ExtremeSwitching

Release Notes for VOSS

Release 8.1.1 (VOSS)
9035868-01 Rev AB
January 2020

© 2017-2020, Extreme Networks, Inc.
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, please 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: www.extremenetworks.com/support/policies/software-licensing

Contents

Chapter 1: About this Document 5
Purpose 5
Conventions 5
Text Conventions 5
Documentation and Training 7
Getting Help 8
Providing Feedback to Us 9
Chapter 2: New in this Release 10
New Hardware 10
New in this Release 11
Endpoint Tracking 11
Fabric Attach and Flex UNI Enhancements 12
Segmented Management Instance for XA1400 Series 15
ExtremeCloud IQ Agent 16
Configuring ExtremeCloud IQ 17
ExtremeCloud IQ CLI Commands 21
Filenames for this Release 26
Chapter 3: Upgrade and Downgrade Considerations 30
Supported Upgrade Paths 31
Upgrading DvR Configurations from Releases 6.0.1.1 and Earlier to 6.0.1.2 and Later 31
Real Time Clock 32
Syslog RFC 5424 and Extreme Management Center Integration 33
Post Upgrade Configuration for Zero Touch Fabric Configuration or Dynamic Nickname Assignment 33
Chapter 4: Hardware and Software Compatibility 37
VSP 4000 Series Hardware 37
VSP 4900 Series Hardware 38
VSP 7200 Series Hardware 39
VSP 7400 Series Hardware 41
VSP 8000 Series Hardware 43
XA1400 Series Hardware 44
Transceivers 44
Power Supply Compatibility 45
Chapter 5: Scaling 49
Layer 2 49
IP Unicast 54
Layer 3 Route Table Size 62
Route Scaling 63
IP Multicast 64

Distributed Virtual Routing (DvR)..... 68
VXLAN Gateway..... 69
Filters, QoS, and Security..... 71
Filter Scaling..... 72
OAM and Diagnostics..... 77
Virtualization Scaling..... 81
Fabric Scaling..... 82
Recommendations..... 89
VRF Scaling..... 90
Chapter 6: Important Notices..... 91
100BASE-FX Support on VSP 4000 Series..... 91
AES-GCM SSH Connection with Open SSH..... 91
Auto Negotiation Settings..... 91
dos-chkdsk..... 91
IKEv2 Digital Certificate Support with Strong Swan..... 92
Feature-Based Licensing..... 92
Subscription Licensing for XA1400 Series..... 93
show vlan remote-mac-table Command Output..... 94
Supported Browsers..... 94
System Name Prompt vs. IS-IS Host Name..... 95
Feature Differences..... 95
VSP 4000 Series Connecting to an ERS 8800 Interoperability Notes..... 95
VSP 4000 Series Notes on Combination Ports..... 96
Chapter 7: Known Issues and Restrictions..... 97
Known Issues..... 97
Restrictions and Expected Behaviors..... 115
Chapter 8: Resolved Issues..... 126
Appendix A: Related Information..... 133
MIB Changes..... 133
Deprecated MIBs..... 133
Modified MIBs..... 133
New MIBs..... 137
Obsolete MIBs..... 151

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: Notice Icons

Icon	Alerts you to...
Important:	A situation that can cause serious inconvenience.
Note:	Important features or instructions.

Table continues…

Icon	Alerts you to...
Tip:	Helpful tips and notices for using the product.
Danger:	Situations that will result in severe bodily injury; up to and including death.
Warning:	Risk of severe personal injury or critical loss of data.
Caution:	Risk of personal injury, system damage, or loss of data.

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	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. If the command syntax is `cfm maintenance-domain maintenance-level <0-7>` , you can enter `cfm maintenance-domain maintenance-level 4`.
Bold text	Bold text indicates the GUI object name you must act upon. Examples: • Click **OK**. • On the **Tools** menu, choose **Options**.
Braces ({ })	Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command. For example, if the command syntax is `ip address {A.B.C.D}`, you must enter the IP address in dotted, decimal notation.
Brackets ([ ])	Brackets ([ ]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command. For example, if the command syntax is `show clock [detail]`, you can enter either `show clock` or `show clock detail`.
Ellipses (…)	An ellipsis (…) indicates that you repeat the last element of the command as needed. For example, if the command syntax is `ethernet/2/1 [ <parameter>`

Table continues…

Convention	Description
	`<value> ]...,` you enter `ethernet/2/1` and as many parameter-value pairs as you need.
Italic Text	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.
`Plain Courier Text`	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. Examples: • `show ip route` • `Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00]`
Separator (>)	A greater than sign (>) shows separation in menu paths. For example, in the Navigation tree, expand the **Configuration** > **Edit** folders.
Vertical Line (|)	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. For example, if the command syntax is `access-policy by-mac action { allow | deny }`, you enter either `access-policy by-mac action allow` or `access-policy by-mac action deny`, but not both.

Documentation and Training

Find Extreme Networks product information at the following locations:

Current Product Documentation
Release Notes
Hardware/software compatibility matrices for Campus and Edge products
Supported transceivers and cables for Data Center products
Other resources, like 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: 1-800-998-2408 (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 and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) 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

Subscribing 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 with your information (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. Click **Submit**.

Providing Feedback to Us

Quality is our first concern at Extreme Networks, and we have 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 especially want to know about:

- Content errors or confusing or conflicting information.
- Ideas for improvements to our documentation so you can find the information you need faster.
- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team, you can do so in two ways:

- Use our short online feedback form at https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@extremenetworks.com.

Please 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 describe what is new in VOSS 8.1.1.

Important:

The following platforms support VOSS 8.1.1:

- VSP 4450 Series
- VSP 4900 Series
- VSP 7200 Series
- VSP 7400 Series
- VSP 8000 Series, which includes VSP 8200 Series and VSP 8400 Series
- XA1400 Series

New Hardware

ExtremeAccess Platform 1400 Series

The XA1400 Series and the associated Fabric Connect VPN (FCVPN) license are now generally available. This release introduces the first phase of integration of ExtremeAccess Platform 1400 Series (XA1400) withExtremeCloud IQ, which supports ExtremeCloud IQ-based monitoring capabilities.

ExtremeAccess Platform 1400 Series (XA1400 Series) are compact general-purpose hardware appliances intended for remote site or branch deployments. Designed for use with Fabric Connect VPN (FCVPN) software, the XA1400 Series enhances the value of an existing Extreme Networks automated campus deployment by expanding the reach of Fabric Connect services to remote sites while delivering a consistent and uniform experience. FCVPN software transparently extends Fabric Connect services, such as L2/L3 VSNs, over third-party provider networks, including MPLS-based WANs or the Internet. In addition to secure Fabric segmentation over the WAN, FCVPN software also supports IPSec for end-to-end traffic encryption.

For more information, see XA1400 Series Switches: Hardware Installation Guide.

Fabric Connect VPN Application Software and Licensing

VOSS on the XA1400 Series platforms is licensed as a Fabric Connect VPN (FCVPN) application software. The primary difference between VOSS software on VSP switches and on XA1400 Series is that the VSP platforms support ASIC-based packet forwarding, and the XA1400 Series supports

software-based packet forwarding on the x86 processor. The XA1400 Series software image supports a subset of VOSS features. FCVPN software is offered as a term-based subscription license. A one-year, three-year, or five-year subscription license is required for each XA1400 Series hardware appliance. The FCVPN application software subscription is available in three service entitlement tiers: ExtremeWorks (EW), PartnerWorks (PW), and ExtremeWorks Premier (EWP). Each includes a right-to-use license, software services, and GTAC support for the term. Two bandwidth tiers of licenses are available. A 100 Mbps connectivity license available for both XA1440 and XA1480 models. A 500 Mbps connectivity license is available for the XA1480 model. The bandwidth license tiers only affect aggregate Wide Area Network (WAN) throughput on the XA1400 Series hardware appliance.

New Transceivers and Components

For more information about optical transceivers and components, see:

Extreme Networks optical transceivers and components	*Extreme Networks Pluggable Transceivers Installation Guide*
Compatibility for Extreme Networks SFP, SFP+, SFP28, QSFP+, and QSFP28 transceiver modules with the VOSS-capable hardware	*Extreme Hardware/Software Compatibility and Recommendation Matrices*

New in this Release

The following sections detail what is new in VOSS 8.1.1.

Endpoint Tracking

Table 3: Endpoint Tracking product support

Feature	Product	Release introduced
For configuration details, see Configuring Fabric Basics and Layer 2 Services for VOSS.		
Endpoint Tracking	VSP 4450 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7200 Series	VOSS 8.1.1
	VSP 7400 Series	VOSS 8.1.1
	VSP 8200 Series	VOSS 8.1.1
	VSP 8400 Series	VOSS 8.1.1
	XA1400 Series	Not Supported

Endpoint Tracking was originally available as a demo feature in VOSS 8.1. It is now generally available and can be used in production environments.

This release adds support for Endpoint Tracking, which provides dynamic assignment of virtual machines (VMs) to IP subnets as they move across servers connected through a Shortest Path Bridging (SPB) cloud.

For more information, see Configuring Fabric Basics and Layer 2 Services for VOSS.

For information about feature support for VOSS products, see VOSS Feature Support Matrix.

Fabric Attach and Flex UNI Enhancements

This release supports the following enhancements for Fabric Attach (FA) and FLEX UNI:

- Dynamically learned Fabric Attach announced VLAN-ISID bindings and manually configured Switched-UNI end-points are supported on the same interface.
- FA and Flex UNI are now supported on the same port, including static MLT, and LACP LAGs at the same time.

A consistency check is in place to ensure that you cannot enable FA and Endpoint Tracking simultaneously on the same interface.

The following table summarizes the combinations of features supported on the same interface.

Feature	Supported feature combinations	
FA	Enable	Not allowed
FLEX UNI	Enable	Enable
Endpoint Tracking	Not allowed	Enable

Note:

You must use Command Line Interface (CLI) or Enterprise Device Manager (EDM) to configure the FA and Flex UNI enhancements. Extreme Management Center support will be available in a future release.

Displaying Switched UNI (ELAN) I-SID Information

Use this procedure to display information on FA-created Switched UNI (ELAN) I-SIDs.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Display all Switched UNI (ELAN) I-SIDs:

   ```
   show i-sid elan
   ```

3. Display ELAN I-SID information on an MLT:

   ```
   show mlt i-sid [<1-512>]
   ```

Note:

Viewing ELAN I-SID information on an MLT is useful to understand the origin of the I-SID when multiple client or proxy devices connecting to the FA Server using SMLT MLT advertise the *same* I-SID-to-VLAN mappings. In the event of a link failure on an MLT, the origin of the I-SID helps determine on which MLT, and thereby from which proxy or client device, the mappings were successfully learnt.

4. Display ELAN I-SID information on ports:

   `show interfaces gigabitEthernet i-sid [{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}]`

Example

Display information on all Switched UNI (ELAN) I-SIDs.

The following sample output displays, for example, the I-SID information on one of the peer switches of the FA Server, in a dual-homed SMLT configuration.

```
Switch:1>en
Switch:1#show i-sid elan

================================================================================================
                                        Isid Info
================================================================================================
ISID        ISID                    PORT               MLT                ORIGIN
ID          TYPE       VLANID       INTERFACES         INTERFACES
------------------------------------------------------------------------------------------------
2002        ELAN       N/A          c2002:1/10         -                  DISC_LOCAL
4000        ELAN       N/A          -                  c4000:1            DISC_BOTH
4001        ELAN       N/A          -                  c4001:1            DISC_LOCAL
4030        ELAN       N/A          -                  c4030:1            DISC_BOTH
4051        ELAN       N/A          -                  c4051:1            DISC_BOTH
10200       ELAN       N/A          -                  c200:1             DISC_REMOTE

c: customer vid    u: untagged-traffic

All 6 out of 6 Total Num of Elan i-sids displayed
```

Note:

The `I-SID TYPE` field displays once for each I-SID. The `I-SID TYPE` of an I-SID that is either learned through FA mapping assignments or configured as an FA management I-SID, is always ELAN. If a platform VLAN has the same I-SID value as that of the I-SID in an FA mapping assignment or in an FA management I-SID configuration, then the platform VLAN is associated with the I-SID endpoint and appears in the `VLANID` column.

Note:

- The `ORIGIN` field displays once for each I-SID. It indicates the origin of the I-SID and *not* the origin of the I-SID endpoint. To view the origin of the I-SID endpoints, execute either the `show mlt i-sid` or the `show interfaces gigabitEthernet i-sid` command.
  - The origin of I-SID `4000` displays as `DISC_BOTH`because it is discovered on both vIST peers.

  - The origin of I-SID 4001 displays as DISC_LOCAL because it is first discovered on the local FA Server switch.
  - The origin of I-SID 10200 displays as DISC_REMOTE because it is first discovered on the peer switch and then synchronized with the local switch.
- If the origin of an I-SID is DISC_LOCAL, DISC_REMOTE, DISC_BOTH or MANAGEMENT, it changes to CONFIG, after you manually configure an endpoint on the I-SID.

Display MLT I-SID information for MLT 1.

In this sample output, the ORIGIN field indicates the origin of the I-SID endpoint.

```
Switch:1#show mlt i-sid 1

================================================================================
                                MLT Isid Info
================================================================================
               ISID                    ISID
MLTID IFINDEX ID       VLANID C-VID  TYPE     ORIGIN           BPDU
--------------------------------------------------------------------------------
1     6144    4000     N/A    4000   ELAN     DISC_BOTH
1     6144    4001     N/A    4001   ELAN     DISC_LOCAL
1     6144    4030     N/A    4030   ELAN     DISC_BOTH
1     6144    4051     N/A    4051   ELAN     DISC_BOTH
1     6144    10200    N/A    200    ELAN     DISC_REMOTE
--------------------------------------------------------------------------------

5 out of 6 Total Num of i-sid endpoints displayed
```

Display I-SID information on the port 1/10:

In this sample output, the ORIGIN field indicates the origin of the I-SID endpoint.

Note:

The origin of I-SID 2002 displays as DISC_LOCAL because it is first discovered on the local FA Server switch. The FA Server switch synchronzies I-SID information only to the peer switch in the case of learning on a port interface.

```
Switch:1#show interfaces gigabitEthernet i-sid 1/10

================================================================================
                                PORT Isid Info
================================================================================
              ISID                    ISID
PORTNUM IFINDEX ID       VLANID C-VID  TYPE     ORIGIN           BPDU
--------------------------------------------------------------------------------
1/10    201     2002     N/A    601    ELAN     DISC_LOCAL
--------------------------------------------------------------------------------

1 out of 6 Total Num of i-sid endpoints displayed
```

Variable Definitions

Use the data in the following table to use the **show i-sid** command.

Variable	Value
elan	Displays all ELAN I-SIDs.

Use the data in the following table to use the **`show mlt i-sid`** command.

Variable	Value
<1–512>	The valid range for MLT ID.

Use the data in the following table to use the **`show interfaces gigabitEthernet i-sid`** command.

Variable	Value
{slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

Segmented Management Instance for XA1400 Series

Table 4: Segmented Management Instance product support

Feature	Product	Release introduced
For configuration details, see Administering VOSS.		
Segmented Management Instance - Management Interface CLIP	VSP 4450 Series	VOSS 7.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 7.0
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 7.0
	VSP 8400 Series	VOSS 7.0
	XA1400 Series	VOSS 8.1.1 - IPv4 only
Segmented Management Instance - Management Interface VLAN	VSP 4450 Series	VOSS 7.0
	VSP 4900 Series	VOSS 8.1
	VSP 7200 Series	VOSS 7.0
	VSP 7400 Series	VOSS 8.0
	VSP 8200 Series	VOSS 7.0
	VSP 8400 Series	VOSS 7.0
	XA1400 Series	VOSS 8.1.1 - IPv4 only
Segmented Management Instance — ability to migrate VLAN or loopback IP address	VSP 4450 Series	Not Supported
	VSP 4900 Series	Not Supported
	VSP 7200 Series	Not Supported

Table continues…

Feature	Product	Release introduced
	VSP 7400 Series	Not Supported
	VSP 8200 Series	Not Supported
	VSP 8400 Series	Not Supported
	XA1400 Series	Not Supported

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.

Note:

In this release, only ping and traceroute use a Segmented Management Instance. For information about configuring ping and traceroute applications, see Troubleshooting VOSS.

All other management applications, such as Telnet, SSH, and RADIUS are supported on regular VLAN, Circuitless IP, and out of band (OOB) interfaces and will use a Segmented Management Instance in a future release.

The Management Instance supports Management Instance CLIP and Management Instance VLAN interface types. When you create the Management Instance, you specify the interface type and the switch automatically creates the appropriate instance ID for that type. Each Management Instance supports a single IPv4 management address for use by management applications.

Note:

In this release, Segmented Management Instance for XA1400 Series supports IPv4 addresses only. For information about feature support for VOSS products, see VOSS Feature Support Matrix.

For information about configuring Segmented Management Instance using CLI and EDM, see Administering VOSS.

For information about viewing Segmented Management Instance statistics, see Monitoring Performance for VOSS.

ExtremeCloud IQ Agent

Table 5: ExtremeCloud IQ Agent product support

Feature	Product	Release introduced
ExtremeCloud IQ Agent	VSP 4450 Series	Not supported
	VSP 4900 Series	VOSS 8.1.1 - demonstration feature
	VSP 7200 Series	Not supported

Table continues…

Feature	Product	Release introduced
	VSP 7400 Series	VOSS 8.1.1 - demonstration feature
	VSP 8200 Series	Not supported
	VSP 8400 Series	Not supported
	XA1400 Series	VOSS 8.1.1 - demonstration feature

Note:

A Segmented Management Instance is required to use ExtremeCloud IQ Agent on ExtremeAccess Platform 1400 Series (XA 1400), VSP 4900 Series, and VSP 7400 Series.

ExtremeCloud IQ provides cloud-managed networking, and delivers unified, full-stack management of wireless access points, switches, and routers. It enables onboarding, configuration, monitoring, troubleshooting, reporting, and more. Using innovative machine learning and artificial intelligence technologies, ExtremeCloud IQ analyzes and interprets millions of network and user data points, from the network edge to the data center, to power actionable business and IT insights, and to deliver new levels of network automation and intelligence.

This release is the first phase of integration with ExtremeCloud IQ. When ExtremeCloud IQ Agent is enabled on supported VOSS devices, the devices can register with and send monitoring and device data, such as port statistics and device health, to ExtremeCloud IQ.

Note:

DEMO FEATURE - VOSS integration with ExtremeCloud IQ is a demonstration feature. Demonstration features are provided for testing purposes. Demonstration features are for lab use only and are not for use in a production environment. Please work with your Extreme Networks sales team to request a demonstration.

The following VOSS devices are supported:

- VSP 7400 Series
- XA1400 Series
- VSP 4900-48P
  - VIM5-4YE
  - VIM5-4X
  - VIM5-4XE
  - VIM5-2Y
  - VIM5-2Q

Configuring ExtremeCloud IQ

After your VOSS device is onboarded (that is, the serial number for the device is associated with your ExtremeCloud IQ account), you are only required to enable the ExtremeCloud IQ Agent. Other

feature configuration, such as configuring proxy parameters and configuring access to ExtremeCloud IQ Manager is optional.

Note:
A Segmented Management Instance is required to use ExtremeCloud IQ Agent on ExtremeAccess Platform 1400 Series (XA 1400), VSP 4900 Series, and VSP 7400 Series.

Enable ExtremeCloud IQ Agent

Use this task to enable ExtremeCloud IQ Agent to communicate with the ExtremeCloud IQ Manager.

Before you begin

You must first onboard the VOSS devices and configure any optional ExtremeCloud IQ Agent parameters before you enable ExtremeCloud IQ Agent.

