 minutes. During this 15 minutes if an SMLT outage occurs and traffic is hashed to B, there will be minimal traffic outage until B creates the distribution tree on the SPBM core.</p>	Use one of the following workarounds: <ul style="list-style-type: none"> • Bounce IP Multicast over Fabric Connect on A's VLAN xxx. • Create VLAN xxx on A and B with no traffic running.
VOSS-11063 VOSS-10628	After deleting and re-creating (or swapping) primary and secondary B-VLANs in a scaled SPBM fabric network with a large number of flows, there might be some unicast and multicast traffic loss on some of the flows.	After deleting and re-creating the B-VLANs, if some of the traffic flows don't recover, then reboot the switch for all the traffic to resume.
VOSS-11414	When IS-IS routes are removed because the next hop is no longer present, you might see COP error messages like the following: <code>COP-SW ERROR ercdProcIpRecMsg: Failed to Delete IP</code>	No workaround, but there is no operational impact.

Table continues...

Known Issues and Restrictions

Issue number	Description	Workaround
	<p>Record. IpAddr:3.0.34.160 IpMask:255.255.255.224 vrfID:9 retStatus: -4</p> <p>This issue has no impact on the switch operation and occurs only when an IS-IS accept policy has been applied.</p>	
VOSS-12399	The system displays continuous LACP/SMLT aggregation transition messages when you configure SMLT on the MLT interface, in a triangular SMLT setup that does not have vIST configured.	In a triangular SMLT setup, if the participant switches are edge switches, do not configure SMLT on the MLT interfaces.
VOSS-12520	An error is displayed on the output console when multicast traffic with source IP address 0.0.0.0 is sent.	Configure an explicit filter rule to drop the packets with source IP address 0.0.0.0.
VOSS-12953 VOSS-13898	<p>The system sometimes displays the following message:</p> <pre>HW INFO Detected Unknown Card Type module in slot</pre>	No workaround, but there is no operational or functional impact.
VOSS-13265	When multiple slots on the switch chassis are powered on at the same time, it causes the cards to power on multiple times.	When using the CLI, power on one card at a time.
VOSS-13789	A link between a channelized 100 Gbps port on the VSP 8600 Series and a 25 Gbps port on the VSP 7400-48Y is not established, when you use a 100 Gbps to 25 Gbps direct-attach copper breakout cable.	To restore the link, you must disable Auto-Negotiation and FEC on the 25 Gbps port on the VSP 7400-48Y.
VOSS-13856	The system sometimes sends traffic packets with the time-to-live (TTL) set to zero.	No workaround, but there is no operational impact or impact to traffic.
VOSS-13921	When you remove a card from the VSP 8600 chassis, you might see the following warning message: <pre>GlobalRouter LMF WARNING lmfUpdateEgressIpmcSession () Null receiver port for subintf null egress vlan 334.</pre>	No workaround, but there is no operational impact or impact to traffic.
VOSS-13924	<p>You might sometimes see the following log messages warning you about SMLT send-queue utilization. These messages appear even if the problem doesn't exist, and can be repetitive.</p> <pre>1 2019-05-30T13:46:03.941-04:00 wolffboro-1 CP1 - 0x00064724 - 00000000 GlobalRouter MLT WARNING SMLT buffer usage over 200M. Low memory warning condition</pre>	No workaround, but there is no functionality impact.

Table continues...