Procedure

1. Enter Application Configuration mode:

   ```
   enable

   configure terminal

   application
   ```

2. Enable ExtremeCloud IQ Agent:

   ```
   cloud-iq enable
   ```

Example

Enable ExtremeCloud IQ Agent:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#cloud-iq enable
```

Configure ExtremeCloud IQ Manager Proxy Parameters

If you use a proxy https server in your network, you must configure proxy parameters so that the ExtremeCloud IQ Agent on the device can communicate with ExtremeCloud IQ Manager through the proxy.

Use this task to configure the proxy parameters for ExtremeCloud IQ Manager on the ExtremeCloud IQ Agent.

Note:

You must onboard the VOSS devices and configure any optional ExtremeCloud IQ Agent parameters on the supported VOSS devices before you enable ExtremeCloud IQ Agent.

Note:

You can configure the ExtremeCloud IQ Agent parameters on the supported VOSS devices first and then onboard the devices (that is, add the serial numbers for the devices in the ExtremeCloud IQ Manager GUI) or vice versa.

Procedure

1. Enter Application Configuration mode:

   ```
   enable
   configure terminal
   application
   ```

2. Configure the ExtremeCloud IQ Manager proxy IPv4 address or DNS name:

   ```
   cloud-iq proxy address <WORD 1-255> tcp-port <1-49151>
   ```

3. Configure the proxy username and password for the ExtremeCloud IQ account:

   ```
   cloud-iq proxy username <WORD 1-64> password <WORD 1-128>
   ```

Example

Configure ExtremeCloud IQ Manager proxy parameters using IPv4 address:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#cloud-iq proxy address 192.0.2.254 tcp-port 21
Switch:1(config-app)#cloud-iq proxy username admin password ****
```

Configure ExtremeCloud IQ Manager proxy parameters using DNS name:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#cloud-iq proxy address hac.extremecouldiq.com tcp-port 21
Switch:1(config-app)#cloud-iq proxy username admin password ****
```

Variable Descriptions

The following table defines parameters for the **`cloud-iq proxy`** command.

Variable	Value
address <*WORD 1-255*>	Specifies the ExtremeCloud IQ Manager proxy IPv4 address or DNS name.
tcp-port <*1-49151*>	Specifies the TCP port for the ExtremeCloud IQ Manager.
username <*WORD 1-64*>	Specifies the ExtremeCloud IQ Manager proxy server username.
password <*WORD 1-128*>	Specifies the ExtremeCloud IQ proxy server password.

Configure Access to ExtremeCloud IQ Manager

Use this task to configure ExtremeCloud IQ Agent parameters to access ExtremeCloud IQ Manager.

Note:

You must onboard the VOSS devices and configure any optional ExtremeCloud IQ Agent parameters on the supported VOSS devices before you enable ExtremeCloud IQ Agent.

Note:

You can configure the ExtremeCloud IQ Agent parameters on the supported VOSS devices first and then onboard the devices (that is, add the serial numbers for the devices in the ExtremeCloud IQ Manager GUI) or vice versa.

Procedure

1. Enter Application Configuration mode:

   ```
   enable
   configure terminal
   application
   ```

2. Configure the ExtremeCloud IQ Manager IPv4 address or DNS name:

   ```
   cloud-iq server address WORD<1-255>
   ```

Example

Configure access to ExtremeCloud IQ Manager using IPv4 address:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#cloud-iq server address 192.0.2.1
```

Configure access to ExtremeCloud IQ Manager using DNS name:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#application
Switch:1(config-app)#cloud-iq server address extremecloudiq.com
```

Variable Descriptions

The following table defines parameters for the **`cloud-iq server`** command.

Variable	Value
address <*WORD 1-255*>	Specifies the ExtremeCloud IQ server IPv4 address or DNS name.

Display ExtremeCloud IQ Information

About this task

Use this task to display ExtremeCloud IQ configuration information and status.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display ExtremeCloud IQ configuration and status:

   ```
   show application cloud-iq
   ```

Display ExtremeCloud IQ Status

About this task

Use this task to display ExtremeCloud IQ status information.

Procedure

1. To enter User EXEC mode, log on to the switch.
2. Display ExtremeCloud IQ status information:

   ```
   show application cloud-iq status
   ```

ExtremeCloud IQ CLI Commands

This section describes the Command Line Interface (CLI) commands for the ExtremeCloud IQ feature configuration.

cloud-iq enable

Enables ExtremeCloud IQ Agent.

Syntax

- **`cloud-iq enable`**
- **`no cloud-iq enable`**

Default

Disabled

Command Mode

Application Configuration

Usage Guidelines

VOSS integration with ExtremeCloud IQ is a demonstration feature on some products. For more information about feature support, see VOSS Feature Support Matrix.

cloud-iq server

Configures the ExtremeCloud IQ Manager IPv4 address or DNS name.

Syntax

- **`cloud-iq server address WORD<1-255>`**
- **`default cloud-iq server`**

- **no cloud-iq server**

Default

The default server address is hac.extremecloudiq.com.

Command Mode

Application Configuration

Usage Guidelines

VOSS integration with ExtremeCloud IQ is a demonstration feature on some products. For more information about feature support, see VOSS Feature Support Matrix.

cloud-iq proxy

Configures the ExtremeCloud IQ Manager proxy parameters.

Syntax

- **cloud-iq proxy address WORD<1-255> tcp-port <1-49151>**
- **cloud-iq proxy username WORD<1-64> password WORD <1-128>**
- **no proxy address WORD<1-255> tcp-port <1-49151>**
- **no proxy username**

Command Paramenters

address WORD<1-255>	Specifies the ExtremeCloud IQ Manager proxy IPv4 address or DNS name.
username <WORD 1-64>	Specifies the ExtremeCloud IQ Manager proxy server username.
password <WORD 1-128>	Specifies the ExtremeCloud IQ proxy server password.
tcp-port <1-49151>	Specifies the TCP port to use for the ExtremeCloud IQ Manager.

Default

None

Command Mode

Application Configuration

Usage Guidelines

VOSS integration with ExtremeCloud IQ is a demonstration feature on some products. For more information about feature support, see VOSS Feature Support Matrix.

cloud-iq notification

Enable the trap notification.

Syntax

- **cloud-iq notification**

- **no cloud-iq notification**

Default

Disabled

Command Mode

Application Configuration

Usage Guidelines

VOSS integration with ExtremeCloud IQ is a demonstration feature on some products. For more information about feature support, see VOSS Feature Support Matrix.

This command will be operational in a future release.

show application cloud-iq

Show ExtremeCloud IQ configuration information and status.

Syntax

- **show application cloud-iq**

Default

None

Command Mode

User EXEC

Command Output

The **show application cloud-iq** command displays the following information:

Table 6:

Output	Description
Agent Admin State	Specifies the administrative state of the ExtremeCloud IQ feature.
Agent Version	Specifies the ExtremeCloud IQ Agent verison running on the device.
Agent Oper State	Specifies the operational status of the ExtremeCloud IQ feature.
Proxy Address	Specifies the proxy address of the ExtremeCloud IQ Manager.
Proxy TCP Port	Specifies the Proxy TCP Port for the ExtremeCloud IQ Manager.
Proxy Username	Specifies the ExtremeCloud IQ Manager proxy server username.
Proxy Password	Specifies the ExtremeCloud IQ Manager proxy server password.

Table continues…

Output	Description
Notification	Enables traps to be sent from the switch to the ExtremeCloud IQ Manager. **Note:** This functionality is not operational in this release.

Example

```
Switch(config-app)#show application cloud-iq

================================================================================
                                Cloud IQ Info
================================================================================
Agent Admin State       : true
Agent Version           : 0.1.0.0
Agent Oper State        : connected
Server Address          : hac.extremecloudiq.com
Proxy Address           : extremeiq.com
Proxy TCP Port          : 21
Proxy Username          : admin
Notification Enable     : false
```

Usage Guidelines

VOSS integration with ExtremeCloud IQ is a demonstration feature on some products. For more information about feature support, see VOSS Feature Support Matrix.

show application cloud-iq status

Show ExtremeCloud IQ status information.

Syntax

- **show application cloud-iq status**

Default

None

Command Mode

User EXEC

Command Output

Output	Description
Connection Status	Specifies that ExtremeCloud IQ Manager and ExtremeCloud IQ Agent are connected.
Last Onboard Time	Specifies the Last Onboard Time for the ExtremeCloud IQ feature.
Agent Version	Specifies the the ExtremeCloud IQ Agent verison running on the device.

Table continues…

Output	Description
Association URL	Specifies the Association URL of the ExtremeCloud IQ feature.
Poll URL	Specifies the Poll URL of the ExtremeCloud IQ feature.
Monitor Frequency	Specifies the Monitor Frequency of the ExtremeCloud IQ feature.
Poll Frequency	Specifies the Poll Frequency of the ExtremeCloud IQ feature.
Last Poll Status	Specifies the Last Poll Status of the ExtremeCloud IQ feature.
Last Poll Success Time	Specifies the Last Poll Success Time of the ExtremeCloud IQ feature.
Last Health Status	Specifies the Last Health Status of the ExtremeCloud IQ feature.
Last Health Success Time	Specifies the Last Health Success Time of the ExtremeCloud IQ feature.
Last Monitor Status	Specifies the Last Monitor Status of the ExtremeCloud IQ feature.
Last Monitor Success Time	Specifies the Last Monitor Success Time of the ExtremeCloud IQ feature.

Example

```
Switch(config-app)#show application cloud-iq status
================================================================================
                                Cloud IQ Status
================================================================================
Connection Status               : Connected
Last Onboard Time               : 18:54:23 11 27 2019 UTC
Agent Version                   : 0.1.0.0
Association URL                 : https://10.16.231.98/hac-webapp/rest/v1/association
Poll URL                        : https://10.16.231.98/hac-webapp/rest/v1/poll/1904Q-20028
Monitor Frequency               : 600
Poll Frequency                  : 30
Last Poll Status                : SUCCESS
Last Poll Success Time          : 14:39:16 11 28 2019 UTC
Last Health Status              : SUCCESS
Last Health Success Time        : 14:38:35 11 28 2019 UTC
Last Monitor Status             : SUCCESS
Last Monitor Success Time       : 14:38:35 11 28 2019 UTC
```

Usage Guidelines

VOSS integration with ExtremeCloud IQ is a demonstration feature on some products. For more information about feature support, see VOSS Feature Support Matrix.

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.

In VOSS 4.2 and later, the encryption modules are included as part of the standard runtime software image file.

Prior to VOSS 4.2.1, image filenames began with VSP, such as, `VSP4K4.1.0.0.tgz`. In VOSS 4.2.1 and later, image filenames start with VOSS, such as, `VOSS8K4.2.1.0.tgz`.

Prior to VOSS 8.1, software image filenames contained either a product family, or a product platform, depending on the product. In VOSS 8.1 and later, all software image filenames contain a product platform, to more accurately and consistently describe the switches that the software applies to.

In VOSS 8.1 and later, when extracting the software image file, the extraction process appends the software version portion of the extracted filenames to include the final full software version. (For example, extracting **`VOSS8400.8.1.0.0.tgz`** results in a software file named **`VOSS8400.8.1.0.0.GA`**.) Ensure that you specify the final full software version (in this case, **`8.1.0.0.GA`**) when using CLI commands that include the software version, such as activating or removing the software.

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.

The following tables provide the filenames and sizes for this release.

Table 7: VSP 4450 Series Software Filenames and Sizes

Description	File	Size
SHA512 Checksum files	VOSS4400.8.1.1.0.sha512	1,549 bytes
MD5 Checksum files	VOSS4400.8.1.1.0.md5	589 bytes
MIB - supported object names	VOSS4400.8.1.1.0_mib_sup.txt	1,363,964 bytes
MIB - zip file of all MIBs	VOSS4400.8.1.1.0_mib.zip	1,155,202 bytes
MIB - objects in the OID compile order	VOSS4400.8.1.1.0_mib.txt	7,662,834 bytes
EDM Help files	VOSSv811_HELP_EDM_gzip.zip	4,328,151 bytes
Logs reference	VOSS4400.8.1.1.0_edoc.tar	66,027,520 bytes

Table continues…

Description	File	Size
Software image	VOSS4400.8.1.1.0.tgz	110,140,699 bytes
Open source software - Master copyright file	VOSS4400.8.1.1.0_oss-notice.html	2,766,416 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 8: VSP 4900 Series Software Filenames and Sizes

Description	File	Size
SHA512 Checksum files	VOSS4900.8.1.1.0.sha512	1,549 bytes
MD5 Checksum files	VOSS4900.8.1.1.0.md5	589 bytes
MIB - supported object names	VOSS4900.8.1.1.0_mib_sup.txt	1,379,452 bytes
MIB - zip file of all MIBs	VOSS4900.8.1.1.0_mib.zip	1,155,202 ytes
MIB - objects in the OID compile order	VOSS4900.8.1.1.0_mib.txt	7,662,834 bytes
EDM Help files	VOSSv811_HELP_EDM_gzip.zip	4,328,151 bytes
Logs reference	VOSS4900.8.1.1.0_edoc.tar	66,027,520 bytes
Software image	VOSS4900.8.1.1.0.tgz	235,909,954 bytes
Open source software - Master copyright file	VOSS4900.8.1.1.0_oss-notice.html	2,766,416 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 9: VSP 7200 Series Software Filenames and Sizes

Description	File	Size
SHA512 Checksum files	VOSS7200.8.1.1.0.sha512	1,549 bytes
MD5 Checksum files	VOSS7200.8.1.1.0.md5	589 bytes
MIB - supported object names	VOSS7200.8.1.1.0_mib_sup.txt	1,372,043 bytes
MIB - zip file of all MIBs	VOSS7200.8.1.1.0_mib.zip	1,155,202 bytes
MIB - objects in the OID compile order	VOSS7200.8.1.1.0_mib.txt	7,662,834 bytes
EDM Help files	VOSSv811_HELP_EDM_gzip.zip	4,328,151 bytes
Logs reference	VOSS7200.8.1.1.0_edoc.tar	66,027,520 bytes
Software image	VOSS7200.8.1.1.0.tgz	124,387,942 bytes
Open source software - Master copyright file	VOSS7200.8.1.1.0_oss-notice.html	2,766,416 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 10: VSP 7400 Series Software Filenames and Sizes

Description	File	Size
SHA512 Checksum files	VOSS7400.8.1.1.0.sha512	1,856 bytes
MD5 Checksum files	VOSS7400.8.1.1.0.md5	704 bytes
MIB - supported object names	VOSS7400.8.1.1.0_mib_sup.txt	1,379,964 bytes
MIB - zip file of all MIBs	VOSS7400.8.1.1.0_mib.zip	1,155,202 bytes
MIB - objects in the OID compile order	VOSS7400.8.1.1.0_mib.txt	7,662,834 bytes
EDM Help files	VOSSv811_HELP_EDM_gzip.zip	4,328,151 bytes
Logs reference	VOSS7400.8.1.1.0_edoc.tar	66,027,520 bytes
Software image	VOSS7400.8.1.1.0.tgz	235,875,191 bytes
Open source software - Master copyright file	VOSS7400.8.1.1.0_oss-notice.html	2,766,416 bytes
YANG model	restconf_yang.tgz	506,020 bytes
Third Party Virtual Machine (TPVM)	TPVM_7400_8.1.1.0.img	1,677,066,240 bytes
Purview Engine Virtual Appliance	purview_7400_8.1.1.0.ova	1,778,386,432 bytes

Table 11: VSP 8200 Series Software Filenames and Sizes

Description	File	Size
SHA512 Checksum files	VOSS8200.8.1.1.0.sha512	1,549 bytes
MD5 Checksum files	VOSS8200.8.1.1.0.md5	589 bytes
MIB - supported object names	VOSS8200.8.1.1.0_mib_sup.txt	1,372,043 bytes
MIB - zip file of all MIBs	VOSS8200.8.1.1.0_mib.zip	1,155,202 bytes
MIB - objects in the OID compile order	VOSS8200.8.1.1.0_mib.txt	7,662,834 bytes
EDM Help files	VOSSv811_HELP_EDM_gzip.zip	4,328,151 bytes
Logs reference	VOSS8200.8.1.1.0_edoc.tar	66,027,520 bytes
Software image	VOSS8200.8.1.1.0.tgz	124,390,394 bytes
Open source software - Master copyright file	VOSS8200.8.1.1.0_oss-notice.html	2,766,416 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 12: VSP 8400 Series Software Filenames and Sizes

Description	File	Size
SHA512 Checksum files	VOSS8400.8.1.1.0.sha512	1,549 bytes
MD5 Checksum files	VOSS8400.8.1.1.0.md5	589 bytes

Table continues…

Description	File	Size
MIB - supported object names	VOSS8400.8.1.1.0_mib_sup.txt	1,372,043 bytes
MIB - zip file of all MIBs	VOSS8400.8.1.1.0_mib.zip	1,155,202 bytes
MIB - objects in the OID compile order	VOSS8400.8.1.1.0_mib.txt	7,662,834 bytes
EDM Help files	VOSSv811_HELP_EDM_gzip.zip	4,328,151 bytes
Logs reference	VOSS8400.8.1.1.0_edoc.tar	66,027,520 bytes
Software image	VOSS8400.8.1.1.0.tgz	185,413,197 bytes
Open source software - Master copyright file	VOSS8400.8.1.1.0_oss-notice.html	2,766,416 bytes
YANG model	restconf_yang.tgz	506,020 bytes

Table 13: XA1400 Series Software Filenames and Sizes

Description	File	Size
SHA512 Checksum files	VOSS1400.8.1.1.0.sha512	1,401 bytes
MD5 Checksum files	VOSS1400.8.1.1.0.md5	537 bytes
MIB - supported object names	VOSS1400.8.1.1.0_mib_sup.txt	1,062,389 bytes
MIB - zip file of all MIBs	VOSS1400.8.1.1.0_mib.zip	1,155,202 bytes
MIB - objects in the OID compile order	VOSS1400.8.1.1.0_mib.txt	7,662,834 bytes
EDM Help files	VOSSv811_HELP_EDM_gzip.zip	4,328,151 bytes
Logs reference	VOSS1400.8.1.1.0_edoc.tar	66,027,520 bytes
Software image	VOSS1400.8.1.1.0.tgz	347,774,661 bytes
Open source software - Master copyright file	VOSS1400.8.1.1.0_oss-notice.html	2,766,416 bytes

Chapter 3: Upgrade and Downgrade Considerations

See the Administering VOSS document for detailed image management procedures that includes information about the following specific upgrade considerations:

- IPv6:
  - Notes for systems using IPv6 static neighbors
  - Considerations for IPv6 VRRP or DHCP Relay configurations saved in VOSS 4.1 or 4.2
- Fabric:
  - Pre-upgrade instructions for IS-IS metric type
  - Upgrade considerations for IS-IS enabled links with HMAC-MD5 authentication
  - The following releases included modified Zero Touch Fabric Configuration support that impacts upgrades from earlier releases: VOSS 7.1.3 and later, VOSS 8.0.6 and later, and VOSS 8.1 and later.
- Upgrade considerations regarding MACsec replay-protect configuration
- Upgrade support for the nni-mstp boot configuration flag
- TACACS+ upgrade consideration
- Considerations for switches running an Extreme Insight virtual service configured prior to VOSS 8.0.5.
- Considerations for VLANs or MLTs where the VLAN or MLT name uses all numbers.
- Considerations for digital certificates configured prior to VOSS 8.1.

If your configuration includes one of the preceding scenarios or features, read the upgrade information in Administering VOSS before you begin an image upgrade.

Important:

Notice for VSP 4450GSX-PWR+, VSP 4450GSX, VSP 4450GTX-HT-PWR+, VSP 7200 Series, and VSP 8000 Series.

For these switch models running VOSS versions earlier than VOSS 6.1.x, you must first upgrade to VOSS 6.1.x before you can upgrade to VOSS 7.0 and later. Ensure that you save and back up your existing configuration before and after you upgrade to the intermediate 6.1.x release.

The same restriction applies to downgrades from VOSS 7.0 and later to releases earlier than VOSS 6.1.x. You must first downgrade to VOSS 6.1.x.

Supported Upgrade Paths

This section identifies the software releases for which upgrades to this release have been validated.

Supported Upgrade Paths for VSP 4450 Series, VSP 7200 Series, and VSP 8000 Series

Validated upgrade paths are VOSS 6.1.x, VOSS 7.x, or VOSS 8.0.x to VOSS 8.1 to 8.1.1.

Release 7.0 introduced a new Linux kernel on these switch models. Upgrades to this release are only supported from VOSS 6.1.x and after.

For these switch models running older VOSS versions, you must first upgrade to 6.1.x before you can upgrade to VOSS 8.1. Ensure that you save and back up your existing configuration before and after you upgrade to the intermediate 6.1.x release.

Important:

If you upgrade to VOSS 8.1 and then need to downgrade to a release earlier than VOSS 6.1.x, you must also do so in steps by first downgrading to a VOSS 6.1.x release before downgrading to the desired release.

Supported Upgrade Paths for VSP 7432CQ

Validated upgrade paths are VOSS 8.0.x to VOSS 8.1 to VOSS 8.1.1.

Supported Upgrade Paths for VSP 7400-48Y

Validated upgrade paths are VOSS 8.0.5.x or later to VOSS 8.1 to VOSS 8.1.1.

Supported Upgrade Paths for XA1400 Series

Validated upgrade paths are VOSS 8.0.50 to VOSS 8.1 to VOSS 8.1.1.

Supported Upgrade Paths for VSP 4900-48P

Validated upgrade paths are VOSS 8.1 to VOSS 8.1.1.

Upgrading DvR Configurations from Releases 6.0.1.1 and Earlier to 6.0.1.2 and Later

Upgrade all DvR nodes to the same release as quickly as possible. Release 6.0.1.2 includes changes to I-SID ranges that are utilized for DvR communication, and thus introduces an incompatibility with DvR nodes running 6.0.1.1 and earlier, with 6.0.1.2 and beyond.

Important:

Because of the change in 6.0.1.2, Extreme Networks recommends a *minimum* software version of 6.0.1.2 in DvR deployments.

In order to minimize the impact of this incompatibility and the resulting loss of connectivity between DvR Controller nodes and Leaf nodes, upgrade all DvR Leaf nodes first. After you upgrade all Leaf nodes, upgrade the Controller nodes, which restores DvR connectivity to the already upgraded Leaf nodes.

Important:

During the time when the Leaf nodes and Controller nodes are running incompatible versions, no DvR connectivity exists between the Controller and Leaf nodes so plan this activity accordingly, such as during a maintenance window.

If you cannot perform the upgrade during a maintenance window, use the following upgrade order to minimize connectivity loss:

1. Upgrade one of the DvR Controller nodes (vIST cluster member).
2. Upgrade the first DvR Leaf vIST cluster member.
3. Upgrade the second DvR Leaf vIST cluster member.
4. Upgrade the other DvR Controller.

By following this upgrade order, you upgrade the first Controller and make it ready for the Leaf nodes as you upgrade them. The other Controller still uses the original software version to accommodate Leaf nodes yet to upgrade, which allows you to upgrade them one at a time. Upgrade the other Controller last. With this upgrade order, only the node you are upgrading experiences a connectivity loss.

Upgrading DvR Configuration from 6.0.1.0 or 6.0.1.1 to 6.1.x.x

To upgrade DvR Leaf nodes:

1. If vIST is configured, use the **`no dvr leaf virtual-ist`** command on the Leaf nodes.
2. Use the **`no dvr leaf`** command on the Leaf nodes.

   Important:

   Do not save the configuration.
3. Upgrade the software to 6.1.x.x on the Leaf nodes, and then reboot the nodes.

To upgrade DvR Controllers:

1. Use the **`no dvr controller`** command on the Controllers.

   Important:

   Do not save the configuration.
2. Upgrade the software to 6.1.x.x on the Controllers, and then reboot the Controllers.

Real Time Clock

The latest VSP switches have an updated real time clock (RTC) component, which is not compatible with some older software releases. If you have the new hardware, the switch prevents you from downgrading to an unsupported release.

The hardware revision number of the affected products has been updated to reflect this change. For each product in the affected product families, the following table identifies the hardware revisions, and higher, that contain the updated RTC component.

Model	Minimum Hardware Revision
VSP 4450GSX	11
VSP 4450GTX-HT-PWR+	11
VSP 7254XSQ and VSP 7254XTQ	13
VSP 8284XSQ	12
VSP 8404	10
VSP 8404C	12

The minimum versions of software required for proper functioning of the product with the new RTC component are as follows:

- 6.x software baseline – 6.1.6.0
- 7.x or later software baseline – 7.1.0.1

All other earlier software versions do not support the new RTC component.

Syslog RFC 5424 and Extreme Management Center Integration

For existing customers with saved configurations prior to VOSS 6.1.2.0 who are parsing the non RFC 5424 syslog format, the device defaults to the old format. When Extreme Management Center registers for syslog, it configures it to the RFC 5424 format and automatically changes the syslog and log formats.

Post Upgrade Configuration for Zero Touch Fabric Configuration or Dynamic Nickname Assignment

If you want to use either, or both, of these features in VOSS 7.0 or later, the following sections identify the possible configuration combinations:

- Option 1: Enable Zero Touch Fabric Connect configuration and Dynamic Nickname on page 34
- Option 2: Enable Dynamic Nickname Assignment on page 35
- Option 3: Enable Zero Touch Fabric Connect configuration on page 35
- Option 4: Disable Zero Touch Fabric Connect configuration and Dynamic Nickname Assignment on page 36

Note:

The following releases included modified Zero Touch Fabric Configuration support that impacts upgrades from earlier releases: VOSS 7.1.3 and later, VOSS 8.0.6 and later, and VOSS 8.1 and later.

- VOSS 7.1.3 and later
- VOSS 8.0.6 and later
- VOSS 8.1 and later

For general steps about how to upgrade the switch software, see Administering VOSS.

Option 1: Enable Zero Touch Fabric Configuration and Dynamic Nickname Assignment

1. Start the nodes with the VOSS 7.0 or later image in factory-default fabric mode.
   - Factory default fabric mode enables Zero Touch Fabric Configuration.
   - The switch configures SPBM and IS-IS to the following default values:
     - SPBM instance 1
     - Primary BVID 4051 and secondary BVID 4052
     - System ID uses default value (derived from the chassis base MAC)
     - Manual area and nickname are zero
   - The switch creates and enables IS-IS interfaces on FAN ports.
2. IS-IS adjacencies are not formed.
3. IS-IS interfaces are in listening mode. These interfaces do not send HELLO PDUs because there is no IS-IS manual area configured. These interfaces listen for incoming HELLO PDUs
4. The node learns the IS-IS manual area from the first HELLO PDU it receives on any IS-IS interface. This learned area is called the Dynamically Learned Area (DLA).
5. The node uses the DLA to send HELLO PDUs on all active IS-IS interfaces and form adjacencies if the IS-IS parameters match.
6. If all nodes in the network started in Zero Touch Fabric Configuration mode, configure the manual area on at least one to them, which has physical connectivity with the rest of the nodes using the FAN interfaces. This node is referred to as the *seed* node. The term seed node describes the starting event to build the SPB network if all nodes start in Zero Touch Fabric Configuration mode.
7. If you insert the new node in a network where SPB is already configured and is connected using the FAN port to the node on its IS-IS interface, the adjacency with that node comes up if it uses the same default BVLANs mentioned above.
8. Because Dynamic Nickname Assignment is not configured yet, nodes become nickname clients. The clients become FAN members and start advertising FAN membership using TLV 147.
9. The FAN is established based on FAN endpoint membership.
10. Select a node and enable the nickname server.

11. After detecting a nickname server exists in the network, the nickname client sends a request for a nickname to the server.
12. The server assigns a nickname, which the client node learns.

Option 2: Enable Dynamic Nickname Assignment

1. Start the nodes with the VOSS 7.0 or later image with the existing configuration.
   - Zero Touch Fabric Configuration is not enabled.
   - The SPBM and IS-IS configuration is based on the configuration file.
   - A manual area is configured.
2. Disable IS-IS.
3. Remove static nicknames on all nodes.
4. Nodes become nickname clients. The clients become FAN members and start advertising FAN membership using TLV 147.
5. The FAN is established based on FAN endpoint membership.
6. Select a node and enable the nickname server.
7. After detecting a nickname server exists in the network, the nickname client sends a request for a nickname to the server.
8. The server assigns a nickname, which the client node learns.

Option 3: Enable Zero Touch Fabric Configuration

1. Start the nodes with the VOSS 7.0 or later image in factory-default fabric mode.
   - Factory default fabric mode enables Zero Touch Fabric Configuration.
   - The switch configures SPBM and IS-IS to the following default values:
     - SPBM instance 1
     - Primary BVID 4051 and secondary BVID 4052
     - System ID uses default value (derived from the chassis base MAC)
     - Manual area and nickname are zero
   - The switch creates and enables IS-IS interfaces on FAN ports.
2. IS-IS adjacencies are not formed.
3. IS-IS interfaces are in listening mode. These interfaces do not send HELLO PDUs because there is no IS-IS manual area configured. These interfaces listen for incoming HELLO PDUs
4. The node learns the IS-IS manual area from the first HELLO PDU it receives on any IS-IS interface. This learned area is called the Dynamically Learned Area (DLA).
5. The node uses the DLA to send HELLO PDUs on all active IS-IS interfaces and form adjacencies if the IS-IS parameters match.
6. If all nodes in the network started in Zero Touch Fabric Configuration mode, configure the manual area on at least one to them, which has physical connectivity with the rest of the nodes using the FAN interfaces. This node is referred to as the *seed* node. The term seed node describes the starting event to build the SPB network if all nodes start in Zero Touch Fabric Configuration mode.

7. If you insert the new node in a network where SPB is already configured and is connected using the FAN port to the node on its IS-IS interface, the adjacency with that node comes up if it uses the same default BVLANs mentioned above.
8. Configure static nicknames on all nodes.

Option 4: Disable Zero Touch Fabric Configuration and Dynamic Nickname Assignment

1. Start the nodes with the VOSS 7.0 or later image with the existing configuration.
   - Zero Touch Fabric Configuration is not enabled.
   - The SPBM and IS-IS configuration is based on the configuration file.
   - A manual area is configured.
   - Static nicknames are configured.
2. Dynamic Nickname Assignment server and clients do not start.

Chapter 4: Hardware and Software Compatibility

This section lists the hardware compatibility for all VOSS platforms.

VSP 4000 Series Hardware

Part number	Model number	Initial release	Supported new feature release				
			7.1	8.0	8.0.5	8.1	8.1.1
EC4400004-E6	VSP 4450GSX-DC	4.0.50	Y	Y	Y	Y	Y
EC4400A03-E6	VSP 4450GTX-HT-PWR+	4.0.50	Y	Y	Y	Y	Y
EC4400A05-E6	VSP 4450GSX-PWR+	4.0	Y	Y	Y	Y	Y
EC4400A05-E6GS	VSP 4450GSX-PWR+ TAA Compliant	4.0.50	Y	Y	Y	Y	Y
EC4800078-E6	VSP 4850GTS-DC	3.0	Y	N	N	N	N
EC4800A78-E6 EC4800A78-E6GS	VSP 4850GTS	3.0	Y	N	N	N	N
EC4800A88-E6 EC4800A88-E6GS	VSP 4850GTS-PWR+	3.0	Y	N	N	N	N

VSP 4000 Series Operational Notes

- **Warning:**

  The USB FLASH drive on all models of VSP 4850 Series (factory built and converted from ERS 4850) is a permanent non-removable part of the switch that you must NEVER remove from the switch to ensure proper operation. Additionally, you must install the USB cover to ensure additional protection against removal. The USB FLASH drive on the VSP 4850 Series switch is uniquely and permanently bound to the operating system of the switch it is first used on and cannot be transferred to a different switch. Removal (and reinsertion) of the USB FLASH drive from the switch is not supported as it can permanently compromise the switch functionality and render it non-functional.

- On a VSP 4450 Series switch, when making the initial connection to the two 10 Gbps SFP+ ports with MACsec-capable PHY (ports 49 and 50), the remote device flaps two times before remaining up due to the MACsec probing done by the VSP 4450 Series switch.