Issue number	Description	Workaround
	<pre> 1 2019-05-30T13:46:03.941-04:00 wolfboro-1 CP1 - 0x00064726 - 00000000 GlobalRouter MLT INFO DBG info for SMLT high mem: istSmltSendBufFullFail 7179242 istSocketWaitingForRestOfMsg 878243 istRxBLearnMacCnt 24456 istTxLearnMacCnt 72888 1 2019-05-30T13:46:46.999-04:00 wolfboro-1 CP1 - 0x00064725 - 00000000 GlobalRouter MLT INFO SMLT buffer usage under 100M, clearing low memory condition </pre>	
VOSS-13954	<p>In a scaled FA environment (with nearly 2000 ELAN I-SIDs on a single homed UNI port), if you disable and enable FA on the chassis in quick succession, there is intermittent loss of traffic on the FA endpoints.</p> <p>This issue is seen only if all endpoints are Layer 2 VSN endpoints and scaled close to the limit of 2000 endpoints on a singled homed UNI port.</p> <p>This issue is not seen if a platform VLAN is associated with an I-SID. This issue is also not seen if you disable FA, wait for the endpoints to be cleaned up, and then enable FA.</p>	<p>In a scaled FA environment, when you bring down a port, ensure that all FA endpoints are cleared using the command <code>show i-sid elan</code>, before you enable the port back up again.</p>
VOSS-14044	<p>When you upload the license file on the switch, the filename length must not exceed 42 characters, including the .xml extension. Otherwise, the file does not load successfully upon system reboot.</p>	<p>Ensure that the length of the license filename is less than or equal to 42 characters, including the .xml extension.</p>
VOSS-14085	<p>When you disable and then enable vIST, the following event synchronization failure log messages are displayed on the hot standby CPU:</p> <pre> Application Sync failed for entity:0x504c5342 representing Module PLSB ,event:5/3 maxNumEvents:6Event sync failed to add MGID <Mgid number> to the mgid map using lanemask 0x1000 as index on standby CP </pre> <p>This failure is observed when you disable and enable vIST and can be observed in other scenarios also. It occurs because the MGID is already updated on the standby CPU when synchronization from the master CPU is received.</p>	<p>No workaround, but there is no functional impact.</p>

Table continues...

Issue number	Description	Workaround
VOSS-14094	<p>When you power on/off a card using the CLI, traffic does not recover on FA-enabled MLTs with member ports on a single I/O slot, because the MLT members are not programmed in the data-path.</p> <p>This issue is seen only when you power on/off the card and is not seen when you reset the card, or physically insert into or remove from the slot.</p> <p>This issue is also not seen on the following:</p> <ul style="list-style-type: none"> • SMLT ports • An MLT that has at least one port on a different slot. • A single non-MLT port. • A port that is not FA-enabled. 	<p>Use the following procedure:</p> <ol style="list-style-type: none"> 1. Disable FA on the MLT first. 2. Power on/off the card on the chassis. Wait for it to finish powering on/off completely. 3. Re-enable FA on the MLT.
VOSS-15017	<p>When you reset the switch chassis, the IO card can sometimes crash and reboot with "out of memory" errors. However, after the reboot, the card operates normally.</p>	<p>No workaround, but there is no operational impact to the card.</p>
VOSS-16056	<p>On DvR Controllers, the output of the command <code>show dvr members</code> can show an incorrect SPB L1 cost.</p> <p>However, there is no functional impact since this value is not used on DvR Controllers. It is used only on DvR Leaf nodes.</p>	<p>On DvR Contollers, use the command <code>show isis spbm unicast-fib</code> to show the correct SPB L1 cost value.</p>
VOSS-18510	<p>Privileged EXEC Authentication does not function on secondary CP in HA warm-standby mode on VSP 8600 Series.</p>	<p>Privileged EXEC password authentication feature is only supported in HA hot-standby mode on VSP 8600 Series.</p>
VOSS-18703	<p>Some DvR host routes might be missed by the <code>clear dvr host-entries</code> command, in a scaled up network with traffic running from all the hosts.</p>	<p>You can bounce IS-IS or the DvR controllers to clear the missed DvR host routes.</p>
VOSS-18881	<p>Flushing MAC address tables multiple times in a very short period of time can cause undefined BCM internal errors on VSP 8600 Series.</p>	<p>Wait at least 10 seconds between MAC flushes on VSP 8600 Series.</p>
VOSS-19465	<p>When in Hot Standby mode with SNMPv3 users configured with SHA, the VSP 8600 Series switch will crash after booting. There is no problem with MD5 SNMP users.</p>	<p>Use Warm Standby instead of Hot Standby mode for configurations with SNMPv3 users confiigured with SHA.</p>
VOSS-26505	<p>The SPBM network does not support multiple multicast streams with the same source address and group address (S,G) learned on a VSP 8600</p>	<p>No workaround.</p>

Table continues...