VSP 4900 Series Hardware

Part number	Model number	Initial release	Supported new feature release	
			8.1	8.1.1
VSP 4900-48P	VSP 4900-48P	8.1	Y	Y
Versatile Interface Modules (VIM) **Note:** Ensure the switch runs, at a minimum, the noted initial software release before you install a VIM.				
VIM5-4X	VIM5-4X	8.1	Y	Y
VIM5-4XE	VIM5-4XE	8.1	Y	Y
VIM5-2Y	VIM5-2Y	8.1	Y	Y
VIM5-4YE	VIM5-4YE	8.1	Y	Y
VIM5-2Q	VIM5-2Q	8.1	Y	Y

Versatile Interface Module Operational Notes

The following table summarizes the operational capabilities of the various VIMs:

	VIM5-4X	VIM5-4XE	VIM5-2Y	VIM5-4YE	VIM5-2Q
Number of supported ports	4	4	2	2	1
Port speeds	• 1 Gbps • 10 Gbps	• 1 Gbps • 10 Gbps	• 10 Gbps or • 25 Gbps Both ports must operate at either 10 Gbps or 25 Gbps (default)	• 10 Gbps or • 25 Gbps Both ports must operate at either 10 Gbps or 25 Gbps (default)	• 40 Gbps • 10 Gbps (with channelization)
PHY present	No	Yes	Yes	Yes	No
Copper transceiver support (1 Gbps/10 Gbps)	10GBASE-T only	Both	10GBASE-T only	10GBASE-T only	Not applicable

Table continues…

	VIM5-4X	VIM5-4XE	VIM5-2Y	VIM5-4YE	VIM5-2Q
MACsec	Not supported	256 bit	Not supported	256 bit	Not supported
Forward Error Correction (FEC)	Not supported	Not supported	Not supported	Auto-FEC enabled	Not supported
1 Gbps Auto-Negotiation	Disabled	Enabled	Not applicable	Not applicable	Not applicable
10 Gbps Auto-Negotiation	Disabled	Disabled	Disabled	Disabled	Not applicable
25 Gbps Auto-Negotiation	Not applicable	Not applicable	Disabled	• Enabled for DACs • Disabled for AOCs, optical transceivers	Not applicable

Note:

Auto-Negotiation values are automatically set based on the type of transceiver detected.

VSP 7200 Series Hardware

Part number	Model number	Initial release	Supported new feature release				
			7.1	8.0	8.0.5	8.1	8.1.1
EC720001F-E6	VSP 7254XSQ DC (front to back airflow)	4.2.1	Y	Y	Y	Y	Y
EC7200A1B-E6 (back-to-front airflow) EC7200A1F-E6 (front-to-back airflow)	VSP 7254XSQ	4.2.1	Y	Y	Y	Y	Y
EC720002F-E6	VSP 7254XTQ DC (Front to back airflow)	4.2.1	Y	Y	Y	Y	Y
EC7200A2B-E6 (back-to-front airflow)	VSP 7254XTQ	4.2.1	Y	Y	Y	Y	Y

Table continues…

Part number	Model number	Initial release	Supported new feature release				
			7.1	8.0	8.0.5	8.1	8.1.1
EC7200A2F-E6 (front-to-back airflow)							
EC7200A3B-E6 (back-to-front airflow) EC7200A3F-E6 (front-to-back airflow)	VSP 7254XSQ Port Licensed	5.1	Y	Y	Y	Y	Y
EC7200A4B-E6 (back-to-front airflow) EC7200A4F-E6 (front-to-back airflow)	VSP 7254XTQ Port Licensed	5.1	Y	Y	Y	Y	Y

VSP 7200 Series Operational Notes

- The VSP 7254XSQ has a PHYless design, which is typical for Data Center top of rack switches. The benefits of a PHYless design are lower power consumption and lower latency. However, due to the PHYless design, the following transceivers that require electronic dispersion compensation (EDC) for proper operation are not supported:
  - AA1403017-E6: 1-port 10GBASE-LRM SFP+
  - AA1403016-E6: 1-port 10GBase-ZR/ZW SFP+

    The AA1403165 10GBASE-ZR CWDM DDI SFP+ transceiver can be substituted for AA1403016-E6 10GBASE-ZR/ZW SFP+
- Software partitions the switch into two logical slots: Slot 1 and Slot 2.
  - Slot 1: 10 Gbps ports: 1 - 48
  - Slot 2: 40 Gbps ports: 1 - 6
- Channelization is supported on the 40 Gbps QSFP+ ports.
- MACsec support:
  - MACsec is only supported on the VSP 7254XTQ 10 Gbps ports.
  - MACsec is not supported on VSP 7254XSQ 10 Gbps ports
  - MACsec is not supported on VSP 7254XTQ and VSP 7254XSQ 40 Gbps ports whether channelization is enabled or not.
- Port licensing support on the port licensed VSP 7254XSQ fiber switch:
  - 24 ports (Slot 1, ports 25 to 48) out of the 48 1/10 GbE SFP/SFP+ ports require a Port License to be unlocked.
  - two ports (Slot 2, ports 5 and 6) out of the six 40 GbE QSFP+ ports require a Port License to be unlocked.

- Port licensing support on the port licensed VSP 7254XTQ copper switch:
  - 24 ports (Slot 1, ports 25 to 48) out of the 48 100 Mbps/1 GbE/10 GbE RJ-45 ports require a Port License to be unlocked.
  - two ports (Slot 2, ports 5 and 6) out of the six 40 GbE QSFP+ ports require a Port License to be unlocked.
- 1000BASE-T SFP (AA1419043-E6) will only operate at 1 Gbps speeds when used on a VSP 7254XSQ.
- When you use 1 Gigabit Ethernet SFP transceivers on VSP 7254XSQ, the software disables auto-negotiation on the port:
  - If you use 1 Gbps fiber SFP transceivers, the remote end must also have auto-negotiation disabled.
  - If you use 1 Gbps copper SFP transceivers, the remote end must have auto-negotiation enabled. If not, the link will not be established.
- When a port on VSP 7254XSQ is disabled or enabled, or a cable replaced, or the switch rebooted, the remote link can flap twice.
- Enable auto-negotiation to ensure proper operation at 100 Mbps speeds on VSP 7254XTQ:
  - Link instability will be seen if both ends are set to 100 Mbps auto-negotiation disabled and you use a straight through cable.
  - If Link instability is seen when you use a cross-over cable, a port disable or enable can fix the issue.

VSP 7400 Series Hardware

Part number	Model Number	Initial release	Supported new feature release			
			8.0	8.0.5	8.1	8.1.1
VSP7400-32C (no power supplies or fans) VSP7400-32C-AC-F (front-to-back airflow) VSP7400-32C-AC-R (back-to-front airflow)	VSP 7432CQ	8.0	Y	Y	Y	Y
VSP7400-48Y-8C (no power supplies or fans)	VSP 7400-48Y	8.0.5	N	Y	Y	Y

Table continues…

Part number	Model Number	Initial release	Supported new feature release			
			8.0	8.0.5	8.1	8.1.1
VSP7400-48Y-8C-AC-F (front-to-back airflow) VSP7400-48Y-8C-AC-R (back-to-front airflow)						

VSP 7400 Series Operational Notes

The VSP 7400 Series has a PHYless design. The benefits of a PHYless design are lower power consumption and lower latency. However, due to the PHYless design, the following transceivers that require electronic dispersion compensation (EDC) for proper operation are not supported:

- AA1403017-E6: 1-port 10GBASE-LRM SFP+
- AA1403016-E6: 1-port 10GBase-ZR/ZW SFP+

  The AA1403165 10GBASE-ZR CWDM DDI SFP+ transceiver can be substituted for AA1403016-E6 10GBASE-ZR/ZW SFP+

The following list provides operational notes for VSP 7432CQ.

- Ports 31 and 32 (low) or ports 29, 30, 31, and 32 (high) are reserved for internal use when Fabric Connect is used.
- The QSFP28 ports support the use of QSFP28 and QSFP+ transceivers:
  - The software detects the transceiver type and sets the port speed as either 100 Gbps for QSFP28 or 40 Gbps for QSFP+.
- Channelization:
  - Channelization is not supported on port 28.
  - Supports 4x10 Gbps when channelization is enabled and QSFP+ transceiver is detected.
  - Supports 4x25 Gbps when channelization is enabled and QSFP28 transceiver is detected.

The following list provides operational notes for VSP 7400-48Y.

- Ports 55 and 56 (low) or ports 53, 54, 55, and 56 (high) are reserved for internal use when Fabric Connect is used.
- The QSFP28 ports support the use of QSFP28 and QSFP+ transceivers:
  - The software detects the transceiver type and sets the port speed as either 100 Gbps for QSFP28 or 40 Gbps for QSFP+.
- The SFP28 ports support the use of SFP28, SFP, and SFP+ transceivers.
  - The software detects the transceiver type and sets the port speed as either 25 Gbps for SFP28,.1 Gbps for SFP, or 10 Gbps for SFP+.
  - Auto-Negotiation is not supported when a 25 Gbps port operates at 1 Gbps. The following log message appears on the switch: `Auto-Negotiation enabled but not applied to port 1/1 since 1G transceiver is present..`

- Channelization is not supported. As a result, you cannot use the following optical components:
  - 40 Gbps or 100 Gbps breakout cables
  - QSFP28 to SFP28 Adapter (PN: 10506)

VSP 8000 Series Hardware

Part number	Model number	Initial release	Supported new feature release				
			7.1	8.0	8.0.5	8.1	8.1.1
EC8200A01-E6 EC8200A01-E6GS	VSP 8284XSQ	4.0	Y	Y	Y	Y	Y
EC8200001-E6	VSP 8284XSQ DC	4.0.50	Y	Y	Y	Y	Y
EC8400001-E6	VSP 8404 DC	4.2.1	Y	Y	Y	Y	Y
EC8400A01-E6 EC8200A01-E6GS	VSP 8404	4.2	Y	Y	Y	Y	Y
EC8400002-E6	VSP 8404C DC	5.3	Y	Y	Y	Y	Y
EC8400A02-E6 EC8200A02-E6GS	VSP 8404C	5.3	Y	Y	Y	Y	Y
Ethernet Switch Modules (ESM) — VSP 8400 Series only **Important:** Ensure the switch runs, at a minimum, the noted initial software release before you install an ESM.							
EC8404001-E6 EC8404001-E6GS	8424XS	4.2	Y	Y	Y	Y	Y
EC8404002-E6 EC8404002-E6GS	8424XT	4.2	Y	Y	Y	Y	Y
EC8404003-E6 EC8404003-E6GS	8408QQ	4.2	Y	Y	Y	Y	Y
EC8404005-E6 EC8404005-E6GS	8418XSQ	4.2	Y	Y	Y	Y	Y
EC8404006-E6 EC8404006-E6GS	8418XTQ	5.0	Y	Y	Y	Y	Y
EC8404007-E6 EC8404007-E6GS	8424GS	5.0	Y	Y	Y	Y	Y
EC8404008-E6	8424GT	5.0	Y	Y	Y	Y	Y

Table continues…

Part number	Model number	Initial release	Supported new feature release				
			7.1	8.0	8.0.5	8.1	8.1.1
EC8404008-E6GS							
EC8404009-E6 EC8404009-E6GS	8402CQ Supported in VSP 8404C only	5.3	Y	Y	Y	Y	Y

XA1400 Series Hardware

Part number	Model number	Initial release	Supported new feature release		
			8.0.50	8.1	8.1.1
XA1440	ExtremeAccess Platform 1440 (XA1440)	8.0.50	Y	Y	Y
XA1480	ExtremeAccess Platform 1480 (XA1480)	8.0.50	Y	Y	Y

Transceivers

The software allows the use of transceivers and direct attach cables from any vendor, which means that the switch will bring up the port operationally when using any transceiver. Extreme Networks does not provide support for operational issues related to the use of non-Extreme Networks branded transceivers and direct attached cables used in the switches.

Extreme Networks supports SFP transceivers with the following part numbers: AA1419013–E5, AA1419014–E5, AA1419015–E5, and AA1419025–E5 to AA1419040–E5. However, Extreme Networks strongly recommends using the newer DDI versions of these SFP transceivers.

Note:

Although VSP 8000 Series and VSP 7200 Series support 10 Gigabit and 40 Gigabit DAC cables in forgiving mode, in releases earlier than VOSS 4.2.1, the command output for **`show pluggable-optical-modules basic`** displays the corresponding vendor name rather than leaving the vendor name field blank.

The following table indicates where to find more information about optical transceivers and components.

Extreme Networks optical transceivers and components	*Extreme Networks Pluggable Transceivers Installation Guide*
Compatibility for Extreme Networks SFP, SFP+, QSFP+, and QSFP28 transceiver modules with the VOSS -capable switches	*Extreme Hardware/Software Compatibility and Recommendation Matrices*

Auto-Negotiation

Use Auto-Negotiation to allow the device to automatically negotiate the best common data rate and duplex mode to use between two Auto-Negotiation-capable Ethernet devices.

When you use a 1 Gigabit SFP transceiver on a 10 Gigabit SFP+ port, you must enable Auto-Negotiation if it is not enabled already. However, if you use 1 Gigabit SFP transceivers on a VSP 4000 Series switch that is connected to third party switches at the remote end, you must have Auto-Negotiation enabled at all times; this applies to SFP transceivers installed in a 1 Gigabit SFP port or a 10 Gigabit SFP+ port.

For VSP 7254XSQ, Auto-Negotiation is always disabled for 1 Gigabit Ethernet transceivers. If using a 1000BASE-T SFP, the remote 1000BASE-T interface must have Auto-Negotiation enabled. If not, the link will not be established. Also note that because the SFP+ ports on the VSP 7254XSQ only support 1 and 10 Gbps speeds, the AA1419043-E6 1000BASE-T SFP will only operate at 1G speeds.

If you use 1 Gbps fiber SFP transceivers, Auto-Negotiation is always disabled so the remote end must also have Auto-Negotiation disabled. Otherwise this is not a supported configuration with VSP 7254XSQ.

Forward Error Correction (FEC)

Forward Error Correction (FEC) is a method of obtaining error control in data transmission over an unreliable or noisy channel in which the source (transmitter) encodes the data in a redundant way by using an error correcting code (ECC). This redundancy enables a destination (receiver) to detect a limited number of errors and correct them without requiring a re-transmission.

For more information about FEC, see Administering VOSS.

Power Supply Compatibility

You can use certain power supplies in more than one platform. This section lists the power supplies and indicates the compatible platforms.

For more specific information on each power supply, see the following documents:

- Installing the Virtual Services Platform 4850GTS Series
- Installing the Virtual Services Platform 4450GTX-HT-PWR+
- Installing the Virtual Services Platform 4450GSX-PWR+
- VSP 4900 Series Switches: Hardware Installation Guide
- Installing the Virtual Services Platform 7200 Series

- VSP 7400 Series Switches: Hardware Installation Guide
- Installing the Virtual Services Platform 8000 Series
- XA1400 Series Switches: Hardware Installation Guide

Table 14: VSP 4000 Series Power Supplies

Platform	300 W AC AL1905A08-E5	300 W DC AL1905005-E5	1,000 W AC AL1905A21-E6	1,000 W AC-HT EC4005A03-E6HT
VSP 4850GTS-DC	—	Y	—	—
VSP 4850GTS-PWR+	—	—	Y	Y
VSP 4850GTS	Y	—	—	—
VSP 4450GTX-HT-PWR+	—	—	—	Y
VSP 4450GSX-DC	—	Y	—	—
VSP 4450GSX-PWR+	—	—	Y	Y

Table 15: VSP 4900 Series Power Supplies

Platform	1100 W AC 10941
VSP 4900-48P	Y

Table 16: VSP 7200 Series and VSP 8000 Series Power Supplies

Platform	460 W AC front-to-back EC7205A1F-E6	460 W AC back-to-front EC7205A1B-E6	800 W AC front-to-back EC8005A01-E6	800 W AC front-to-back EC7205A0F-E6	800 W AC back-to-front EC7205A0B-E6	800 W DC front-to-back EC8005001-E6
VSP 8284XSQ	—	—	Y	—	—	—
VSP 8284XSQ DC	—	—	—	—	—	Y
VSP 8404	—	—	Y	—	—	—
VSP 8404 DC	—	—	—	—	—	Y
VSP 8404C	—	—	Y	—	—	—
VSP 8404C DC	—	—	—	—	—	Y
VSP 7254XSQ front-to-back	Y	—	—	—	—	—
VSP 7254XSQ back-to-front	—	Y	—	—	—	—

Table continues…

Platform	460 W AC front-to-back EC7205A1F-E6	460 W AC back-to-front EC7205A1B-E6	800 W AC front-to-back EC8005A01-E6	800 W AC front-to-back EC7205A0F-E6	800 W AC back-to-front EC7205A0B-E6	800 W DC front-to-back EC8005001-E6
VSP 7254XTQ front-to-back	—	—	—	Y	—	—
VSP 7254XTQ back-to-front	—	—	—	—	Y	—
VSP 7254XSQ DC	—	—	—	—	—	Y
VSP 7254XTQ DC	—	—	—	—	—	Y

The following table for VSP 7400 Series includes the orderable part number as well as the model number or model name, as it appears on the power supply.

Table 17: VSP 7400 Series Power Supplies

Platform	750 W AC front-to-back XN-ACPWR-750W-F	750 W AC back-to-front XN-ACPWR-750W-R	750 W DC front-to-back XN-DCPWR-750W-F	750 W DC back-to-front XN-DCPWR-750W-R
Model Number/ Model Name	700-013684-0100/ MC75A4-3	700-013917-0000/ MC75A4-3-001	700-013670-0000	700-013670-0100
VSP 7432CQ front-to-back	Y	—	—	—
VSP 7432CQ back-to-front	—	Y	—	—
VSP 7432CQ front-to-back DC	—	—	Y	—
VSP 7432CQ back-to-front DC	—	—	—	Y
VSP 7400-48Y front-to-back	Y	—	—	—
VSP 7400-48Y back-to-front	—	Y	—	—
VSP 7400-48Y front-to-back DC	—	—	Y	—
VSP 7400-48Y back-to-front DC	—	—	—	Y

Table 18: XA1400 Series Power Supplies

Platform	12 V DC XA1400-PWR-ADPT
XA1440	Y
XA1480	Y

Chapter 5: Scaling

This section documents scaling capabilities of the VOSS platforms.

The scaling and performance information shown in the following tables is provided for the purpose of assisting with network design. It is recommended that network architects and administrators design and manage networks with an appropriate level of network scaling “head room.” The scaling and performance figures provided have been verified using specific network topologies using limited switch configurations. There is no guarantee that the scaling and performance figures shown are applicable to all network topologies and switch configurations and are provided as a realistic estimation only. If you experience scaling and performance characteristics that you feel are sufficiently below what has been documented, contact Extreme Networks technical support for additional assistance.

Note:

If your switch uses Advanced Feature Bandwidth Reservation in Full Feature mode, this affects scaling information that is based on the number of available ports. If you enable the boot configuration flag for this feature, remember to deduct the number of reserved ports from the documented scaling maximum. Not all hardware platforms require this feature to provide full feature support. For more information, see Administering VOSS.

Layer 2

Table 19: Layer 2 Maximums

Attribute	Product	Maximum number supported
Directed Broadcast interfaces **Note:** The number of Directed Broadcast interfaces must be less than or equal to 200. However, if you configure VLANs with both NLB and Directed Broadcast, you can only scale up to 100 VLANs.	VSP 4450 Series	n/a
	VSP 4900 Series	200 See Note.
	VSP 7200 Series	200 See Note.
	VSP 7400 Series	200 See Note.
	VSP 8000 Series	200

Table continues…

Attribute	Product	Maximum number supported
		See Note.
	XA1400 Series	n/a
MAC table size (without SPBM)	VSP 4450 Series	32,000
	VSP 4900 Series	80,000
	VSP 7200 Series	224,000
	VSP 7400 Series	160,000
	VSP 8000 Series	224,000
	XA1400 Series	2,000 for XA1440 4,000 for XA1480
MAC table size (with SPBM)	VSP 4450 Series	16,000
	VSP 4900 Series	40,000
	VSP 7200 Series	112,000
	VSP 7400 Series	80,000
	VSP 8000 Series	112,000
	XA1400 Series	2,000 for XA1440 4,000 for XA1480
Endpoint Tracking MAC addresses per switch	VSP 4450 Series	n/a
	VSP 4900 Series	n/a
	VSP 7200 Series	8,000
	VSP 7400 Series	8,000
	VSP 8000 Series	8,000
	XA1400 Series	n/a
Port-based VLANs	VSP 4450 Series	4,059
	VSP 4900 Series	2,000 with RESTCONF 4,000 without RESTCONF
	VSP 7200 Series	4,059
	VSP 7400 Series	4,059
	VSP 8000 Series	4,059
	XA1400 Series	500
Private VLANs	VSP 4450 Series	200
	VSP 4900 Series	200
	VSP 7200 Series	200
	VSP 7400 Series	200

Table continues…

Attribute	Product	Maximum number supported
	VSP 8000 Series	VSP 8404C = 400 Other VSP 8000 Series platforms = 200
	XA1400 Series	n/a
Protocol-based VLANs (IPv6 only)	VSP 4450 Series	1
	VSP 4900 Series	1
	VSP 7200 Series	1
	VSP 7400 Series	1
	VSP 8000 Series	1
	XA1400 Series	n/a
RSTP instances	VSP 4450 Series	1
	VSP 4900 Series	1
	VSP 7200 Series	1
	VSP 7400 Series	1
	VSP 8000 Series	1
	XA1400 Series	1
MSTP instances	VSP 4450 Series	12
	VSP 4900 Series	12
	VSP 7200 Series	12
	VSP 7400 Series	64
	VSP 8000 Series	12
	XA1400 Series	12
LACP aggregators	VSP 4450 Series	24
	VSP 4900 Series	52 (48 fixed ports + 4 VIM ports)
	VSP 7200 Series	54 (up to 72 with channelization)
	VSP 7400 Series	VSP 7432CQ = 32 (up to 125 with channelization) configured in Full Port mode VSP 7400-48Y = 56 configured in Full Port mode
	VSP 8000 Series	84 (up to 96 with channelization)

Table continues…

Attribute	Product	Maximum number supported
	XA1400 Series	8
Ports per LACP aggregator	VSP 4450 Series	8 active
	VSP 4900 Series	8 active
	VSP 7200 Series	8 active
	VSP 7400 Series	8 active
	VSP 8000 Series	8 active
	XA1400 Series	8
MLT groups	VSP 4450 Series	50
	VSP 4900 Series	52 (48 fixed ports + 4 VIM ports)
	VSP 7200 Series	54 (up to 72 with channelization)
	VSP 7400 Series	VSP 7432CQ = 32 (up to 125 with channelization) configured in Full Port mode VSP 7400-48Y = 56 configured in Full Port mode
	VSP 8000 Series	84 (up to 96 with channelization)
	XA1400 Series	8
Ports per MLT group	VSP 4450 Series	8
	VSP 4900 Series	8
	VSP 7200 Series	8
	VSP 7400 Series	8
	VSP 8000 Series	8
	XA1400 Series	8
LST groups	VSP 4450 Series	48
	VSP 4900 Series	48
	VSP 7200 Series	48
	VSP 7400 Series	48
	VSP 8000 Series	48
	XA1400 Series	n/a
Interfaces per LST group	VSP 4450 Series	8 upstream 128 downstream

Table continues…

Attribute	Product	Maximum number supported
	VSP 4900 Series	8 upstream 128 downstream
	VSP 7200 Series	8 upstream 128 downstream
	VSP 7400 Series	8 upstream 128 downstream
	VSP 8000 Series	8 upstream 128 downstream
	XA1400 Series	n/a
SLPP VLANs	VSP 4450 Series	128
	VSP 4900 Series	128
	VSP 7200 Series	128
	VSP 7400 Series	500
	VSP 8000 Series	128
	XA1400 Series	128
VLACP interfaces	VSP 4450 Series	50
	VSP 4900 Series	52 (48 fixed ports + 4 VIM ports)
	VSP 7200 Series	54 (up to 72 with channelization)
	VSP 7400 Series	VSP 7432CQ = 32 (up to 125 with channelization) configured in Full Port mode VSP 7400-48Y = 56 configured in Full Port mode
	VSP 8000 Series	84 (up to 96 with channelization)
	XA1400 Series	8
Microsoft NLB cluster IP interfaces **Note:** The number of NLB cluster IP interfaces multiplied by the number of configured clusters must be less than or equal to 200. The number of NLB cluster IP interfaces is the key, not the number of VLANs. You	VSP 4450 Series	n/a
	VSP 4900 Series	200 See Note.
	VSP 7200 Series	200 See Note.

Table continues…

Attribute	Product	Maximum number supported
can configure 1 VLAN with up to 200 NLB cluster IP interfaces or configure up to 200 VLANs with 1 NLB cluster IP interface per VLAN. For example: 1 virtual interface per cluster x 200 clusters = 200 or 2 virtual interfaces per cluster x 100 clusters = 200 However, if you configure VLANs with both NLB and Directed Broadcast, you can only scale up to 100 VLANs assuming there is only 1 NLB cluster IP interface per VLAN.	VSP 7400 Series	200 See Note.
	VSP 8000 Series	200 See Note.
	XA1400 Series	n/a

IP Unicast

Table 20: IP Unicast Maximums

Attribute	Product	Maximum number supported
Note: The maximum number of IP interfaces is based on the following formulas: • For VSP 4900 Series : - If you disable the VRF scaling boot configuration flag: • = 500 – (# of VRRP IPv4 interfaces) - (# of VRRP IPv6 interfaces) – (# of RSMLT interfaces) – 2 (if IP Shortcuts is enabled) – 3x(# of VRFs) - If you enable the VRF scaling boot configuration flag: • = 500 – (# of VRRP IPv4 interfaces) – (# of VRRP IPv6 interfaces) - (# of RSMLT interfaces) – 2 (if IP Shortcuts is enabled) – 3 • For VSP 7200 Series, VSP 8200 Series, and VSP 8400 Series: - If you disable the VRF scaling boot configuration flag: • = 505 – (# of VRRP IPv4 interfaces) - (# of VRRP IPv6 interfaces) – (# of RSMLT interfaces) – 2 (if IP Shortcuts is enabled) – 3x(# of VRFs) - If you enable the VRF scaling boot configuration flag: • = 505 – (# of VRRP IPv4 interfaces) – (# of VRRP IPv6 interfaces) - (# of RSMLT interfaces) – 2 (if IP Shortcuts is enabled) – 3		

Table continues…