Issue number	Description	Workaround
	Series source BEB on multiple VLANs in the same VRF.	

Filter Restrictions and Expected Behaviors

This section lists known restrictions and expected behaviors that may first appear to be issues.

The following list describes the expected behavior with filters:

- ACL: The incoming packets must be tagged to hit an entry of port-based ACLs containing a VLAN based qualifier in the ACE.
- ACL: InVlan ACLs can match tagged or untagged traffic, with the port-default VLAN considered if the incoming packet is untagged. However, if an ACE of an InVlan ACL contains the qualifier `vlan-tag-prio`, it can be used to filter only tagged traffic and not the untagged traffic.
- ACL: The outPort ACLs cannot match on the fields that are changed in the packet during forwarding decisions. Hence, the fields (Destination MAC, Source MAC, VLAN ID, etc.), which get modified during Layer 3 routing, cannot be used to match on the new contents of these fields in the outgoing packet.
- ACL: The outPort ACLs cannot match on a destination port that is a member of an MLT. So if port 1/5 is a member of an MLT (static or via LACP), an ACE of an outPort filter with member 1/5 will not be hit.
- ACL: In an outPort ACL, the ACEs containing Layer 3 qualifiers will only be hit for packets that are routed. So the qualifiers such as `src-ip` and `dst-ip` (in the `filter acl ace ip <acl><ace>` command) does not work for Layer 2 switched packets.
- ACL: Each filter member port uses a separate TCAM entry, which impacts the overall ACE scaling number. For example, an inPort filter with 5 members that has one ACE configured uses 10 different TCAM entries (with at least 5 each for the user and default ACEs).
- ACL: For outPort ACLs, the use of the `ethertype` qualifier results in two TCAM entries being used internally instead of one (one each for single tagged and untagged packets). The packets with multiple tags are unsupported as we cannot match on Ethertype field of such packets. If VLAN qualifiers are present in ACE (for example, `vlan-id` or `vlan-tag-prio`), the entry for untagged packets is not created internally. So a single TCAM entry is used that matches the tagged packets alone. This impacts the overall ACE scaling number.
- There can be a single ACE hit for a packet. Port-based ACLs have precedence over VLAN based ACLs. However, the default ACEs have a lower priority than the user ACEs.
 1. User ACE of InPort ACL
 2. User ACE of InVlan ACL
 3. Default ACE of InPort ACL
 4. Default ACE of InVlan ACL

*** Note:**

If a packet matches a user ACE in both an inPort and inVLAN ACL, the inVLAN ACL is ignored.

If a packet matches a user ACE in VLAN-based ACL and the default ACE of an inPort ACL, the user ACE in the inVLAN ACL is hit and the inPort ACL is ignored.

- ACL: The monitor actions (monitor-dst-port or monitor-dst-mlt) are not supported for outPort ACLs. They are only applicable to Ingress ACLs (InPort or InVlan). For flow-based mirroring, you can configure these monitor actions at the ACE level.
- ACE: When an ACE with action count is disabled, the statistics associated with the ACE are reset.
- For ACEs of port-based ACLs, you can configure VLAN qualifiers. Configuring Port qualifiers are not permitted.

For ACEs of VLAN-based ACLs, you can configure port qualifiers. Configuring VLAN qualifiers are not permitted.

Filters and QoS

Note the following VSP 8600 filters:

- VSP 8600 does not support the following qualifiers in the egress direction (outPort). However, ingress support (inVlan/InPort) for these qualifiers are available.
 - `arprequest` and `arpresponse`
 - `ip-frag-flag`
 - `tcp-flags`
- The `ip-options` qualifier is not supported.
- The QoS ACE action `remark-dot1p` on ingress (for port and VLAN ACLs) is not supported.

For more information, see [Configuring QoS and ACL-Based Traffic Filtering for VOSS](#).

Chapter 10: Resolved Issues

This section details the issues that are resolved in this release.

Fixes from previous releases

VSP 8600 Series Release 8.0 incorporates all fixes from prior releases, up to and including VSP 8600 Series Release 6.3.5.

Resolved issues in VSP 8600 Series 8.0

Issue number	Description
VOSS-8110	CLI does not display the SSL certificate information.
VOSS-10802	<p>The <code>show app-telemetry status</code> command incorrectly displays the Collector status as <code>not reachable</code>.</p> <p>There is no functional impact with this issue and it occurs only in configurations where the sFlow collector is reachable through both the management port (MGMT) and GRT. sFlow can reach the collector through MGMT, but Application Telemetry does not route through MGMT. Application Telemetry routes through the GRT.</p>
VOSS-10876	Application Telemetry reachability information is not available from EDM.
VOSS-12573	If a Layer 3 S-UNI port has untagged BPDU enabled and at the same time has a platform VLAN configured on a different I-SID, ARP Learning does not work on that platform VLAN.
VOSS-12679	VLACP flap on SMLT link is causing loop triggering slpp.
VOSS-13029	Default routing not installed on some IOC modules, but is installed on others in the same chassis.
VOSS-14481	100G port on 8606CQ IOC module do not receive packets intermittently.
VOSS-15059	ACL with action deny on ipv6 dropping ipv4 packets.
VOSS-15770	Command <code>vlan mlt <vlan-id> <mlt-id></code> does not show in running config when port(s) in LACP/MLT.
VOSS-15959	ECMP route addition and deletion failures with cyclical error messages in log.
VOSS-16842	mDNS packets (224.0.0.251) being copied to the CPU instead of being flooded to VLAN.
VOSS-16929	IOC Modules do not receive keep-alive messages from master CPU module and become non-operational.
VOSS-18029	Card type is not recognized after chassis power cycle.

Table continues...

Resolved Issues

Issue number	Description
VOSS-18257	When <code>dbg enable</code> and <code>show io 12-tables <slot></code> commands are executed, the session hangs and the ability to get io output is not possible.
VOSS-18555	LED on 10GB module (VSP 8624XS) does not properly indicate the link status after power off/on the module by sys power slot.
VOSS-18614	SNMPwalk stops and timeout if DNS server is configured on switch.
VOSS-18833	Identification field in IPv4 header incorrectly set to zero for IS-IS hello packets.
VOSS-19016	<p>In some situations where the smlt-virtual-bmac and per bmac have the same last byte the VSP 8600 might erroneously terminate packets that are destined for the peer.</p> <p>For example, when the peer bmac is 02:be:b1:00:01:21 and virtual bmac is 02:be:b1:00:ca:21, packets destined for 02:be:b1:00:01:21 are incorrectly terminated.</p>

Appendix A: Related Information

The following section contains information related to the current release.

New MIBs

Table 13:

Object Name	Object OID
rcDvr	1.3.6.1.4.1.2272.1.219
rcDvrGlobal	1.3.6.1.4.1.2272.1.219.1
rcDvrGlobalDomainId	1.3.6.1.4.1.2272.1.219.1.1
rcDvrGlobalRole	1.3.6.1.4.1.2272.1.219.1.2
rcDvrGlobalEnable	1.3.6.1.4.1.2272.1.219.1.3
rcDvrVirtualIstLocalAddr	1.3.6.1.4.1.2272.1.219.1.4
rcDvrVirtualIstLocalMask	1.3.6.1.4.1.2272.1.219.1.5
rcDvrVirtualIstPeerAddr	1.3.6.1.4.1.2272.1.219.1.6
rcDvrVirtualIstClusterId	1.3.6.1.4.1.2272.1.219.1.7
rcDvrGlobalDomainSid	1.3.6.1.4.1.2272.1.219.1.8
rcDvrGlobalBackboneSid	1.3.6.1.4.1.2272.1.219.1.9
rcDvrGlobalGatewayMac	1.3.6.1.4.1.2272.1.219.1.10
rcDvrGlobalInbandMgmtIp	1.3.6.1.4.1.2272.1.219.1.11
rcDvrGlobalInjectDefaultRouteDisable	1.3.6.1.4.1.2272.1.219.1.12
rcDvrGlobalOperState	1.3.6.1.4.1.2272.1.219.1.13
rcDvrGlobalSystemIdAsMac	1.3.6.1.4.1.2272.1.219.1.14
rcDvrGlobalHostEntriesClear	1.3.6.1.4.1.2272.1.219.1.15
rcDvrGlobalHostEntriesClearIpv4	1.3.6.1.4.1.2272.1.219.1.16
rcDvrGlobalHostEntriesClearL2Sid	1.3.6.1.4.1.2272.1.219.1.17
rcDvrGlobalHostEntriesClearL3Sid	1.3.6.1.4.1.2272.1.219.1.18
rcDvrVirtualIstSid	1.3.6.1.4.1.2272.1.219.1.19

Table continues...

Related Information

Object Name	Object OID
rcDvrRouteTable	1.3.6.1.4.1.2272.1.219.2
rcDvrRouteEntry	1.3.6.1.4.1.2272.1.219.2.1
rcDvrRouteDestIpAddrType	1.3.6.1.4.1.2272.1.219.2.1.1
rcDvrRouteDestIpAddr	1.3.6.1.4.1.2272.1.219.2.1.2
rcDvrRouteDestMask	1.3.6.1.4.1.2272.1.219.2.1.3
rcDvrRouteL3Isid	1.3.6.1.4.1.2272.1.219.2.1.4
rcDvrRouteEcmpIndex	1.3.6.1.4.1.2272.1.219.2.1.5
rcDvrRouteNextHopMac	1.3.6.1.4.1.2272.1.219.2.1.6
rcDvrRouteL2Isid	1.3.6.1.4.1.2272.1.219.2.1.7
rcDvrRouteCost	1.3.6.1.4.1.2272.1.219.2.1.8
rcDvrRouteNextHopName	1.3.6.1.4.1.2272.1.219.2.1.9
rcDvrRouteType	1.3.6.1.4.1.2272.1.219.2.1.10
rcDvrRouteVrfId	1.3.6.1.4.1.2272.1.219.2.1.11
rcDvrMembersTable	1.3.6.1.4.1.2272.1.219.3
rcDvrMembersEntry	1.3.6.1.4.1.2272.1.219.3.1
rcDvrMemberMacAddress	1.3.6.1.4.1.2272.1.219.3.1.1
rcDvrMemberSysId	1.3.6.1.4.1.2272.1.219.3.1.2
rcDvrMemberNickName	1.3.6.1.4.1.2272.1.219.3.1.3
rcDvrMemberRole	1.3.6.1.4.1.2272.1.219.3.1.4
rcDvrMemberDomainId	1.3.6.1.4.1.2272.1.219.3.1.5
rcDvrInterfacesTable	1.3.6.1.4.1.2272.1.219.4
rcDvrInterfacesEntry	1.3.6.1.4.1.2272.1.219.4.1
rcDvrInterfaceVlanIpAddrType	1.3.6.1.4.1.2272.1.219.4.1.1
rcDvrInterfaceVlanIpAddr	1.3.6.1.4.1.2272.1.219.4.1.2
rcDvrInterfaceL3Isid	1.3.6.1.4.1.2272.1.219.4.1.3
rcDvrInterfaceL2Isid	1.3.6.1.4.1.2272.1.219.4.1.4
rcDvrInterfaceVlanIpMask	1.3.6.1.4.1.2272.1.219.4.1.5
rcDvrInterfaceVrfId	1.3.6.1.4.1.2272.1.219.4.1.6
rcDvrInterfaceVlanId	1.3.6.1.4.1.2272.1.219.4.1.7
rcDvrInterfaceGwIpAddrType	1.3.6.1.4.1.2272.1.219.4.1.8
rcDvrInterfaceGwIpAddr	1.3.6.1.4.1.2272.1.219.4.1.9
rcDvrInterfaceAdminState	1.3.6.1.4.1.2272.1.219.4.1.10
rcDvrInterfaceSpbmcState	1.3.6.1.4.1.2272.1.219.4.1.11
rcDvrInterfaceIcmpVersion	1.3.6.1.4.1.2272.1.219.4.1.12
rcDvrHostEntriesTable	1.3.6.1.4.1.2272.1.219.5