Attribute	Product	Maximum number supported
• For VSP 7400 Series: - If you disable the VRF scaling boot configuration flag: • = 1000 – (# of VRRP IPv4 interfaces) - (# of VRRP IPv6 interfaces) – (# of RSMLT interfaces) – 2 (if IP Shortcuts is enabled) – 3x(# of VRFs) - If you enable the VRF scaling boot configuration flag: • = 1000 – (# of VRRP IPv4 interfaces) – (# of VRRP IPv6 interfaces) - (# of RSMLT interfaces) – 2 (if IP Shortcuts is enabled) – 3		
IP interfaces (IPv4 or IPv6 or IPv4+IPv6)	VSP 4450 Series	256
	VSP 4900 Series	500 See Note.
	VSP 7200 Series	505 See Note.
	VSP 7400 Series	1,000 See Note.
	VSP 8000 Series	VSP 8404C = 500 Other VSP 8000 Series platforms = 505 See Note.
	XA1400 Series	500 (IPv4 only)
VRRP interfaces (IPv4 or IPv6)	VSP 4450 Series	64
	VSP 4900 Series	252 See Note.
	VSP 7200 Series	252 See Note.
	VSP 7400 Series	500 See Note.
	VSP 8000 Series	252 See Note.
	XA1400 Series	64 (IPv4 only)
Routed Split Multi-Link Trunking (RSMLT) interfaces (IPv4 or IPv6 or IPv4+IPv6)	VSP 4450 Series	252
	VSP 4900 Series	252
	VSP 7200 Series	252 See Note.
	VSP 7400 Series	500

Table continues…

Attribute	Product	Maximum number supported
		See Note.
	VSP 8000 Series	252 See Note.
	XA1400 Series	n/a
VRRP interfaces with fast timers (200ms) - IPv4/IPv6	VSP 4450 Series	24
	VSP 4900 Series	24
	VSP 7200 Series	24
	VSP 7400 Series	24
	VSP 8000 Series	24
	XA1400 Series	24
DvR Virtual IP interfaces	VSP 4450 Series	501 with vIST 502 without vIST
	VSP 4900 Series	501 with vIST 502 without vIST
	VSP 7200 Series	501 with vIST 502 without vIST
	VSP 7400 Series	999 with vIST 1,000 without vIST
	VSP 8000 Series	501 with vIST 502 without vIST
	XA1400 Series	n/a
ECMP groups/paths per group	VSP 4450 Series	500/4
	VSP 4900 Series	1,000/8
	VSP 7200 Series	1,000/8
	VSP 7400 Series	1,000/8
	VSP 8000 Series	1,000/8
	XA1400 Series	500/8
OSPF v2/v3 interfaces	VSP 4450 Series	100
	VSP 4900 Series	500
	VSP 7200 Series	500
	VSP 7400 Series	500
	VSP 8000 Series	500
	XA1400 Series	48 (v2 only)
OSPF v2/v3 neighbors (adjacencies)	VSP 4450 Series	100

Table continues…

Attribute	Product	Maximum number supported
	VSP 4900 Series	500
	VSP 7200 Series	500
	VSP 7400 Series	500
	VSP 8000 Series	500
	XA1400 Series	24 (v2 only)
OSPF areas	VSP 4450 Series	12 for each VRF 64 for the switch
	VSP 4900 Series	12 for each VRF 80 for the switch
	VSP 7200 Series	12 for each VRF 80 for the switch
	VSP 7400 Series	12 for each VRF 80 for the switch
	VSP 8000 Series	12 for each VRF 80 for the switch
	XA1400 Series	12 for each VRF 64 for each switch
IPv4 ARP table	VSP 4450 Series	6,000
	VSP 4900 Series	32,000
	VSP 7200 Series	32,000
	VSP 7400 Series	56,000 non-SPB deployments 40,000 SPB deployments
	VSP 8000 Series	32,000
	XA1400 Series	2,000 for XA1440 4,000 for XA1480
IPv4 CLIP interfaces	VSP 4450 Series	64
	VSP 4900 Series	64
	VSP 7200 Series	64
	VSP 7400 Series	64
	VSP 8000 Series	64
	XA1400 Series	64
IPv4 RIP interfaces	VSP 4450 Series	200
	VSP 4900 Series	200

Table continues…

Attribute	Product	Maximum number supported
	VSP 7200 Series	200
	VSP 7400 Series	200
	VSP 8000 Series	200
	XA1400 Series	200
IPv4 BGP peers	VSP 4450 Series	12
	VSP 4900 Series	256
	VSP 7200 Series	256
	VSP 7400 Series	256
	VSP 8000 Series	256
	XA1400 Series	12
IPv4 VRFs with iBGP	VSP 4450 Series	16
	VSP 4900 Series	16
	VSP 7200 Series	16
	VSP 7400 Series	16
	VSP 8000 Series	16
	XA1400 Series	n/a
IPv4 VRF instances For additional information, see VRF Scaling on page 90.	VSP 4450 Series	128 including GRT
	VSP 4900 Series	258 including mgmt VRF and GRT
	VSP 7200 Series	256 including mgmt VRF and GRT
	VSP 7400 Series	256 including mgmt VRF and GRT
	VSP 8000 Series	256 including mgmt VRF and GRT
	XA1400 Series	24 including GRT
IPv4 static ARP entries	VSP 4450 Series	200 for each VRF 1,000 for the switch
	VSP 4900 Series	2,000 for each VRF 10,000 for the switch
	VSP 7200 Series	2,000 for each VRF 10,000 for the switch
	VSP 7400 Series	2,000 for each VRF 10,000 for the switch
	VSP 8000 Series	2,000 for each VRF

Table continues…

Attribute	Product	Maximum number supported
		10,000 for the switch
	XA1400 Series	200 for each VRF 1,000 for the switch
IPv4 static routes	VSP 4450 Series	1,000 for each VRF 1,000 for the switch
	VSP 4900 Series	1,000 for each VRF 5,000 for the switch
	VSP 7200 Series	1,000 for each VRF 5,000 for the switch
	VSP 7400 Series	1,000 for each VRF 5,000 for the switch
	VSP 8000 Series	1,000 for each VRF 5,000 for the switch
	XA1400 Series	1,000 for each VRF 5,000 for the switch
IPv4 route policies	VSP 4450 Series	500 for each VRF 5,000 for the switch
	VSP 4900 Series	500 for each VRF 5,000 for the switch
	VSP 7200 Series	500 for each VRF 5,000 for the switch
	VSP 7400 Series	500 for each VRF 5,000 for the switch
	VSP 8000 Series	500 for each VRF 5,000 for the switch
	XA1400 Series	500 for each VRF 5,000 for the switch
IPv4 UDP forwarding entries	VSP 4450 Series	128
	VSP 4900 Series	512
	VSP 7200 Series	512
	VSP 7400 Series	1,024
	VSP 8000 Series	512
	XA1400 Series	128

Table continues…

Attribute	Product	Maximum number supported
IPv4 DHCP Relay forwarding entries	VSP 4450 Series	128
	VSP 4900 Series	1,024
	VSP 7200 Series	1,024
	VSP 7400 Series	1,024
	VSP 8000 Series	1,024
	XA1400 Series	128
IPv6 DHCP Snoop entries in Source Binding Table	VSP 4450 Series	1,024
	VSP 4900 Series	1,024
	VSP 7200 Series	1,024
	VSP 7400 Series	1,024
	VSP 8000 Series	1,024
	XA1400 Series	n/a
IPv6 Neighbor table	VSP 4450 Series	4,000
	VSP 4900 Series	8,000
	VSP 7200 Series	8,000
	VSP 7400 Series	32,000
	VSP 8000 Series	8,000
	XA1400 Series	n/a
IPv6 static entries in Source Binding Table	VSP 4450 Series	256
	VSP 4900 Series	256
	VSP 7200 Series	256
	VSP 7400 Series	256
	VSP 8000 Series	256
	XA1400 Series	n/a
IPv6 static neighbor records	VSP 4450 Series	128
	VSP 4900 Series	128 per VRF 512 per system
	VSP 7200 Series	128 per VRF 512 per system
	VSP 7400 Series	128 per VRF 512 per system
	VSP 8000 Series	128 per VRF 512 per system
	XA1400 Series	n/a

Table continues…

Attribute	Product	Maximum number supported
IPv6 CLIP interfaces	VSP 4450 Series	64
	VSP 4900 Series	64
	VSP 7200 Series	64
	VSP 7400 Series	64
	VSP 8000 Series	64
	XA1400 Series	n/a
IPv6 static routes	VSP 4450 Series	1,000
	VSP 4900 Series	1,000
	VSP 7200 Series	1,000
	VSP 7400 Series	1,000
	VSP 8000 Series	1,000
	XA1400 Series	n/a
IPv6 6in4 configured tunnels	VSP 4450 Series	64
	VSP 4900 Series	64
	VSP 7200 Series	64
	VSP 7400 Series	64
	VSP 8000 Series	64
	XA1400 Series	n/a
IPv6 DHCP Relay forwarding	VSP 4450 Series	128
	VSP 4900 Series	512 per switch 10 per VRF
	VSP 7200 Series	512 per switch 10 per VRF
	VSP 7400 Series	512
	VSP 8000 Series	512
	XA1400 Series	n/a
IPv6 BGP peers	VSP 4450 Series	12 Up to 8,000 IPv6 prefixes for BGPv6 peering
	VSP 4900 Series	256 Up to 8,000 IPv6 prefixes for BGPv6 peering
	VSP 7200 Series	256 Up to 8,000 IPv6 prefixes for BGPv6 peering

Table continues…

Attribute	Product	Maximum number supported
	VSP 7400 Series	256
	VSP 8000 Series	256 Up to 8,000 IPv6 prefixes for BGPv6 peering
	XA1400 Series	n/a
IPv6 VRFs with iBGP	VSP 4450 Series	16
	VSP 4900 Series	16
	VSP 7200 Series	16
	VSP 7400 Series	16
	VSP 8000 Series	16
	XA1400 Series	n/a
BFD VRF instances	VSP 4450 Series	16
	VSP 4900 Series	16
	VSP 7200 Series	16
	VSP 7400 Series	16
	VSP 8000 Series	16
	XA1400 Series	n/a
BFD sessions per switch (IPv4/IPv6) with default values	VSP 4450 Series	16
	VSP 4900 Series	16
	VSP 7200 Series	16
	VSP 7400 Series	16
	VSP 8000 Series	16
	XA1400 Series	n/a

Layer 3 Route Table Size

Table 21: Layer 3 Route Table Size Maximums

Attribute	Maximum number supported
IPv4 RIP routes	See Route Scaling on page 63.
IPv4 OSPF routes	
IPv4 BGP routes	
IPv4 SPB shortcut routes	

Table continues…

Attribute	Maximum number supported
IPv4 SPB Layer 3 VSN routes	
IPv6 OSPFv3 routes - GRT only	
IPv6 SPB shortcut routes - GRT only	
IPv6 RIPng routes	

Route Scaling

The following table provides information on IPv4 and IPv6 route scaling. The route table is a shared hardware resource where IPv4 routes consume one entry and IPv6 routes with a prefix length less than 64 consume two entries.

The route scaling does not depend on the protocol itself, but rather the general system limitation in the following configuration modes:

- URPF check mode - Enable this boot configuration flag to support Unicast Reverse Path Forwarding check mode.
- IPv6 mode - Enable this boot configuration flag to support IPv6 routes with prefix-lengths greater than 64 bits. When the IPv6-mode is enabled, the maximum number of IPv4 routing table entries decreases. This flag does not apply to all hardware platforms.

Table 22: VSP 4450 Series, VSP 4900 Series, VSP 7200 Series, and VSP 8000 Series

URPF mode	IPv6 mode	VSP 4450 Series			VSP 7200 Series, VSP 4900 Series, and VSP 8000 Series		
		IPv4	IPv6		IPv4	IPv6	
			Prefix less than 64	Prefix greater than 64		Prefix less than 64	Prefix greater than 64
No	No	15,744	7,887	256	15,488	7,744	n/a
No	Yes	n/a	n/a	n/a	7,488	3,744	2,000
Yes	No	7,744	3,872	256	7,488	3,744	n/a
Yes	Yes	n/a	n/a	n/a	3,488	1,744	1,000

Note:

The stated numbers in the preceding rows are one-dimensional where the given number implies that *only* routes for that address family or type are present. For a given row in the table, the maximum scaling number is ‘x’ IPv4 routes *OR* ‘y’ ipv6 <= 64 routes *OR* ‘z’ ipv6 >64 routes (not a combination of all).

Table 23: VSP 7400 Series

URPF mode	IPv6 mode	VSP 7400 Series		
		IPv4	IPv6	
			Prefix less than 64	Prefix greater than 64
No	No	15,000	7,000	n/a
No	Yes	7,000	3,500	2,000
Yes	No	7,000	3,500	n/a
Yes	Yes	3,000	1,500	1,000

Note:

The stated numbers in the preceding rows are one-dimensional where the given number implies that *only* routes for that address family or type are present. For a given row in the table, the maximum scaling number is ‘x’ IPv4 routes *OR* ‘y’ ipv6 <= 64 routes *OR* ‘z’ ipv6 >64 routes (not a combination of all).

Table 24: XA1400 Series

IPv4 BGP routes (control plane only)	15,488
IPv4 OSFP routes	15,488
IPv4 RIP routes	15,488
IPv4 routes	15,488
IPv4 SPB Shortcut routes	15,488

IP Multicast

Table 25: IP Multicast Maximums

Attribute	Product	Maximum number supported
Combination of VLANs + number of IPv4 senders + IPv6 senders (non-SPBM mode)	VSP 4450 Series	4,059
	VSP 4900 Series	8192
	VSP 7200 Series	8,192
	VSP 7400 Series	8,192
	VSP 8000 Series	8,192
	XA1400 Series	n/a
Combination of Layer 2 VSNs + number of IPv4 senders + number of IPv6 senders (SPBM mode)	VSP 4450 Series	4,059
	VSP 4900 Series	8192

Table continues…

Attribute	Product	Maximum number supported
	VSP 7200 Series	8,192
	VSP 7400 Series	8,192
	VSP 8000 Series	8,192
	XA1400 Series	n/a
IGMP/MLD interfaces (IPv4/IPv6)	VSP 4450 Series	4,059
	VSP 4900 Series	4,059
	VSP 7200 Series	4,059
	VSP 7400 Series	4,059
	VSP 8000 Series	4,059
	XA1400 Series	n/a
PIM interfaces (IPv4/IPv6)	VSP 4450 Series	128 Active
	VSP 4900 Series	128 Active
	VSP 7200 Series	128 Active
	VSP 7400 Series	128 Active
	VSP 8000 Series	128 Active
	XA1400 Series	n/a
PIM Neighbors (IPv4/IPv6) (GRT Only)	VSP 4450 Series	128
	VSP 4900 Series	128
	VSP 7200 Series	128
	VSP 7400 Series	128
	VSP 8000 Series	128
	XA1400 Series	n/a
PIM-SSM static channels (IPv4/IPv6)	VSP 4450 Series	512
	VSP 4900 Series	4,000
	VSP 7200 Series	4,000
	VSP 7400 Series	4,000
	VSP 8000 Series	4,000
	XA1400 Series	n/a
Multicast receivers/IGMP joins (IPv4/IPv6) (per switch)	VSP 4450 Series	1,000
	VSP 4900 Series	6000
	VSP 7200 Series	6,000
	VSP 7400 Series	6,000
	VSP 8000 Series	6,000
	XA1400 Series	n/a

Table continues…

Attribute	Product	Maximum number supported
Total multicast routes (S,G,V) (IPv4/IPv6) (per switch)	VSP 4450 Series	1,000
	VSP 4900 Series	6,000
	VSP 7200 Series	6,000
	VSP 7400 Series	6,000
	VSP 8000 Series	6,000
	XA1400 Series	n/a
Total multicast routes (S,G,V) (IPv4) on an SPB-PIM Gateway configured switch	VSP 4450 Series	1,000
	VSP 4900 Series	3000
	VSP 7200 Series	3,000
	VSP 7400 Series	3,000
	VSP 8000 Series	3,000
	XA1400 Series	n/a
Static multicast routes (S,G,V) (IPv4/IPv6)	VSP 4450 Series	512
	VSP 4900 Series	4000
	VSP 7200 Series	4,000
	VSP 7400 Series	4,000
	VSP 8000 Series	4,000
	XA1400 Series	n/a
Multicast enabled Layer 2 VSN (IPv4)	VSP 4450 Series	1,000
	VSP 4900 Series	2,000
	VSP 7200 Series	2,000
	VSP 7400 Series	2,000
	VSP 8000 Series	2,000
	XA1400 Series	n/a
Multicast enabled Layer 3 VSN (IPv4)	VSP 4450 Series	128 including mgmt VRF and GRT
	VSP 4900 Series	256 including mgmt VRF and GRT
	VSP 7200 Series	256 including mgmt VRF and GRT
	VSP 7400 Series	256 including mgmt VRF and GRT
	VSP 8000 Series	256 including mgmt VRF and GRT
	XA1400 Series	n/a

Table continues…

Attribute	Product	Maximum number supported
SPB-PIM Gateway controller S,Gs (source announcements) with MSDP (IPv4)	VSP 4450 Series	6,000
	VSP 4900 Series	6,000
	VSP 7200 Series	6,000
	VSP 7400 Series	6,000
	VSP 8000 Series	6,000
	XA1400 Series	n/a
SPB-PIM Gateway controllers per SPB fabric (IPv4)	VSP 4450 Series	5
	VSP 4900 Series	n/a
	VSP 7200 Series	5
	VSP 7400 Series	5
	VSP 8000 Series	5
	XA1400 Series	n/a
SPB-PIM Gateway nodes per SPB fabric (IPv4)	VSP 4450 Series	64
	VSP 4900 Series	64
	VSP 7200 Series	64
	VSP 7400 Series	64
	VSP 8000 Series	64
	XA1400 Series	n/a
SPB-PIM Gateway interfaces per BEB (IPv4)	VSP 4450 Series	64
	VSP 4900 Series	64
	VSP 7200 Series	64
	VSP 7400 Series	64
	VSP 8000 Series	64
	XA1400 Series	n/a
PIM neighbors per SPB-PIM Gateway node (IPv4)	VSP 4450 Series	64
	VSP 4900 Series	64
	VSP 7200 Series	64
	VSP 7400 Series	64
	VSP 8000 Series	64
	XA1400 Series	n/a

Distributed Virtual Routing (DvR)

Table 26: DvR Maximums

Attribute	Product	Maximum number supported
Note: • On the DvR leaf, you must enable the VRF scaling boot configuration flag if more than 24 VRFs are required in the DvR domain. • Scaling of the VSP 4450 Series controls the scaling of the DvR domain it is in. For example, if a VSP 4450 Series switch is in a DvR domain with other platforms such as VSP 7200 Series and VSP 8000 Series, the scaling of the entire domain is limited to the scaling of the VSP 4450 Series.		
DvR Virtual IP interfaces	VSP 4450 Series	501 with vIST 502 without vIST
	VSP 4900 Series	501 with vIST 502 without vIST
	VSP 7200 Series	501 with vIST 502 without vIST
	VSP 7400 Series	999 with vIST 1,000 without vIST
	VSP 8000 Series	501 with vIST 502 without vIST
	XA1400 Series	n/a
DvR domains per SPB fabric	VSP 4450 Series	16
	VSP 4900 Series	16
	VSP 7200 Series	16
	VSP 7400 Series	16
	VSP 8000 Series	16
	XA1400 Series	n/a
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.	VSP 4450 Series	n/a
	VSP 4900 Series	n/a
	VSP 7200 Series	8
	VSP 7400 Series	8
	VSP 8000 Series	8
	XA1400 Series	n/a
Leaf nodes per DvR domain	VSP 4450 Series	250

Table continues…

Attribute	Product	Maximum number supported
	VSP 4900 Series	250
	VSP 7200 Series	250
	VSP 7400 Series	250
	VSP 8000 Series	250
	XA1400 Series	n/a
DvR enabled Layer 2 VSNs	VSP 4450 Series	501 with vIST 502 without vIST
	VSP 4900 Series	501 with vIST 502 without vIST
	VSP 7200 Series	501 with vIST 502 without vIST
	VSP 7400 Series	999 with vIST 1,000 without vIST
	VSP 8000 Series	501 with vIST 502 without vIST
	XA1400 Series	n/a
DvR host route scaling per DvR domain (scaling number includes local as well as foreign hosts of the Layer 2 VSN that are members of the domain) If DvR Layer 2 VSNs span DvR domains, and all DvR Controllers have an IP interface on the Layer 2 VSNs, then the DvR host scaling is network-wide, as DvR Controllers will consume as many host routes as there are hosts across all DvR domains.	VSP 4450 Series	6,000
	VSP 4900 Series	32,0000
	VSP 7200 Series	32,000
	VSP 7400 Series	40,000
	VSP 8000 Series	32,000
	XA1400 Series	n/a

VXLAN Gateway

Table 27: VXLAN Gateway Maximums

Attribute	Product	Maximum number supported
MAC addresses in base interworking mode	VSP 4450 Series	n/a
	VSP 4900 Series	n/a
	VSP 7200 Series	112,000

Table continues…

Attribute	Product	Maximum number supported
	VSP 7400 Series	80,000
	VSP 8000 Series	112,000
	XA1400 Series	n/a
MAC addresses in full interworking mode	VSP 4450 Series	n/a
	VSP 4900 Series	n/a
	VSP 7200 Series	74,000
	VSP 7400 Series	50,000
	VSP 8000 Series	74,000
	XA1400 Series	n/a
VNI IDs per node	VSP 4450 Series	n/a
	VSP 4900 Series	n/a
	VSP 7200 Series	2,000
	VSP 7400 Series	2,000
	VSP 8000 Series	VSP 8404C = 4,000 Other VSP 8000 Series platforms = 2,000
	XA1400 Series	n/a
VTEP destinations per node or VTEP	VSP 4450 Series	n/a
	VSP 4900 Series	n/a
	VSP 7200 Series	500
	VSP 7400 Series	500
	VSP 8000 Series	500
	XA1400 Series	n/a

The following table provides maximum numbers for OVSDB protocol support for VXLAN Gateway.

Table 28: OVSDB protocol support for VXLAN Gateway Maximums

Attribute	Product	Maximum number supported
Maximum controllers to which a single VTEP switch can connect	VSP 4450 Series	n/a
	VSP 4900 Series	n/a
	VSP 7200 Series	3
	VSP 7400 Series	3
	VSP 8000 Series	3
	XA1400 Series	n/a

Filters, QoS, and Security

Table 29: Filters, QoS, and Security Maximums

Attribute	Product	Maximum number supported
For more information, see Filter Scaling on page 72.		
Total IPv4 Ingress rules/ACEs (Port/VLAN/ InVSN based, Security/QoS filters)	VSP 4450 Series	1,020
	VSP 4900 Series	1,536
	VSP 7200 Series	766
	VSP 7400 Series	1,536
	VSP 8000 Series	VSP 8404C = 3,070 Other VSP 8000 Series platforms = 766
	XA1400 Series	500
Total IPv4 Egress rules/ACEs (Port based, Security filters)	VSP 4450 Series	255 200 if you enable the ipv6-egress-filter boot configuration flag
	VSP 4900 Series	248
	VSP 7200 Series	248 200 if you enable the ipv6-egress-filter boot configuration flag
	VSP 7400 Series	783 271 if you enable the ipv6-egress-filter boot configuration flag
	VSP 8000 Series	VSP 8404 and VSP 8404C = 251 Other VSP 8000 Series platforms = 252 200 if you enable the ipv6-egress-filter boot configuration flag
	XA1400 Series	500
Total IPv6 Ingress rules/ACEs (Port/VLAN/ InVSN based, Security filters)	VSP 4450 Series	255
	VSP 4900 Series	1024
	VSP 7200 Series	256

Table continues…

Attribute	Product	Maximum number supported
	VSP 7400 Series	767
	VSP 8000 Series	VSP 8404 = 511 VSP 8404C = 2,047 Other VSP 8000 Series platforms = 256
	XA1400 Series	n/a
Total IPv6 egress rules/ACEs (Port based, Security filters)	VSP 4450 Series	256
	VSP 4900 Series	256
	VSP 7200 Series	256
	VSP 7400 Series	511
	VSP 8000 Series	256
	XA1400 Series	n/a
EAP and NEAP (clients per port) **Note:** The total of EAP clients plus NEAP clients per port or per switch cannot exceed 8,192.	VSP 4450 Series	32 for EAP 8,192 for NEAP
	VSP 4900 Series	32 for EAP 8,192 for NEAP
	VSP 7200 Series	32 for EAP 8,192 for NEAP
	VSP 7400 Series	32 for EAP 8,192 for NEAP
	VSP 8000 Series	32 for EAP 8,192 for NEAP
	XA1400 Series	n/a

Filter Scaling

This section provides more details on filter scaling numbers for the VOSS platforms.

VSP 4450 Series

The switch supports the following maximum limits:

- 220 IPv4 ingress ACLs
- 50 IPv4 egress ACLs
- 128 IPv6 ingress ACLs
- 1,020 IPv4 ingress ACEs
- 252 IPv4 egress ACEs

- 255 IPv6 ingress ACEs
- 255 IPv6 egress ACEs

VSP 4900 Series

The switch supports the following maximum limits:

- 512 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
  - 512 ACLs with 1 security ACE each OR
  - 256 ACLs with 1 QoS ACE each OR
  - a combination based on the following rule:
    - `( (num ACLs + num security ACEs) <= 1024) && ((num ACLs + num QoS ACEs) <= 512)`

  This maximum implies a VLAN member count of 1 for inVlan ACLs
- 512 IPv6 ingress ACLs (inPort):
  - 512 ACLs with 1 security ACE each OR
  - a combination based on the following rule:
    - `(num ACLs + num security ACEs) <= 512`
- 124 egress ACLs (outPort only):
  - 124 ACLs with 1 security ACE each (one of these ACLs can have 2 ACEs) OR
  - a combination based on the following rule:
    - `(num ACLs + num ACEs) <= 248`

  This maximum implies a port member count of 1 for outPort ACLs.
- 1534 ingress ACEs:

  Theoretical maximum of 1534 implies 1 ingress ACL with 1023 security ACEs and 511 QoS ACEs
  - Ingress ACEs supported: `(1024 (security) - # of ACLs) + (512 (QoS) - # of ACLs).`

  This maximum also implies a VLAN member count of 1 for an inVlan ACL.
- 247 egress ACEs:

  Theoretical maximum of 247 implies 1 egress ACL with 247 security ACEs
  - Egress ACEs supported: `248 - # of ACLs`.

  This maximum also implies a port member count of 1 for the outPort ACL.

VSP 7400 Series

The switch supports the following maximum limits for *ACL* scaling:

- 512 non-IPv6 ingress ACLs (inPort or inVlan):
  - 256 ACLs with 1 Security ACE each + 256 ACLs with 1 QoS ACE each OR

- 384 ACLs with 1 Security ACE each and/or 1 QoS ACE each OR
- a combination based on the following rule:
  - `num ACLs <= 512 && (num ACLs + num Security ACEs) <= 512 && (num ACLs + num QoS ACEs) <= (512 – X)` where X = num IPv6 ACLs + num IPv6 ACEs

This maximum implies a single port on inPort ACLs, and a single VLAN on inVlan ACLs.

- 384 IPv6 ingress ACLs (inPort):
  - 384 IPv6 ACLs with 1 Security ACE each OR
  - A combination based on the following rule:
    - `num IPv6 ACLs <= 384 && (num IPv6 ACLs + num Security ACEs) <= (768 – X)` where X = num non-IPv6 ACLs + num non-IPv6 QoS ACEs

  This maximum implies a single port on inPort ACLs.
- 254 non-IPv6 egress ACLs (outPort):
  - 254 ACLS with 1 Security ACE each OR
    - A combination based on the following rule:
      - `num ACLs <= 254 && (num ACLs + num Security ACEs) <= 508`

  This maximum implies a single port on outPort ACLs.
- 256 IPv6 Egress ACLs (outPort):
  - 256 ACLS with 1 Security ACE each OR
  - A combination based on the following rule:
    - `num ACLs <= 256 && (num ACLs + num Security ACEs) <= 512`

  This maximum implies a single port on outPort ACLs.

The switch supports the following maximum limits for *ACE* scaling:

- 1,536 non-IPv6 ingress ACEs

  This theoretical maximum implies
  - 1 non-IPv6 ingress ACL with 768 Security ACEs and 768 QoS ACEs
  - no IPv6 ACLs configured