Table continues...

Object Name	Object OID
rcDvrHostEntriesEntry	1.3.6.1.4.1.2272.1.219.5.1
rcDvrHostEntriesIpAddrType	1.3.6.1.4.1.2272.1.219.5.1.1
rcDvrHostEntriesIpAddr	1.3.6.1.4.1.2272.1.219.5.1.2
rcDvrHostEntriesMask	1.3.6.1.4.1.2272.1.219.5.1.3
rcDvrHostEntriesL3lsid	1.3.6.1.4.1.2272.1.219.5.1.4
rcDvrHostEntriesMacAddr	1.3.6.1.4.1.2272.1.219.5.1.5
rcDvrHostEntriesL2lsid	1.3.6.1.4.1.2272.1.219.5.1.6
rcDvrHostEntriesPort	1.3.6.1.4.1.2272.1.219.5.1.7
rcDvrHostEntriesDomainId	1.3.6.1.4.1.2272.1.219.5.1.8
rcDvrHostEntriesType	1.3.6.1.4.1.2272.1.219.5.1.9
rcDvrHostEntriesNextHopName	1.3.6.1.4.1.2272.1.219.5.1.10
rcDvrHostEntriesNextHopMac	1.3.6.1.4.1.2272.1.219.5.1.11
rcDvrHostEntriesClear	1.3.6.1.4.1.2272.1.219.5.1.12
rcDvrHostEntriesVrflid	1.3.6.1.4.1.2272.1.219.5.1.13
rcDvrL3vsnTable	1.3.6.1.4.1.2272.1.219.6
rcDvrL3vsnEntry	1.3.6.1.4.1.2272.1.219.6.1
rcDvrL3vsnVrflid	1.3.6.1.4.1.2272.1.219.6.1.1
rcDvrL3vsnlsid	1.3.6.1.4.1.2272.1.219.6.1.2
rcDvrL3vsnVrfName	1.3.6.1.4.1.2272.1.219.6.1.3
rcDvrL3vsnInjectDefaultRouteDisable	1.3.6.1.4.1.2272.1.219.6.1.4
rcDvrDatabaseTable	1.3.6.1.4.1.2272.1.219.7
rcDvrDatabaseEntry	1.3.6.1.4.1.2272.1.219.7.1
rcDvrDatabaseDestIpAddrType	1.3.6.1.4.1.2272.1.219.7.1.1
rcDvrDatabaseDestIpAddr	1.3.6.1.4.1.2272.1.219.7.1.2
rcDvrDatabaseDestMask	1.3.6.1.4.1.2272.1.219.7.1.3
rcDvrDatabaseL3lsid	1.3.6.1.4.1.2272.1.219.7.1.4
rcDvrDatabaseEcmpIndex	1.3.6.1.4.1.2272.1.219.7.1.5
rcDvrDatabaseNextHop	1.3.6.1.4.1.2272.1.219.7.1.6
rcDvrDatabaseL2lsid	1.3.6.1.4.1.2272.1.219.7.1.7
rcDvrDatabaseOutgoingInterface	1.3.6.1.4.1.2272.1.219.7.1.8
rcDvrDatabaseSpbCost	1.3.6.1.4.1.2272.1.219.7.1.9
rcDvrDatabasePrefixCost	1.3.6.1.4.1.2272.1.219.7.1.10
rcDvrDatabaseNextHopName	1.3.6.1.4.1.2272.1.219.7.1.11
rcDvrDatabaseAge	1.3.6.1.4.1.2272.1.219.7.1.12
rcDvrDatabaseVrflid	1.3.6.1.4.1.2272.1.219.7.1.13

Table continues...

Object Name	Object OID
rcDvrBackboneEntriesTable	1.3.6.1.4.1.2272.1.219.8
rcDvrBackboneEntriesEntry	1.3.6.1.4.1.2272.1.219.8.1
rcDvrBackboneEntriesIpAddrType	1.3.6.1.4.1.2272.1.219.8.1.1
rcDvrBackboneEntriesIpAddr	1.3.6.1.4.1.2272.1.219.8.1.2
rcDvrBackboneEntriesL3Isid	1.3.6.1.4.1.2272.1.219.8.1.3
cDvrBackboneEntriesDomainId	1.3.6.1.4.1.2272.1.219.8.1.4
rcDvrBackboneEntriesEcmpIndex	1.3.6.1.4.1.2272.1.219.8.1.5
rcDvrBackboneEntriesHostMacAddr	1.3.6.1.4.1.2272.1.219.8.1.6
rcDvrBackboneEntriesL2Isid	1.3.6.1.4.1.2272.1.219.8.1.7
rcDvrBackboneEntriesAdvControllerName	1.3.6.1.4.1.2272.1.219.8.1.8
rcDvrBackboneEntriesAdvController	1.3.6.1.4.1.2272.1.219.8.1.9
rcDvrBackboneEntriesNextHopName	1.3.6.1.4.1.2272.1.219.8.1.10
rcDvrBackboneEntriesNextHopMac	1.3.6.1.4.1.2272.1.219.8.1.11
rcDvrBackboneMembersTable	1.3.6.1.4.1.2272.1.219.9
rcDvrBackboneMembersEntry	1.3.6.1.4.1.2272.1.219.9.1
rcDvrBackboneMemberMacAddress	1.3.6.1.4.1.2272.1.219.9.1.1
rcDvrBackboneMemberSysid	1.3.6.1.4.1.2272.1.219.9.1.2
rcDvrBackboneMemberNickName	1.3.6.1.4.1.2272.1.219.9.1.3
rcDvrBackboneMemberRole	1.3.6.1.4.1.2272.1.219.9.1.4
rcDvrBackboneMemberDomainId	1.3.6.1.4.1.2272.1.219.9.1.5
rc2kBootConfigEnableDvrLeafMode	1.3.6.1.4.1.2272.1.100.5.1.54
rcVlanDvrEnable	1.3.6.1.4.1.2272.1.3.2.1.76
rcVlanDvrGwIpv4Addr	1.3.6.1.4.1.2272.1.3.2.1.77
rclsisGlobalBackboneEnable	1.3.6.1.4.1.2272.1.63.1.22
rclgmpStreamTimeout	1.3.6.1.4.1.2272.1.30.11.7.0
rcVossSystemPrivExecPasswordEnable	1.3.6.1.4.1.2272.1.101.1.1.1.5

Modified MIBs

Object Name	Object OID
rclpRedistributeInterVrfProtocol	1.3.6.1.4.1.2272.1.8.100.22.1.2
rclpRedistributeInterVrfRouteSource	1.3.6.1.4.1.2272.1.8.100.22.1.4
rc2kCardFrontOperStatus	1.3.6.1.4.1.2272.1.100.6.1.5
rc2kCardTemperatureZoneInfo	1.3.6.1.4.1.2272.1.204.1.0.14