  - a single port on inPort ACLs, and a single VLAN on inVLAN ACLs
- 768 IPv6 ingress ACEs

  This theoretical maximum implies
  - 1 IPv6 ingress ACL with 768 Security ACEs
  - no non-IPv6 ACLs configured
  - a port member count of 1 for inPort ACLs
- 783 non-IPv6 egress ACEs.

This theoretical maximum implies

- 1 egress ACL with 783 Security ACEs
- a port member count of 1 for outPort ACLs
- Non IPv6 egress ACEs supported: `784 - num non-IPv6 egress ACLs`

- 511 IPv6 egress ACEs

  This theoretical maximum implies

  - 1 egress ACL with 511 Security ACEs
  - a port member count of 1 for ourPort ACLs
  - `511 - num IPv6 egress ACLs`

VSP 7200 Series, VSP 8200 Series, and VSP 8404

The switch supports the following maximum limits:

- 256 non-IPv6 ingress ACLs (inPort, inVSN, or inVlan):
  - 256 ACLs with 1 security ACE each OR
  - 128 ACLs with 1 QoS ACE each OR
  - a combination based on the following rule:
    - `( (num ACLs + num security ACEs) <= 512) && ((num ACLs + num QoS ACEs) <= 256)`

  This maximum implies a VLAN member count of 1 for inVlan ACLs
- 256 IPv6 ingress ACLs (inPort,):
  - 256 ACLs with 1 security ACE each OR
  - 256 ACLs with 1 QoS ACE each OR
  - a combination based on the following rule:
    - `(num ACLs + num security ACEs) <= 256`
- 124 egress ACLs (outPort only):
  - 124 ACLs with 1 security ACE each (one of these ACLs can have 2 ACEs)

  This maximum implies a port member count of 1 for outPort ACLs.
- 766 ingress ACEs:

  Theoretical maximum of 766 implies 1 ingress ACL with 511 security ACEs and 255 QoS ACEs

  - Ingress ACEs supported: `(512 (security) - # of ACLs) + (256(QoS) - # of ACLs).`

  This maximum also implies a VLAN member count of 1 for an inVlan ACL.
- 252 egress ACEs:

  Theoretical maximum of 252 implies 1 egress ACL with 252 security ACEs

  - Egress ACEs supported: `253 - # of ACLs`.

  This maximum also implies a port member count of 1 for the outPort ACL.

VSP 8404C

The switch supports a maximum 3,070 non-IPv6 ingress ACEs, 2,047 IPv6 ingress ACEs, and 251 non-IPv6 egress ACEs.

IPv6 ingress and IPv6 egress QoS ACL/Filters are not supported. If you disable an ACL, the ACL state affects the administrative state of all of the ACEs within it.

The switch supports the following maximum limits for *ACL* scaling:

- 1,024 non-IPv6 ingress ACLs (inPort, inVlan, or InVSN):
  - 1,024 ACLs with 1 security ACE each OR
  - a combination based on the following rule:
    - `num of ACLs <= 1,024 AND (num of ACLs + Security ACEs) <= 2,048 AND (num of ACLs + QoS ACEs) <= 1,024`

  This maximum implies a VLAN member count of 1 for inVlan ACLs.
- 1,024 IPv6 ingress ACLs (inPort):
  - 1,024 IPv6 ACLs with 1 security ACE each OR
  - a combination based on the following rule:
    - `num of IPv6 ACLs <= 1,024 AND (num of IPv6 ACLs + Security ACEs) <= 2,048`
- 126 non-IPv6 egress ACLs (outPort):
  - 126 ACLs with 1 Security ACE each OR
  - a combination based on the following rule:
    - `num ACLs <= 126 AND num ACLs + num security ACEs) <= 252`

  This maximum implies a port member counter of 1 for outPort ACLs.

The switch supports the following maximum limits for *ACE* scaling:

- 3,070 non-IPv6 ingress ACEs:

  The theoretical maximum implies the following configuration:
  - 1 non-IPv6 ingress ACL with 2,047 security ACEs and 1,023 QoS ACEs
  - a VLAN member count of 1 for inVlan ACLs
  - Non-IPv6 Ingress ACEs supported: `[2,048(security) - (num of ACLs)] + [1,024(QoS) - (num of ACLs)]`
- 2,047 IPv6 ingress ACEs:

  The theoretical maximum implies the following configuration:
  - 1 IPv6 ingress ACL with 2,047 security ACEs
  - IPv6 Ingress ACEs supported: `[2,048(security) - (num of ACLs)]`
- 251 non-IPv6 egress ACEs:

The theoretical maximum implies the following configuration:

- 1 egress ACL with 251 security ACEs
- a port member count of 1 for outPort ACLs
- Non IPv6 egress ACEs supported: `252 - (num egress ACLs)`

XA1400 Series

The switch supports the following maximum limits:

- 500 IPv4 ingress ACLs
- 500 IPv4 egress ACLs
- 500 IPv4 ingress ACEs
- 500 IPv4 egress ACEs

OAM and Diagnostics

Table 30: OAM and Diagnostics Maximums

Attribute	Product	Maximum number supported
EDM sessions	VSP 4450 Series	5
	VSP 4900 Series	5
	VSP 7200 Series	5
	VSP 7400 Series	5
	VSP 8000 Series	5
	XA1400 Series	5
FTP sessions (IPv4/IPv6)	VSP 4450 Series	8 total (4 for IPv4 and 4 for IPv6)
	VSP 4900 Series	8 total (4 for IPv4 and 4 for IPv6)
	VSP 7200 Series	8 total (4 for IPv4 and 4 for IPv6)
	VSP 7400 Series	8 total (4 for IPv4 and 4 for IPv6)
	VSP 8000 Series	8 total (4 for IPv4 and 4 for IPv6)
	XA1400 Series	4 (IPv4 only)
Rlogin sessions (IPv4/IPv6)	VSP 4450 Series	16 total (8 for IPv4 and 8 for IPv6)

Table continues…

Attribute	Product	Maximum number supported
	VSP 4900 Series	16 total (8 for IPv4 and 8 for IPv6)
	VSP 7200 Series	16 total (8 for IPv4 and 8 for IPv6)
	VSP 7400 Series	16 total (8 for IPv4 and 8 for IPv6)
	VSP 8000 Series	16 total (8 for IPv4 and 8 for IPv6)
	XA1400 Series	8 (IPv4 only)
SSH sessions (IPv4/IPv6)	VSP 4450 Series	8 total (any combination of IPv4 and IPv6)
	VSP 4900 Series	8 total (any combination of IPv4 and IPv6)
	VSP 7200 Series	8 total (any combination of IPv4 and IPv6)
	VSP 7400 Series	8 total (any combination of IPv4 and IPv6)
	VSP 8000 Series	8 total (any combination of IPv4 and IPv6)
	XA1400 Series	8 (IPv4 only)
Telnet sessions (IPv4/IPv6)	VSP 4450 Series	16 total (8 for IPv4 and 8 for IPv6)
	VSP 4900 Series	16 total (8 for IPv4 and 8 for IPv6)
	VSP 7200 Series	16 total (8 for IPv4 and 8 for IPv6)
	VSP 7400 Series	16 total (8 for IPv4 and 8 for IPv6)
	VSP 8000 Series	16 total (8 for IPv4 and 8 for IPv6)
	XA1400 Series	8 (IPv4 only)
TFTP sessions (IPv4/IPv6)	VSP 4450 Series	2 total (any combination of IPv4 and IPv6)
	VSP 4900 Series	2 total (any combination of IPv4 and IPv6)
	VSP 7200 Series	2 total (any combination of IPv4 and IPv6)
	VSP 7400 Series	2 total (any combination of IPv4 and IPv6)

Table continues…

Attribute	Product	Maximum number supported
	VSP 8000 Series	2 total (any combination of IPv4 and IPv6)
	XA1400 Series	n/a
Mirrored ports (source)	VSP 4450 Series	49
	VSP 4900 Series	51 (52 ports per chassis, 48 fixed ports plus up to 4 ports on the VIMs)
	VSP 7200 Series	53 (up to 71 with channelization)
	VSP 7400 Series	31 (up to 125 with channelization) with Advanced Feature Bandwidth Reservation configured in Full Port mode
	VSP 8000 Series	83 (up to 95 with channelization)
	XA1400 Series	7
Mirroring ports (destination)	VSP 4450 Series	4
	VSP 4900 Series	4
	VSP 7200 Series	4
	VSP 7400 Series	4
	VSP 8000 Series	4
	XA1400 Series	4
Fabric RSPAN Port mirror instances per switch (Ingress only)	VSP 4450 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	VSP 4900 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	VSP 7200 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	VSP 7400 Series	Port mirror sessions can be mapped to 24 unique I-

Table continues…

Attribute	Product	Maximum number supported
		SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	VSP 8000 Series	Port mirror sessions can be mapped to 24 unique I-SID offsets for Ingress Mirror. Only one I-SID offset for Egress Mirror.
	XA1400 Series	n/a
Fabric RSPAN Flow mirror instances per switch (Ingress only)	VSP 4450 Series	Filter ACL ACE sessions can be mapped to only 1 mirror I-SID offset.
	VSP 4900 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	VSP 7200 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	VSP 7400 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	VSP 8000 Series	Filter ACL ACE sessions can be mapped to 24 unique I-SID offsets.
	XA1400 Series	n/a
Fabric RSPAN Monitoring I-SIDs (network value)	VSP 4450 Series	1,000 Monitoring I-SIDs across SPB network
	VSP 4900 Series	1,000 Monitoring I-SIDs across SPB network
	VSP 7200 Series	1,000 Monitoring I-SIDs across SPB network
	VSP 7400 Series	1,000 Monitoring I-SIDs across SPB network
	VSP 8000 Series	1,000 Monitoring I-SIDs across SPB network
	XA1400 Series	n/a
sFlow sampling limit	VSP 4450 Series	125 samples per second
	VSP 4900 Series	3,100 samples per second
	VSP 7200 Series	3,100 samples per second
	VSP 7400 Series	9,000 samples per second

Table continues…

Attribute	Product	Maximum number supported
	VSP 8000 Series	3,100 samples per second
	XA1400 Series	n/a
IPFIX flows	VSP 4450 Series	n/a
	VSP 4900 Series	n/a
	VSP 7200 Series	n/a
	VSP 7400 Series	32,767
	VSP 8000 Series	n/a
	XA1400 Series	n/a
Application Telemetry host monitoring - maximum number of monitored hosts **Note:** These resources are shared with the IPv4 Filter Ingress rules/ACEs.	VSP 4450 Series	509 hosts
	VSP 4900 Series	382 hosts
	VSP 7200 Series	382 hosts
	VSP 7400 Series	767 hosts
	VSP 8000 Series	VSP 8404C = 1,534 hosts Other VSP 8000 Series platforms = 382 hosts
	XA1400 Series	n/a

Virtualization Scaling

Note:

The scaling attributes in this section do not apply to the following products:

- VSP 4450 Series
- VSP 4900 Series
- VSP 7200 Series
- VSP 8200 Series
- VSP 8400 Series
- XA1400 Series

Table 31: Virtualization Maximums

Attribute	Product	Maximum number supported
Simultaneous Virtual Machines	VSP 7400 Series	5
CPU cores available to VMs	VSP 7400 Series	6

Table continues…

Attribute	Product	Maximum number supported
Memory available to VMs	VSP 7400 Series	12 GB
Storage available to VMs	VSP 7400 Series	100 GB
Total SRIOV vports available to VMs	VSP 7400 Series	16
Vports available to single VM	VSP 7400 Series	16

Fabric Scaling

This section lists the fabric scaling information.

Table 32: Fabric Maximums

Attribute	Product	Maximum number supported (with and without vIST)
Number of SPB regions	VSP 4450 Series	1
	VSP 4900 Series	1
	VSP 7200 Series	1
	VSP 7400 Series	1
	VSP 8000 Series	1
	XA1400 Series	1
Number of B-VIDs	VSP 4450 Series	2
	VSP 4900 Series	2
	VSP 7200 Series	2
	VSP 7400 Series	2
	VSP 8000 Series	2
	XA1400 Series	2
Maximum number of Physical and Logical (Fabric Extend) NNI interfaces/adjacencies	VSP 4450 Series	255
	VSP 4900 Series	255
	VSP 7200 Series	255
	VSP 7400 Series	255
	VSP 8000 Series	255
	XA1400 Series	255 without IPsec 64 with IPsec
SPBM enabled nodes per area (BEB + BCB)	VSP 4450 Series	550
	VSP 4900 Series	550

Table continues…

Attribute	Product	Maximum number supported (with and without vIST)
	VSP 7200 Series	800
	VSP 7400 Series	2,000
	VSP 8000 Series	800
	XA1400 Series	550
Number of BEBs this node can share services with (Layer 2 VSNs, Layer 3 VSNs, E-Tree, Multicast, Transparent Port UNI). **Note:** vIST clusters are counted as 3 nodes. Each Fabric Extend IS-IS adjacency or VXLAN remote VTEP reduces this number by 1.	VSP 4450 Series	500
	VSP 4900 Series	500
	VSP 7200 Series	500
	VSP 7400 Series	2,000
	VSP 8000 Series	500
	XA1400 Series	n/a
Maximum number of vIST/IST clusters this node can share I-SIDs with	VSP 4450 Series	500
	VSP 4900 Series	330
	VSP 7200 Series	330
	VSP 7400 Series	2,000
	VSP 8000 Series	330
	XA1400 Series	n/a
Layer 2 MAC table size (with SPBM)	VSP 4450 Series	16,000
	VSP 4900 Series	40,000
	VSP 7200 Series	112,000
	VSP 7400 Series	80,000
	VSP 8000 Series	112,000
	XA1400 Series	2,000 for XA1440 4,000 for XA1480
I-SIDs supported	VSP 4450 Series	See Number of I-SIDs supported on page 87
	VSP 4900 Series	See Number of I-SIDs supported on page 87
	VSP 7200 Series	See Number of I-SIDs supported on page 87
	VSP 7400 Series	See Number of I-SIDs supported on page 87
	VSP 8000 Series	See Number of I-SIDs supported on page 87

Table continues…

Attribute	Product	Maximum number supported (with and without vIST)
	XA1400 Series	See Number of I-SIDs supported on page 87
Maximum number of Layer 2 VSNs per switch	VSP 4450 Series	1,000
	VSP 4900 Series	4,059
	VSP 7200 Series	4,059
	VSP 7400 Series	4,000
	VSP 8000 Series	4,059
	XA1400 Series	124
Maximum number of Switched UNI I-SIDs per switch	VSP 4450 Series	See Number of I-SIDs supported on page 87
	VSP 4900 Series	See Number of I-SIDs supported on page 87
	VSP 7200 Series	See Number of I-SIDs supported on page 87
	VSP 7400 Series	See Number of I-SIDs supported on page 87
	VSP 8000 Series	See Number of I-SIDs supported on page 87
	XA1400 Series	n/a
Maximum number of Transparent Port UNIs per switch	VSP 4450 Series	48
	VSP 4900 Series	52
	VSP 7200 Series	54 (up to 72 with channelization)
	VSP 7400 Series	VSP 7432CQ = 32 (up to 125 with channelization) configured in Full Port mode VSP 7400-48Y = 56 configured in Full Port mode
	VSP 8000 Series	84 (up to 96 with channelization)
	XA1400 Series	n/a
Maximum number of E-Tree PVLAN UNIs per switch	VSP 4450 Series	200
	VSP 4900 Series	200
	VSP 7200 Series	200
	VSP 7400 Series	200

Table continues…

Attribute	Product	Maximum number supported (with and without vIST)
	VSP 8000 Series	VSP 8404C = 400 Other VSP 8000 Series platforms = 200
	XA1400 Series	n/a
Maximum number of Layer 3 VSNs per switch See VRF Scaling on page 90.	VSP 4450 Series	128 including mgmt VRF and GRT
	VSP 4900 Series	256 including mgmt VRF and GRT
	VSP 7200 Series	256 including mgmt VRF and GRT
	VSP 7400 Series	256 including mgmt VRF and GRT
	VSP 8000 Series	256 including mgmt VRF and GRT
	XA1400 Series	23
Maximum number of SPB Layer 2 multicast UNI I-SIDs	VSP 4450 Series	See Number of I-SIDs supported on page 87
	VSP 4900 Series	See Number of I-SIDs supported on page 87
	VSP 7200 Series	See Number of I-SIDs supported on page 87
	VSP 7400 Series	See Number of I-SIDs supported on page 87
	VSP 8000 Series	See Number of I-SIDs supported on page 87
	XA1400 Series	n/a
Maximum number of SPB Layer 3 multicast UNI I-SIDs	VSP 4450 Series	Maximum 1,000 for a BEB: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90% of available resources are reached.
	VSP 4900 Series	Maximum 6,000 for a BEB: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90%

Table continues…

Attribute	Product	Maximum number supported (with and without vIST)
		of available resources are reached.
	VSP 7200 Series	Maximum 6,000 for a BEB: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90% of available resources are reached.
	VSP 7400 Series	Maximum 6,000 for a BEB: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90% of available resources are reached.
	VSP 8000 Series	Maximum 6,000 for a BEB: Due to internal resource sharing IP Multicast scaling depends on network topology. Switch will issue warning when 85 and 90% of available resources are reached.
	XA1400 Series	n/a
Maximum number of FA ISID/VLAN assignments per port	VSP 4450 Series	94
	VSP 4900 Series	94
	VSP 7200 Series	94
	VSP 7400 Series	94
	VSP 8000 Series	94
	XA1400 Series	n/a
Maximum number of IP multicast S,Gs when operating as a BCB	VSP 4450 Series	1,000
	VSP 4900 Series	16,000
	VSP 7200 Series	16,000
	VSP 7400 Series	50,000
	VSP 8000 Series	16,000
	XA1400 Series	2,000

Number of I-SIDs Supported for the Number of Configured IS-IS Interfaces and Adjacencies (NNIs)

The number of I-SIDs supported depends on the number of IS-IS interfaces and adjacencies (NNIs) configured.

The following table shows the number of UNI I-SIDs supported per BEB. UNI I-SIDs are used for Layer 2 VSN, Layer 3 VSN, Transparent-UNI, E-Tree, Switched-UNI and S, G for Multicast.

Number of IS-IS interfaces (NNIs)	Product	I-SIDs with vIST configured on the platform	I-SIDs without vIST configured on the platform
4	VSP 4450 Series	1,000	1,000
	VSP 4900 Series	4,000	4,000
	VSP 7200 Series	4,000	4,000
	VSP 7400 Series	4,000	4,000
	VSP 8000 Series	4,000	4,000
	XA1400 Series	n/a	150
6	VSP 4450 Series	1,000	1,000
	VSP 4900 Series	3,500	4,000
	VSP 7200 Series	3,500	4,000
	VSP 7400 Series	3,500	4,000
	VSP 8000 Series	3,500	4,000
	XA1400 Series	n/a	150
10	VSP 4450 Series	650	1,000
	VSP 4900 Series	2,900	4,000
	VSP 7200 Series	2,900	4,000
	VSP 7400 Series	2,900	4,000
	VSP 8000 Series	2,900	4,000
	XA1400 Series	n/a	150
20	VSP 4450 Series	350	700
	VSP 4900 Series	2,000	4,000
	VSP 7200 Series	2,000	4,000
	VSP 7400 Series	2,000	4,000
	VSP 8000 Series	2,000	4,000
	XA1400 Series	n/a	150
48	VSP 4450 Series	n/a	n/a
	VSP 4900 Series	1,000	2,000
	VSP 7200 Series	1,000	2,000
	VSP 7400 Series	1,000	2,000

Table continues…

Number of IS-IS interfaces (NNIs)	Product	I-SIDs with vIST configured on the platform	I-SIDs without vIST configured on the platform
	VSP 8000 Series	1,000	2,000
	XA1400 Series	n/a	150
72	VSP 4450 Series	n/a	n/a
	VSP 4900 Series	750	1,500
	VSP 7200 Series	750	1,500
	VSP 7400 Series	750	1,500
	VSP 8000 Series	750	1,500
	XA1400 Series	n/a	150
100	VSP 4450 Series	n/a	n/a
	VSP 4900 Series	550	1,100
	VSP 7200 Series	550	1,100
	VSP 7400 Series	550	1,100
	VSP 8000 Series	550	1,100
	XA1400 Series	n/a	150
128	VSP 4450 Series	n/a	n/a
	VSP 4900 Series	450	900
	VSP 7200 Series	450	900
	VSP 7400 Series	450	900
	VSP 8000 Series	450	900
	XA1400 Series	n/a	150
250	VSP 4450 Series	n/a	n/a
	VSP 4900 Series	240	480
	VSP 7200 Series	240	480
	VSP 7400 Series	240	480
	VSP 8000 Series	240	480
	XA1400 Series	n/a	150

Interoperability Considerations for IS-IS External Metric

BEBs running VOSS 5.0 can advertise routes into IS-IS with the metric type as external. They can also correctly interpret route advertisements with metric type external received via IS-IS. In an SPB network with a mix of products running different versions of software releases, you must take care to ensure that turning on the ability to use metric-type external does not cause unintended loss of connectivity.

Note the following before turning on IS-IS external metric if the SPB network has switches running a release prior to VOSS 5.0:

- There are no special release or product type implications if the switch does not have IP Shortcuts or Layer 3 VSN enabled. For example, this applies to Layer 2 only BEBs and BCBs.
- There are no special release or product type implications if the Layer 3 VSN in which routes are being advertised with a metric-type of external is not configured on the switch.
- If a switch running a VOSS release that is prior to VOSS 5.0 but VOSS 4.2.1 or later, it will treat all IS-IS routes as having metric-type internal, regardless of the metric-type (internal or external) used by the advertising BEB in its route advertisement.
- Switches running VSP 9000 Series release 4.1.0.0 or later will treat all IS-IS routes as having metric-type internal, regardless of the metric-type (internal or external) used by the advertising BEB in its route advertisement.
- Switches running VOSS releases prior to 4.2.1.0 may not correctly install IS-IS routes in a Layer 3 VSN if any routes advertised with metric-type external are advertised in that Layer 3 VSN by other BEBs in the network. Layer 3 VSNs in which there are no routes with an external metric-type will not be impacted. Similar note applies to the GRT.
- Switches running VSP 9000 Series releases prior to 4.1.0.0 may not correctly install IS-IS routes in a Layer 3 VSN if any routes advertised with metric-type external are advertised in that Layer 3 VSN by other BEBs in the network. Layer 3 VSNs in which there are no routes with an external metric-type will not be impacted. Similar note applies to GRT.
- Switches running any ERS 8800 release may not correctly install IS-IS routes in a Layer 3 VSN if any routes advertised with metric-type external are advertised in that Layer 3 VSN by other BEBs in the network. Layer 3 VSNs in which there are no routes with an external metric-type will not be impacted. Similar note applies to GRT.

Recommendations

This section provides recommendations that affect feature configuration.

Pay special attention to the expected scaling of routes in the network and the number of OSPF neighbors in a single VRF when you select configuration values for the **`isis l1-hellointerval`** and **`isis l1-hello-multiplier`** commands on IS-IS interfaces. The default values for these commands work well for most networks, including those using moderately-scaled routes.

VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, and VSP 8000 Series

The default values work well for 16,000 routes and 64 OSPF neighbors in a single VRF. However, in highly-scaled networks, you may need to configure higher values for these commands.

For example, if the total number of non IS-IS routes on a given BEB exceeds 16,000 in combination with approximately 128 OSPF neighbors in a single VRF, you should configure a value of 12 for **`isis l1-hellomultiplier`**, instead of using the default value of 3.

VSP 4450 Series

If the total number of non IS-IS routes on a given BEB exceeds 25,000 in combination with approximately 60,000 IS-IS routes that the BEB receives from other BEBs in the network, you should configure a value of 12 for **`isis l1-hellomultiplier`**, instead of using the default value of 3.

VRF Scaling

By default, the system reserves VLAN IDs 4060 to 4094 for internal use.

If you enable both the VRF scaling and the SPBM mode boot configuration flags, the system reserves additional VLAN IDs (3500 to 3998) for internal use.

By default, VRF scaling is disabled and SPBM mode is enabled.

Chapter 6: Important Notices

Unless specifically stated otherwise, the notices in this section apply to all VOSS platforms.

100BASE-FX Support on VSP 4000 Series

VSP 4000 Series supports 100BASE-FX transceivers on the VSP 4450GSX or VSP 4850 Series models in SFP ports only. These models do not support 100BASE-FX in SFP+ ports.

AES-GCM SSH Connection with Open SSH

Switch side encryption and authentication type must be set to the AES-GCM-128/256 methods and needs at least one hmac method in the authentication list in addition for the connection to work.

Auto Negotiation Settings

VOSS 4.1 and later software requires the same auto negotiation settings on link partners to avoid incorrect declaration of link status. Mismatched settings can cause the links to stay down as well as unpredictable behavior. Ensure the auto negotiation settings between local ports and their remote link partners match before upgrading software to VOSS 4.1 or later.

dos-chkdsk

If at the end of the **dos-chkdsk WORD<1-99>** command output you see the following choice:

```
    1) Correct
    2) Don't correct
```

Then, you should run the **dos-chkdsk WORD<1-99> repair** command.

IKEv2 Digital Certificate Support with Strong Swan

Strong Swan server must be customized to get IKEv2 Digital Certificate connection between switch and server for RFCs that Strong Swan is compliant and switch is not. This includes SHA256 signing check, IPv6 identifier check and others.

Feature-Based Licensing

The following VOSS platforms support a licensing model that includes Base and Premier licenses:

- VSP 4450 Series
- VSP 4900 Series
- VSP 7200 Series
- VSP 7400 Series
- VSP 8200 Series
- VSP 8400 Series

The Base License, which is included with the purchase of the switch, enables the basic networking capabilities of the device. You can purchase Premier Licenses separately to enable advanced features on the switch.

Premier Licenses enable advanced features not available in the Base License. The following table provides information on the Premier Licenses that the switch supports.

License type	Supported features
Premier License	• DvR Controller • DvR interfaces on more than 24 VRFs/Layer 3 VSNs on Leaf nodes **Note:** DvR Leaf functionality is part of the base software license and the software allows you to create DvR interfaces on Layer 3 VSNs on Leaf nodes. Because a Premier license is required to configure more than 24 VRFs, for deployments where DvR Controllers have more than 24 VRFs configured with DvR, then Leaf nodes only create the first 24 Layer 3 VSNs (VRFs) and no more, unless you install a Premier or Premier with MACsec license. • Extreme Insight • Fabric Connect Layer 3 Virtual Services Networks (VSNs) • Greater than 16 BGP peers • Greater than 24 VRFs

Table continues…

License type	Supported features
	• VXLAN Gateway
Premier with MACsec License	• DvR Controller • DvR interfaces on more than 24 VRFs/Layer 3 VSNs on Leaf nodes **Note:** DvR Leaf functionality is part of the base software license and the software allows you to create DvR interfaces on Layer 3 VSNs on Leaf nodes. Because a Premier license is required to configure more than 24 VRFs, for deployments where DvR Controllers have more than 24 VRFs configured with DvR, then Leaf nodes only create the first 24 Layer 3 VSNs (VRFs) and no more, unless you install a Premier or Premier with MACsec license. • Extreme Insight • Fabric Connect Layer 3 Virtual Services Networks (VSNs) • Greater than 16 BGP peers • Greater than 24 VRFs • IEEE 802.1AE MACsec • VXLAN Gateway

For information about licensing including order codes and how to load a license file, see Administering VOSS.

Subscription Licensing for XA1400 Series

Each XA1400 Series device requires a subscription license.

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

Note:

VOSS Release 8.0.50 or later is required to support subscription licenses generated through the Extreme Networks Support Portal.

The following sections detail the different categories of licenses supported on the XA1400 Series switch.

Factory Default Trial License

A new switch includes a 60-day Factory Default Trial License starting from the time the switch is first booted. You can configure all features (except MACsec), without restrictions and save the configuration. No license file is required.

The system generates warning messages to inform you about the time remaining in the license period. The alerts appear once every 5 days for the first 55 days, and then once daily for the last 5 days. If you reboot the switch after the 60-day period, and a valid software license is not present, the licensed features in the configuration are not loaded. You must install a valid license to enable the licensed features.

Subscription License

All subscription licenses support all VOSS features on the switch, plus software upgrades and technical support services entitlement during the license term. A one, three, or five year subscription license is required for each XA1400 Series device. Three services entitlement tiers of license are available: ExtremeWorks, PartnerWorks, and ExtremeWorks Premier.

A Subscription License is available in two bandwidth tiers of licenses: Small License and Medium License. A Small License enables up to 100 Mbps aggregate throughput Fabric Extend WAN tunneling connectivity, and a Medium License enables up to 500 Mbps aggregate throughput Fabric Extend WAN tunneling connectivity.

License expiry notifications are sent to the console and management station every 30 days until the last 30 days of the subscription. Then every 5 days until the last 9 days of the subscription, and then daily until the Subscription License expires.

Once the Subscription License expires, you get a limited 30 day grace period. When a subscription expires, notification messages are shown as the grace period counts down. Messages are shown on the console and in the alarms database indicating that the license is expired. If the system reboots after a license expiration, the grace period immediately ends and the system does not load or support any saved configurations or software services. License expired messages continue to show on the console and in the alarms database until a valid subscription license is installed.

Important:

The 30 day grace period is lost if the system reboots after a license expires. You must renew your Subscription License to allow the software features to continue to function.

show vlan remote-mac-table Command Output

The output for the **`show vlan remote-mac-table`** command can be different than what appears for the same command on VSP 9000 Series.

Because all MinM packets that originate from the IST switch use the virtual B-MAC as the source BMAC, the remote BEB learns the C-MAC against the virtual B-MAC. Because the remote BEB uses the shortest path to the virtual B-MAC, the remote BEB can show the IST peer as a tunnel in the **`show vlan remote-mac-table`** command output.

Supported Browsers

Use the following browser versions to access Enterprise Device Manager (EDM):

- Microsoft Edge 41+
- Microsoft Internet Explorer 11.0+
- Mozilla Firefox 58.0+
- Google Chrome 64+

System Name Prompt vs. IS-IS Host Name

Beginning with VOSS 6.1.2, 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.

Feature Differences

Extreme Networks has implemented feature parity between the VOSS platforms with a few exceptions. Some features are supported on one platform and not another to maintain compatibility with previous releases. In other cases, the difference is between of the role of the switch in the network.

For information about feature support across all VOSS platforms, see VOSS Feature Support Matrix.

VSP 4000 Series Connecting to an ERS 8800 Interoperability Notes

- For customers running ERS 8800 version 7.1.x:
  - The minimum software release is 7.1.3.1, however the recommended ERS 8800 software release is 7.1.5.4 or later.
  - On switches using 8612 XLRS or 8812XL modules for the links connecting to the VSP 4000 Series, the minimum software version is 7.1.5.4.
  - The “spbm version” on the ERS 8800 must be “802.1aq”.
- For customers running ERS 8800 version 7.2.x:
  - The minimum software release is 7.2.0.2, however the recommended ERS 8800 software release is 7.2.1.1 or later.
  - On switches using 8612 XLRS or 8812XL modules for the links connecting to the VSP 4000 Series switch, the minimum software version is 7.2.1.1.

- Diffserv is enabled in the VSP 4000 Series port settings, and is disabled in the ERS 8800 port settings, by default.

VSP 4000 Series Notes on Combination Ports

When the VSP 4000 Series is reset, the peer connections for all ports, including combination ports 47 and 48 on VSP 4450GTX-HT-PWR+, will transition down. During the reset, the fiber ports remain down, but only the copper ports 47 and 48 come up periodically throughout the reset. The copper ports 47 and 48 come up approximately 15 seconds into the reset, remain up for approximately 60 seconds, and then transition down until the boot sequence is complete and all ports come back up.

The following is an example of the status of the combination ports during reset.

```
CP1 [03/18/70 09:55:35.890] 0x0000c5e7 00300001.238 DYNAMIC SET GlobalRouter HW INFO Link
Down(1/47)
CP1 [03/18/70 09:55:35.903] 0x0000c5e7 00300001.239 DYNAMIC SET GlobalRouter HW INFO Link
Down(1/48)
CP1 [03/18/70 09:55:49.994] 0x0000c5ec 00300001.239 DYNAMIC CLEAR GlobalRouter HW INFO
Link Up(1/48)
CP1 [03/18/70 09:55:50.322] 0x0000c5ec 00300001.238 DYNAMIC CLEAR GlobalRouter HW INFO
Link Up(1/47)
CP1 [03/18/70 09:56:43.131] 0x0000c5e7 00300001.238 DYNAMIC SET GlobalRouter HW INFO Link
Down(1/47)
CP1 [03/18/70 09:56:43.248] 0x0000c5e7 00300001.239 DYNAMIC SET GlobalRouter HW INFO Link
Down(1/48)
```

Cabled Connections for Both Copper and Fiber Ports

The following limitations apply when the combination ports have cabled connections for both the copper and fiber ports.

Do not use the fiber port and do not insert an SFP into the optical module slot in the following situations:

- a copper speed setting of either 10M or 100M is required
- a copper duplex setting of half-duplex is required

Note:

These limitations apply only when auto-negotiation is disabled. To avoid this limitation, use auto-negotiation to determine the speed to 10/100/1000 and to determine the duplex.

The 100M-FX SFP requires auto-negotiation to be disabled. Therefore, auto-negotiation will also be disabled for the copper port. Configure the peer switch to disable auto-negotiation.

Chapter 7: Known Issues and Restrictions

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

Known Issues

This section identifies the known issues in this release.

Known Issues for VOSS 8.1.1

Issue number	Description	Workaround
-	HTTPS connection fails for CA-signed certificate with certificate inadequate type error on FF.	Ensure End-Entity, Intermediate CA and Root CA certificates are all SHA256 based and RSA2048 key signed, and Extended key usage field is set to TLS webserver Auth only for subject and root. For intermediate, it must be set with other required bits to avoid this issue. Add the root, intermediate CAs in the trust store of the browser for accessing the EDM with HTTPS.
-	VRF provisioning is restricted to 127 VRFs on VSP 4000 Series.	None.
VOSS-1265	On the port that is removed from a T-UNI LACP MLT, non T-UNI configuration is blocked as a result of T-UNI consistency checks.	When a port is removed from a T-UNI LACP MLT, the LACP key of the port must be set to default.
VOSS-1278	SLA Mon tests fail (between 2% and 8% failure) between devices when you have too many agents involved with scaled configurations.	This happens only in a scaled scenario with more than seven agents, otherwise the failure does not occur. The acceptable failure percentage is 5%, but you may see failures of up to 8%.
VOSS-1280	The following error message occurs when performing shutdown/no-shutdown commands continuously: `IO1 [05/02/14 06:59:55.178:UTC] 0x0011c525`	None. When this issue occurs, the port in question can go down, then performs a shutdown/no-shutdown of the port to bring it up and resumes operation.

Table continues…

Issue number	Description	Workaround
	`00000000 GlobalRouter COP-SW ERROR vsp4kTxEnable Error changing TX disable for SFP module: 24, code: -8`	
VOSS-1285	CAKs are not cleared after setting the device to factory-default.	None. Currently this is the default behavior and does not affect functionality of the MACsec feature.
VOSS-1288	Shutting down the T1 link from one end of the link does not shut down the link at the remote end. You may experience traffic loss if the remote side of the link is not shut down.	This issue occurs only when a T1 SFP link from one end is shutdown. Enable a dynamic link layer protocol such as LACP or VLACP on both ends to shut the remote end down too. As an alternative, administratively disable both ends of the T1 SFP link to avoid the impact.
VOSS-1289	On a MACsec-enabled port, you can see delayed packets when the MACsec port is kept running for more than 12 hours. This delayed packet counter can also increment when there is complete reordering of packets so that the application might receive a slow response. But in this second case, it is a marginal increase in the packet count, which occurs due to PN mismatch sometimes only during Key expiry, and does not induce any latency.	None.
VOSS-1309	You cannot use EDM to issue **`ping`** or **`traceroute`** commands for IPv6 addresses.	Use CLI to initiate **`ping`** and **`traceroute`** commands.
VOSS-1310	You cannot use EDM to issue **`ping`** or **`traceroute`** commands for IPv4 addresses.	Use CLI to initiate **`ping`** and **`traceroute`** commands.
VOSS-1312	On the 40-gigabit ports, the small metallic fingers that surround the ports are fragile and can bend out of shape during removal and insertion of the transceivers. When the fingers are bent, they prevent the insertion of the QSFP+ transceiver.	Insert the QSFP+ carefully. If the port gets damaged, it needs to be repaired.
VOSS-1335	In an IGMP snoop environment, after dynamically downgrading the IGMP version to version 2 (v2), when you revert back to version 3 (v3), the following is observed: • The multicast traffic does not flow. • The sender entries are not learned on the local sender switch.	Use a v3 interface as querier in a LAN segment that has snoop-enabled v2 and v3 interfaces.

Table continues…

Issue number	Description	Workaround
	• The Indiscard packet count gets incremented on the **`show int gig error`** statistics command.	
VOSS-1340	From EDM, you cannot perform a Layer 2 IP ping for an IPv6 address. EDM displays the following error: `No next Hop address found for ip address provided`	Use the CLI to perform a Layer 2 IP ping.
VOSS-1344	In EDM, you cannot select multiple 40 gigabit ports or a range of ports that includes 40 gigabit ports to graph or edit. You need to select them and edit them individually.	None.
VOSS-1348	In the COM EDM Plugin command, the Layer 2 Traceroute IPv6 does not work properly and displays the error: `No Such Name`.	Use the CLI to initiate the Layer 2 Traceroute for IPv6.
VOSS-1349	On EDM, the port LED for channelized ports only shows the status of sub-port #1, but not the rest of the sub-ports. When you remove sub-port #1, and at least one other sub-port is active and online, the LED color changes to amber, when it should be green because at least one other sub-ports is active and online. The LED only shows the status of sub-port #1.	None.
VOSS-1354	An intermittent link-flap issue can occur in the following circumstance for the copper ports. If you use a crossover cable and disable auto-negotiation, the port operates at 100 Mbps. A link flap issue can occur intermittently and link flap detect will shutdown the port.	Administratively shutdown, and then reenable the port. Use auto-negotiation. Disabling auto-negotiation on these ports is not a recommended configuration.
VOSS-1358	Traffic is forwarded to IGMP v2 SSM group, even after you delete the IGMP SSM-map entry for the group.	If you perform the delete action first, you can recreate the SSM-map record, and then disable the SSM-map record. The disabled SSM-map record causes the receiver to timeout because any subsequent membership reports that arrive and match the disabled SSM-map record are dropped. You can delete the SSM-map record after the receivers time out.
VOSS-1359	The 4 byte AS confederation identifier and peers configuration are not retained across	Reconfigure the 4 byte AS confederation identifier and peers on the device, and reboot.

Table continues…

Issue number	Description	Workaround
	a reboot. This problem occurs when 4 Byte AS is enabled with confederation.	
VOSS-1360	After you enable enhanced secure mode, and log in for the first time, the system prompts you to enter a new password. If you do not meet the minimum password requirements, the following system output message appears: `Password should contain a minimum of 2 upper and lowercase letters, 2 numbers and 2 special characters like !@#$%^*(). Password change aborted. Enter the New password:` The system output message does not display the actual minimum password requirements you need to meet, which are configured on your system. The output message is an example of what the requirements may need to meet. The actual minimum password requirements you need to meet are configured on your system by the administrator.	None.
VOSS-1367	The router ospf entry always appears in the configuration file regardless of whether OSPF is configured. This line does not perform any configuration and has no impact on the running software.	None.
VOSS-1368	When you use Telnet or SSH to connect to the switch, it can take up to 60 seconds for the login prompt to appear. However, this situation is very unlikely to happen, and it does not appear in a standard normal operational network.	Do not provision DNS servers on a switch to avoid this issue altogether.
VOSS-1370	If you configure egress mirroring on NNI ports, you do not see the MAC-in-MAC header on captured packets.	Use an Rx mirror on the other end of the link to see the packets.
VOSS-1371	A large number of IPv6 VRRP VR instances on the same VLAN can cause high CPU utilization.	Do not create more than 10 IPv6 VRRP VRs on a single VLAN.
VOSS-1389	If you disable IPv6 on one RSMLT peer, the switch can intermittently display `COP-SW ERROR` and `RCIP6 ERROR` error messages. This issue has no impact.	None.

Table continues…

Issue number	Description	Workaround
VOSS-1390	If you delete the SPBM configuration and re-configure SPBM using the same nickname but a different IS-IS system ID without rebooting, the switch displays an error message.	Reboot the switch after you delete the SPBM configuration.
VOSS-1403	EDM displays the user name as Admin, even though you login using a different user name.	None.
VOSS-1406	When you re-enable insecure protocols in the CLI SSH secure mode, the switch does not display a warning message.	None.
VOSS-1418	EDM displays the IGMP group entry that is learned on a vIST MLT port as TX-NNI.	Use CLI to view the IGMP group entry learned on a vIST MLT port.
VOSS-1428	When port-lock is enabled on the port and re-authentication on the EAP client fails, the port is removed from the RAIUS-assigned VLAN. This adds the port to the default VLAN and displays an error message. This issue has no impact.	The error message is incorrect and can be ignored.
VOSS-1433	When you manually enable or disable IS-IS on 40 Gbps ports with CR4 direct attach cables (DAC), the port bounces once.	Configure IS-IS during the maintenance period. Bring the port down, configure the port and then bring the port up.
VOSS-1438	In a rare scenario in Simplified vIST configuration when vIST state is toggled immediately followed by vIST MLT ports are toggled, one of the MLT ports will go into blocking state resulting in failure to process data packets hashing to that link.	Before enabling vIST state ensure all vIST MLT ports are shut and re-enabled after vIST is enabled on the DUT.
VOSS-1440 VOSS-1441	When you configure a scaled Layer 3 VSN (24 Layer 3 VSN instances), route leaking from GRT to VRF on the local DUT does not happen. The switch displays an incorrect error message: `Only 24 Layer 3 VSNs can be configured.`	None.
VOSS-1459 VOSS-1463 VOSS-1471	When you use Fabric Extend over IP (FE-IP) and Fabric Extend over Layer 2 VLAN (FE-VID) solution, if you change the ingress and egress .1p map, packets may not follow correct internal QoS queues for FE tunnel to FE tunnel, or FE tunnel to regular NNI traffic.	Do not change the default ingress and egress .1p maps when using Fabric Extend. With default ingress and egress .1p maps, packets follow the correct internal QoS when using the Fabric Extend feature.
VOSS-1473	If the I-SID associated with a Switched UNI or Fabric Attach port does not have a platform VLAN association and you disable	None.

Table continues…

Issue number	Description	Workaround
	Layer 2 Trusted, then the non IP traffic coming from that port does not take the port QoS and still uses the .1p priority in the packet.	
VOSS-1530	If you improperly close an SSH session, the session structure information does not clear and the client can stop functioning.	Disable and enable SSH.
VOSS-1560	If you apply an ipv6-out-route-map on a BGP peer to filter a particular IPv6 prefix range with a match network condition, it does not filter the full prefix range.	Configure the incoming policy to filter incoming advertised routes on BGP+ peers.
VOSS-1584	The **`show debug-file all`** command is missing.	None.
VOSS-1585	The system does not generate a log message, either in the log file or on screen, when you run the **`flight-recorder`** command.	None.
VOSS-1608	If you use an ERS 4850 FA Proxy with a VOSS FA Server, a mismatch can exist in the show output for tagged management traffic. The ERS device always sends traffic as tagged. The VOSS FA Server can send both tagged and untagged. For untagged, the VOSS FA Server sends VLAN ID 4095 in the management VLAN field of the FA element TLV. The ERS device does not recognize this VLAN ID and so still reports the traffic as tagged.	There is no functional impact.
VOSS-16271	If you have an untagged S-UNI associated with platform VLAN configured on port and you enable Endpoint Tracking on the same port, untagged traffic will get dropped.	Remove untag S-UNI and reconfigure it.
VOSS-1706	EAPOL: Untagged traffic is not honoring the port QOS for Layer 2 trusted/ Layer 3 untrusted. This issue is only seen on EAPOL-enabled ports.	None.
VOSS-2014	IPV6 MLD Group is learned for Link-Local Scope Multicast Addresses. This displays additional entries in the Multicast routing tables.	None.
VOSS-2033	The following error messages appear when you use the **`shutdown`** and **`no shutdown`** commands on the MLT interface with ECMP and BGP+ enabled:	Disable the alternate path.

Table continues…

Issue number	Description	Workaround

	`CP1 [01/23/16 11:10:16.474:UTC] 0x00108628 00000000 GlobalRouter RCIP6 ERROR rcIpReplaceRouteNotifyIpv6:FAIL ReplaceTunnelRec conn_id 2`
	`CP1 [12/09/15 12:27:02.203:UTC] 0x00108649 00000000 GlobalRouter RCIP6 ERROR ifyRpcOutDelFibEntry: del FIB of Ipv6Route failed with 0: ipv6addr: 201:6:604:0:0:0:0:0, mask: 96, nh: 0:0:0:0:0:0:0:0 cid 6657 owner BGP`
	`CP1 [12/09/15 12:20:30.302:UTC] 0x00108649 00000000 GlobalRouter RCIP6 ERROR ifyRpcOutDelFibEntry: del FIB of Ipv6Route failed with 0: ipv6addr: 210:6:782:0:0:0:0:0, mask: 96, nh: fe80:0:0:0:b2ad:aaff:fe55:5088 cid 2361 owner OSPF`	

VOSS-2036	IPsec statistics for the management interface do not increment for inESPFailures or InAHFailures.	None.
VOSS-2117	If you configure static IGMP receivers on an IGMPv3 interface and a dynamic join and leave are received on that device from the same destination VLAN or egress point, the device stops forwarding traffic to the static receiver group after the dynamic leave is processed on the device. The end result is that the IGMP static groups still exist on the device but traffic is not forwarded.	Disable and re-enable IGMP Snooping on the interface.
VOSS-2128	EAP Security and Authentication EDM tabs display additional information with internal values populated, which is not useful for the end user.	There is no functional impact. Ignore the additional information in EDM. Use the CLI command **`show eapol port interface`** to see port status.
VOSS-2207	You cannot configure an SMTP server hostname that begins with a digit. The system displays the following error: `Error: Invalid IP Address or Hostname for SMTP server`	None.

Table continues…

Issue number	Description	Workaround
VOSS-2208	While performing CFM Layer 2 traceroute between two BEBs via a transit BCB, the transit BCB hop is not seen, if the transit BCB has ISIS adjacencies over FE l3core with both source BEB and destination BEB.	None.
VOSS-2253	Trace level command does not list module IDs when '**?**' is used.	To get the list of all module IDs, type **`trace level,`** and then press **`Enter.`**
VOSS-2270	The packet internal CoS is derived incorrectly for packets sourced from a brouter port when the CoS should be derived from the port level QoS. The following list identifies scenarios that derive the internal CoS from the port QoS: • Untagged non-IP packet • Untagged IP packet, and the source port is Layer 3 untrusted • Tagged non-IP packet and the source port is Layer 2 untrusted • Tagged IP packet and the source port is Layer 3 untrusted and Layer 2 untrusted	Use the port default QoS configuration for the brouter port. The port default configuration is Layer 2 trusted and Layer 3 trusted, and under this configuration, only the first scenario in the list is still an issue. The other scenarios do not occur.
VOSS-2279	When an IPv6 neighbor device boots, the following error message occurs in the peer device console: `GlobalRouter COP-SW ERROR ercdProcIpv6RouteMsg: Failed to Delete IPV6 Record - Ip: fe80:0:0:8dc:b2ad:aaff:fe55:1b91, NextHop:0:0:0:0:0:0:0:0, mask: 128`	There is no functional impact. Port **`shutdown`** and **`no shutdown`** commands, which recovers the traffic, works even when the switch is in an error state.
VOSS-2285	When on BEB, continuously pinging IPv6 neighbor address using CLI command **`ping -s`**, ping packets do not drop, but instead return no answer messages.	Restart the ping. Avoid intensive CPU processing.
VOSS-2333	Layer 2 ping to Virtual BMAC (VBMAC) fails, if the VBMAC is reachable via Layer 2 core.	None.
VOSS-2411	On a VSP 4450GSX-DC, the https-port info is not displayed or saved into the configuration.	None.
VOSS-2418	When you configure and enable the SLA Mon agent, the SLA Mon server is able to	None.

Table continues…

Issue number	Description	Workaround
	discover it but the agent registration on the switch does not occur.	
VOSS-2422	When a BGP Neighbor times out, the following error message occurs: `CP1 [03/11/16 13:43:39.084:EST] 0x000b45f2 00000000 GlobalRouter SW ERROR ip_rtdeleteVrf: orec is NULL!`	There is no functional impact. Ignore the error message.
VOSS-2859	You cannot modify the port membership on a protocol-based VLAN using EDM, after it has been created.	Use CLI to provision the port membership on the protocol-based VLAN or delete the protocol-based VLAN, and then re-create it with the correct port member setting.
VOSS-3393	When the SLA Mon agent IP is created on a CLIP interface, the switch provides the CLIP-id as the agent MAC.	There is no functional impact. Use different CLIP IDs to differentiate the SLA Mon agents from the SLA Mon server.
VOSS-4255	If you run IP traceroute from one end host to another end host with a DvR Leaf in between, an intermediate hop will appear as not responding because the Leaf does not have an IP interface to respond. The IP traceroute to the end host will still work.	None.
VOSS-4728	If you remove and recreate an IS-IS instance on an NNI port with autonegotiation enabled in addition to vIST and R/SMLT enabled, it is possible that the NNI port will briefly become operationally down but does recover quickly. This operational change can lead to a brief traffic loss and possible reconvergence if non-ISIS protocols like OSPF or BGP are also on the NNI port.	If you need to remove and recreate an IS-IS instance on an autonegotiation enabled NNI port that also has non-ISIS traffic, do so during a maintenance window to minimize possible impact to other non-ISIS traffic.
VOSS-4840	If you run the **`show fulltech`** command in an SSH session, do not disable SSH on the system. Doing so can block the SSH session.	None.
VOSS-4912	The VSP 4000 Series does not advertise an LLDP Management TLV.	None.
VOSS-5130	Disabling and immediately enabling IS-IS results in the following log message: `PLSBFIB ERROR: /vob/cb/nd_protocols/plsb/lib/plsbFib.cpp(line 1558) unregisterLocalInfo() local`	There is no functional impact. Ignore the error message.

Table continues…

Issue number	Description	Workaround
	`entry does not exist. key(0xfda010000fffa40)`	
VOSS-5159 & VOSS-5160	If you use a CLIP address as the management IP address, the switch sends out 127.1.0.1 as the source IP address in both SMTP packets and TACACS+ packets.	None.
VOSS-5173	A device on a DvR VLAN cannot authenticate using RADIUS if the RADIUS server is on a DvR VLAN on a DvR Leaf using an in-band management IP address.	Place the RADIUS server in a non-DvR VLAN off a DvR Leaf or DvR Controller.
VOSS-5331	When you enable FHS ND inspection on a VLAN, and an IPv6 interface exists on the same VLAN, the IPv6 host client does not receive a ping response from the VLAN.	None.
VOSS-5603	In a scaled DvR environment (scaled DvR VLANs), you may see a higher CPU utilization while deleting a DvR leaf node from the DvR domain (no dvr leaf). The CPU utilization stays higher for several minutes on that node only and then returns to normal after deleting all the internal VLANs on the leaf node.	It is recommended to use a maintenance window when removing leaf(s) from a DvR domain.
VOSS-5627	The system does not currently restrict the number of VLANs on which you can simultaneously configure NLB and Directed Broadcast, resulting in resource hogging.	Ensure that you configure NLB and Directed Broadcast on not more that 100 VLANs simultaneously, assuming one NLB cluster for each VLAN. Also, ensure that you configure NLB on a VLAN first, and then Directed Broadcast, so as to not exhaust the NLB and Directed Broadcast shared resources. The shared resources are NLB interfaces and VLANs with Directed Broadcast enabled. The permissible limit for the shared resources is 200.
VOSS-6189	When you connect to EDM using HTTPS in Microsoft Edge or Mozilla FireFox, the configured values for the RADIUS KeepAliveTimer and CFM SBM MepId do not appear.	Use Internet Explorer when using an HTTPS connection.
VOSS-6822	If the IPsec/IKE software used in the Radius server side is strongSwan, there is a compatibility issue between VOSS and	None.

Table continues…

Issue number	Description	Workaround
	strongSwan in terms of IPv6 Digicert (IKEv1/v2) authentication.	
VOSS-6928	On VSP 8000 Series platforms, IPv4 Filters with redirect next hop action do not forward when a default route is not present or a VLAN common to ingress VLAN of the filtered packet is not present.	Configure a default route if possible.
VOSS-7006	SMLT MACs are not synced correctly when you create a new VLAN on one of the vIST peers.	After you create a VLAN, enter the following command: **vlan mac-address-entry <vlan id> re-sync**
VOSS-7139	DHCPv6 Snooping is not working in an SPB network as the DHCPv6 Snooping entries are not being displayed.	Administrator should add manual entries.
VOSS-7457	The switch can experience an intermittent traffic loss after you disable a Fabric Extend tunnel.	Bounce the tunnel between the devices.
VOSS-7472	EDM shows incorrect guidance for ACL TCP flag mask. EDM reports `0…63` as hexadecimal. CLI correctly shows `<0-0x3F | 0-63> Mask value <Hex | Decimal>`. This is a display issue only with no functional impact.	Use CLI to see the correct unit values.
VOSS-7495	The VSP 4000 Series CLI Help text shows an incorrect port for **`boot config flags linerate-directed-broadcast`**. The Help text shows 1/48. The correct port is 1/46.	None
VOSS-8424	A fragmented ping from an external device to a switch when the VLAN IP interface is tied to a non-default VRF fails.	None.
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-9206	Interface statistics InDiscard counter in **`show interfaces gigabitEthernet error`** output does not increment consistently when IPv6 packets are dropped when uRPF checks fail. This issue applies only to VSP 4000 Series.	None.
VOSS-9516	When you connect to EDM using HTTPS, you can see multiple `SSL negotiation with client successful` messages during your EDM session. This message	None.

Table continues…

Issue number	Description	Workaround
	appears each time a successful SSL_Handshake occurs between the web browser and the web server. The log file may not show as many messages as the console and the timing between messages can be different because logging does not occur in real time.	
VOSS-9589	Dynamic Nickname Assignment is not supported over Fabric Extend tunnels.	None.
VOSS-9621	For VOSS products, 1G Copper Pluggable auto-negotiation is always enabled after a reboot, despite configuration settings.	If you do not want to use auto-negotation, disable it after the reboot.
VOSS-9917	The log message `INFO Switch Externally Rebooted with CoreDump` does not consistently appear on the console port before reboot when you select the **softResetCoreDump** option from EDM.	None.
VOSS-9921	Bootup redirection timeout is longer than the UNI port (SMLT) unlock timer. If both vIST nodes boot together in factory default configuration fabric mode or without a nickname, the vIST ports will not enable for up to 4 minutes. During the delay the nickname server is unreachable and vIST is not online.	None.
VOSS-10380	If you enable and configure IPv6 Source Guard and EAPoL on a port, and create and configure a Guest VLAN on the same port without DHCP Snooping and ND-inspection, no error is shown. The port is not added to the Guest VLAN.	None.
VOSS-10381	If you enable and configure IPv6 Source Guard and EAPoL MHSA on a port, and create and configure RAVs for Non-EAP clients on the same port without DHCP Snooping and ND-inspection, no error is shown. The client displays as authenticated into RAV, even when port is not a member of RAV.	None.
VOSS-10412	Removal of the QSFP+ to SFP+ adapter with a 10G pluggable is not detected on the VSP 8404 and VSP 8404C when in non channelized mode.	The QSFP+ to SFP+ adapter and detection works only on ports with channelization enabled.

Table continues…

Issue number	Description	Workaround
VOSS-10574	IS-IS sys-name output is not truncated for **`show isis spbm nick-name`** or **`show ip route`** commands. If a long character sys-name is in use, the full sys-name display can cause misalignment of the output columns.	None.
VOSS-10815	DvR over SMLT: Traffic is lost at failover on SMLT towards EXOS switches. DvR hosts are directly connected to the DvR controllers vIST pair on SMLT LAG and switched-UNIs are dynamically added using Fabric Attach. Only occurs when the access SMLT is LACP MLT and all the ports in the MLT are down. When all ports in the MLT down and an ARP request is received over an NNI link, there is no physical port that can be associated with the ARP request. The ARP entry is learned against NNI link, and MAC syncs from vIST peer or from a non-vIST peer when bouncing vIST.	None.
VOSS-10891	DvR leaf vIST: Wrong rarSmltCheckSmltPeerMac MLT warning displays when the peer vIST MAC address is learned from local	None. rarSmltCheckSmltPeerMac MLT warning has no functional impact. You can ignore the error message.
VOSS-11895	In a vIST SMLT environment where streams are both local and remote, if source and receiver port links are removed and reinserted several times, eventually traffic will not be forwarded to local single-homed receivers on one peer if the traffic is ingressing from the vIST peer over the NNI link. If the stream ingresses locally, it is received by the local UNI receivers.	Disable and renable Fabric Multicast (**`spbm <1-100> multicast enable`**) on the source VLAN to allow the streams to be deleted and come back in properly.
VOSS-11943	This release does not support per-port configuration of Application Telemetry. Because the feature is enabled globally and VSP 7432CQ supports 32 100 Gbps ports, an undesirable condition may be encountered when an exceeded amount of Application Telemetry mirrored packets are sent to the collector.	None.
VOSS-12330	When accesssing the on-switch RESTCONF API documentation in a web	Ensure you include the trailing slash (/) in the URL: `http(s)://<ip-address>:8080/apps/restconfdoc/`. For more

Table continues…

Issue number	Description	Workaround
	browser, the page does not render correctly.	information, see Configuring User Interfaces and Operating Systems for VOSS.
VOSS-12405	To reach a VM, all front panel traffic must travel through an Insight port, which is a 10 Gbps port. If front panel port traffic is over 10 Gbps, this situation represents an oversubscription on the Insight port and some of the packets will be dropped. As a result, Extreme Management Center can lose connectivity to the Analytics engine if Application Telemetry is enabled.	None.
VOSS-13159	The ixgbevf Ethernet device driver within the TPVM does not correctly handle the interface MTU setting. Specifically, if you configure the interface in SR-IOV mode, packets larger than the MTU size are allowed.	To avoid this problem, configure the desired MTU size on both the relevant front-panel port and Insight port from VOSS.
VOSS-13463	Out port statistics for MLT port interfaces are not accurate.	Use the command `show io nic-counters` to display detailed port stats and error info on XA1400 Series.
VOSS-13667	An intermittent issue in SMLT environments, where ARPs or IPv6 neighbors are resolved with delay can cause a transient traffic loss for the affected IPv6 nieghbors. The situation auto-corrects.	None.
VOSS-13680	Interface error statistics display is inaccurate in certain scenarios.	Use the command `show io nic-counters` to display detailed port stats and error info on XA1400 Series.
VOSS-13681	QoS: show qos cosq-stats cpu-port command output is not supported.	Use the command `show io cpu-cosq-counters` to display detailed cosq-stats on XA1400 Series.
VOSS-13693	QoS: Traffic can egress out of the queue at a different ratio than the default configuration. After the guaranteed traffic rate is served to all egress port queues, any excess bandwidth is shared equally to all queues instead of distributing on weight assigned to each queue.	None.
VOSS-13702	Do not use the ACE actions of deny and mirror-to-isid together on VSP 7400 Series.	None.

Table continues…

Issue number	Description	Workaround
VOSS-13717 VOSS-14393 VOSS-14972	Link on remote side doesn't go down after admin shut on XA1400 while using 10G DAC or a 4x10 - 40 G breakout DAC. On the XA1400 side link goes down but Link LED shows as up. Both 10G and 4x10G DAC are not fully supported because of this issue	None for DAC and breakout cables. Because of this issue, the following optical transceivers are not supported: • AA1404036-E6 • AA1404042-E6 • C9799X4-5M
VOSS-13789	A link between a 25 Gbps port on VSP 7400-48Y and a channelized 100 Gbps port on VSP 8600 Series is not established using a 100 Gb to 25 Gb direct-attach copper breakout cable.	You must disable Auto-Negotiation and FEC on the 25 Gbps port.
VOSS-13794	You cannot use SFTP to transfer files larger than 2 GB to a VSP switch.	Use SCP.
VOSS-13938	You can configure LLDP-MED on an FA-enabled port, and **`show lldp`** commands show the configuration as applied but the information is not advertized and it does not appear in **`show running-config`** output nor in config.cfg if you save the configuration	None.
VOSS-13947	After you enable MSTP-Fabric Connect Multi Homing (**`spbm 1 stp-multi-homing enable`**), you cannot view the configuration, role, or statistics for the STP virtual port.	None.
VOSS-13948	After you enable MSTP-Fabric Connect Multi Homing (**`spbm 1 stp-multi-homing enable`**), MSTP resiliency times are 30 to 40 seconds because the internal SPB-STP port is not fast-aging remote CMAC entries after a topology change occurs.	None.
VOSS-13974	When an 8408QQ ESM has more than two channelized ports and is rebooted, the MKA MACsec sessions on the other cards in the same box may toggle. This issue is not seen if one or two ports are channelized on the same card.	None.
VOSS-14150	CLI remote console might stop wrapping text after some usage.	Reset the CLI window or open a new remote console window.
VOSS-14391	On an VSP 8404C switch using an 8424XT ESM, on a port with MACsec connectivity, if you set Auto-Negotiation advertisements	To avoid this issue, set the Auto-Negotiation advertisements directly to 10000-full.

Table continues…

Issue number	Description	Workaround
	to 1000-full, and then subsequently set the advertisement to 10000-full, the link will not come up.	If you have experienced the issue, shut the port down and bring it back up.
VOSS-14494	Layer 2 VSN and Layer 3 VSN UNI to NNI traffic between two Backbone Edge Bridges does not hash to different ports of a MLT network-to-network interface. MLT hashing for XA1400 devices occurs after the mac-in-mac encapsulation is done. The hash keys used are the Backbone destination and Backbone source MAC addresses (BMAC DA and BMAC SA) in the Mac-in-Mac header. Even for the Transit BCB case on XA 1400 devices for NNI to NNI traffic, the MLT hash keys used are the Backbone destination and Backbone source MAC addresses (BMAC DA and BMAC SA) in the Mac-in-Mac header.	None.
VOSS-14515	Console output errors and warnings are shown during an XA1400 Series reboot, such as: • error: no such device: ((hd0,gpt1)/EFI/BOOT)/EFI/BOOT/grub.cfg. error: file `/EFI/BOOT/grubenv' not found • error: no suitable video mode found. • vfio-pci 0000:05:00.0: Invalid PCI ROM header signature: expecting 0xaa55, got 0xbeef • [0.727012] ACPI: No IRQ available for PCI Interrupt Link [LNKS]. Try pci=noacpi or acpi=off • exportfs: can't open /etc/exports for reading • KCORE: WARNING can't find /boot/b/uImage-gemini.bin. No kexec kernel will be configured.	None. The errors or warnings are host OS or guest OS related with no functional impact and can be ignored.
VOSS-14590	ISIS logical-interface displays the same egress port for different tunnels when the underlay reachability is from different port interfaces.	None.
VOSS-14592	Operation down log messages display on console for an already shutdown port. Log	None.

Table continues…

Issue number	Description	Workaround
	messages are printed whenever a shut CLI command executes even if the port was previously shutdown.	
VOSS-14597	Ping (originated from local CP) fails for jumbo frames on Layer 3 VSN interface.	None.
VOSS-14616	Seeing Queue buffer usage logs when changing the logical interface source IP with 64 tunnels. When changing the source IP with 64 tunnels, seeing "GlobalRouter CPU INFO CPP: 60 percent of fbufs are in use: 0 in Tx queue,1843 in RxQueue0 0 in RxQueue1 0 in RxQueue2 0 in RxQueue3 0 in RxQueue4 0 in RxQueue5 0 in RxQueue6 0 in RxQueue7 ".	None.
VOSS-14639	Packets ingressing on the front panel ports with source mac as switch's local VLAN MAC are forwarded instead of being dropped.	You can create an ACL filter rule to match packets based on the source MAC and drop the packets when this condition is encountered.
VOSS-14647	For unknown unicast, broadcast and multicast traffic packets of 144 bytes or less, there will be egress drops on the VSP 4900 Series in the case of congestion/ oversubscription on the egress.	None. Note that this issue is not seen with packets over 144 bytes.
VOSS-14656	Console output "ErrLog: Error Level=2 [(null)] seen during OpenVas testing. No functional impact.	None.
VOSS-14805 VOSS-15305	The following transceivers are not supported on XA1400 Series switches: • 10 Gb Bidirectional 40 km SFP+ Module (10GB-BX40-D and 10GBBX40-U) • 1000BASE-BX10 Bidirectional 10 km DDI SFP Modules (AA1419069-E6 and AA1419070-E6)	Use only supported transceivers.
VOSS-15079	The Extreme Networks 10 meter SFP+ passive copper DAC (Model Number 10307) does not function on ports 2/3 and 2/4 of the VIM5-4X.	Use the Extreme Networks SFP+ active optical DAC (Model Number AA1403018-E6) with the VIM5-4X.
VOSS-15112	BFD sessions associated with static routes may flap once before remaining up, when shutting down and bringing back up a BFD peer port.	None. Ignore the extra BFD session flap.

Table continues…

Issue number	Description	Workaround
VOSS-15313	On an VSP 8404C switch using an 8424XT ESM, on a link with MACsec connectivity on both ends, and Auto-Negotiation advertisements set to 10000-full, the link will not come back up if the ESM is hot-swapped or the slot is reset.	To avoid this issue, disable MACsec prior to the hot swap or reset, and then re-enable. If you have experienced the issue, shut either one of the link ports down and bring it back up.
VOSS-15350	In SMLT configurations with LACP, when shutting down and bringing back up SMLT links, there may be momentary traffic loss for VLANs with VRRP configured.	Enable VRRP backup-master for all ports in VRRP-enabled VLANs.
VOSS-15391	An SNMP walk on the **rcIgmpSnoopTraceTable** table will fail with an **OID not increasing** error. CLI and EDM are unaffected by this issue.	None.
VOSS-15463	XA1440 and XA1480 switches may experience intermittent Link Up and Link Down transitions on the 10/100/1000BASE-T Ethernet ports upon booting.	No workaround, but there is no functional impact.
VOSS-15525	If you enable the **filter-untagged-frame** option before you enable the **untag-port-default-vlan** option, the default VLAN untagging for the port does not function correctly.	Disable the **filter-untagged-frame** option, disable and then re-enable the **untag-port-default-vlan** option, and then re-enable the **filter-untagged-frame** option.
VOSS-15541	You may experience temporary traffic loss when shutting down an LACP SMLT port (and therefore causing the local SMLT to go down), in a network with scaled Multicast traffic over an SPB cloud, while the datapath processes all dpm letter messages during LCAP recovery. This slow LACP recovery situation is only seen with scaled Multicast traffic over an SPB cloud.	Use static MLTs.
VOSS-15605	When you delete the VLAN IDs from the assigned I-SID of two vIST peers, the second VLAN ID deletion triggers log report **0x0013851e** from the first peer, indicating that a Layer 3 MAC address deletion has failed.	No workaround, but there is no functional impact—the MAC address was deleted when the VLAN:ISID association was deleted.
VOSS-15720	During key refresh events for MKA dynamic SAKs, you may experience 3-4 msec packet drops, depending on the interface line rate, incoming packet size,	None. Packet loss during key refresh events is very minimal (approximately 3-4 msec). If applications cannot tolerate this

Table continues…

Issue number	Description	Workaround
	and incoming packet rate. Under average conditions, on a 1 Gbps port, there may be an average packet loss of 20 msec over a 24 hour period, while on a 10 Gbps port, there may be an average packet loss of 160 msec over a 24 hour period.	amount of traffic loss also, it is advisable to use Static SAKs instead of dynamic SAKs.
VOSS-15812	L3VSN IPv4 BGP (and static) routes having their next-hops resolved via IS-IS routes may result in traffic loss.	Choose the following workarounds, based on your deployment and needs: • Use static routes to reach the loopbacks used as BGP peers, (static routes having better preference than IS-IS); use static routes with next-hops reachable on the UNI side (L2VSN). • Use OSPF to reach the loopbacks used as BGP peers, but take care to ensure that the OSPF route towards the BGP peer is chosen as the “best route” (as IS-IS has a better preference than OSPF). There are several ways to accomplish this—either don’t redistribute that route in IS-IS if it is not needed, or control the redistribution with a route-map, etc. • Have BGP peers reachable directly via a C-VLAN; do not use loopback interfaces as BGP peer addresses. • If none of the above workaround scenarios are suitable for your deployment, do not use internal Border Gateway Protocol (iBGP) peering.

Restrictions and Expected Behaviors

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

For Port Mirroring considerations and restrictions, see Troubleshooting VOSS.

General Restrictions and Expected Behaviors

The following table provides a description of the restriction or behavior.

Issue number	Description	Workaround
—	If you access the Extreme Insight virtual machine using `virtual-`	None.

Table continues…

Issue number	Description	Workaround
	`service tpvm console` and use the Nano text editor inside the console access, the command **`^o<cr>`** does not write the file to disk.	
VOSS-7	Even when you change the LLDP mode of an interface from CDP to LLDP, if the remote side sends CDP packets, the switch accepts them and refreshes the existing CDP neighbor entry.	Disable LLDP on the interface first, and then disable CDP and re-enable LLDP.
VOSS-687	EDM and CLI show different local preference values for a BGP IPv6 route. EDM displays path attributes as received and stored in the BGP subsystem. If the attribute is from an eBGP peer, the local preference appears as zero. CLI displays path attributes associated with the route entry, which can be modified by a policy. If a route policy is not configured, the local preference shows the default value of 100.	None.
VOSS-1954	After you log in to EDM, if you try to refresh the page by clicking on the refresh button in the browser toolbar, it will redirect to a blank page. This issue happens only for the very first attempt and only in Firefox.	To refresh the page and avoid this issue, use the EDM refresh button instead of the browser refresh button. If you do encounter this issue, place your cursor in the address bar of the browser, and press **Enter**. This will return you to the EDM home page.
VOSS-2166	The IPsec security association (SA) configuration has a NULL Encryption option under the **Encrpt-algo** parameter. Currently, you must fill the **encrptKey** and **keyLength** sub-parameters to set this option; however, these values are not used for actual IPsec processing as it is a NULL encryption option. The NULL option is required to interoperate with other vendors whose IPsec solution only supports that mode for encryption.	There is no functional impact due to this configuration and it only leads to an unnecessary configuration step. No workaround required.

Table continues…

Issue number	Description	Workaround
VOSS-2185	MAC move of the client to the new port does not automatically happen when you move a Non-EAP client authenticated on a specific port to another EAPoL or Non-EAP enabled port.	As a workaround, perform one of the following tasks: • Clear the non-EAP session on the port that the client is first authenticated on, before you move the client to another port. • Create a VLAN on the switch with the same VLAN ID as that dynamically assigned by the RADIUS server during client authentication. Use the command **`vlan create <2-4059> type port-mstprstp <0-63>`**. Ensure that the new port is a member of this VLAN.
VOSS-5197	A BGP peer-group is uniquely identified by its name and not by its index. It is possible that the index that is configured for a peer-group changes between system reboots; however this has no functional impact.	None.
VOSS-7553	Option to configure the default queue profile rate-limit and weight values are inconsistent between EDM and CLI. Option to configure default values is missing in EDM.	None.
VOSS-7640	The same route is learned via multiple IPv6 routing protocols (a combination of two of the following : RIPng, OSPFv3 and BGPv6). In this specific case, an eBGP (current best – preference 45) route is replaced by and iBGP (preference 175) which in turn is replaced by and OSPFv3 (external 2) route (preference 125).	None.
VOSS-7647	With peer group configuration, you cannot configure Update Source interface with IPv6 loopback address in EDM.	Use CLI.
VOSS-9174	OVSDB remote VTEP and MAC details can take between 5 to 10 minutes to populate and display after a HW-VTEP reboots.	Known issue in VMware NSX 6.2.4. You can upgrade to NSX 6.4 to resolve this issue.

Table continues…

Issue number	Description	Workaround
VOSS-9462	OVSDB VNID I-SID MAC bindings are not populated on HW-VTEPs after configuration changes.	Known issue in VMware NSX 6.2.4. You can upgrade to NSX 6.4 to resolve this issue.
VOSS-10168	The system CLI does not prevent you from using the same IP address for theVXLAN Gateway hardware VTEP replication remote peer IP and OOB Management IP.	Manually check the IP configured as the OOB Management IP. Do not use the OOB Management IP address as the replication remote peer IP address.
VOSS-11817	The OVS connect-type for virtual service Vports is designed in such a way that it connects to any generic virtual machine (VM) guest OS version using readily available Ethernet device drivers. This design approach provides initial connectivity to the VM in a consistent manner. A consequence of this approach is that Vports created with connect-type OVS will show up as 1 Gbps interfaces in the VM even though the underlying Ethernet connection supports 10 Gbps .	If additional performance is desired, upgrade the VM guest OS with an Ethernet device driver that supports 10 Gbps interfaces.
VOSS-12151	If logical switch has only hardware ports binding, and not VM behind software VTEP, Broadcast, Unknown Unicast, and Multicast (BUM) traffic does not flow between host behind two hardware VTEP. The NSX replicator node handles the BUM traffic. NSX does not create the replicator node unless a VM is present. In an OVSDB topology, it is expected that at least one VM connects to the software VTEP. This issue is an NSX-imposed limitation.	After you connect the VM to the software VTEP, the issue is not seen.
VOSS-12395	You cannot use the following cables on 10 Gb fiber interfaces, or 40 Gb channelized interfaces, with the QSA28 adapter: • 1, 3, and 5 meter QSFP28 25 Gb DAC • 20 meter QSFP28 25 Gb AOC	n/a
wi01068569	The system displays a warning message that routes will not inject until	n/a

Table continues…

Issue number	Description	Workaround
	the apply command is issued after the enable command. The warning applies only after you enable redistribution, and not after you disable redistribution. For example: **`Switch:1(config)#isis apply redistribute direct vrf 2`**	
wi01112491	IS-IS enabled ports cannot be added to an MLT. The current release does not support this configuration.	n/a
wi01122478	Stale SNMP server community entries for different VRFs appear after reboot with no VRFs. On a node with a valid configuration file saved with more than the default vrf0, SNMP community entries for that VRF are created and maintained in a separate text file, `snmp_comm.txt`, on every boot. The node reads this file and updates the SNMP communities available on the node. As a result, if you boot a configuration that has no VRFs, you may still see SNMP community entries for VRFs other than the globalRouter vrf0 .	n/a
wi01137195	A static multicast group cannot be configured on a Layer 2 VLAN before enabling IGMP snooping on the VLAN. After IGMP snooping is enabled on the Layer 2 VLAN for the first time, static multicast group configuration is allowed, even when IGMP snooping is disabled later on that Layer 2 VLAN.	n/a
wi01138851	Configuring licenses using EDM is not supported.	n/a
wi01141638	When a VLAN with 1000 multicast senders is deleted, the console or Telnet session stops responding and SNMP requests time out for up to 2 minutes.	n/a
wi01142142	When a multicast sender moves from one port to another within the same BEB or from one vIST peer BEB to another, with the old port operationally up, the source port information in the output of the **`show ip igmp sender`**	You can perform one of the following workarounds: • On an IGMP snoop-enabled interface, you can flush IGMP sender records.

Table continues…

Issue number	Description	Workaround
	command is not updated with new sender port information.	**Caution:** Flushing sender records can cause a transient traffic loss. • On an IGMP-enabled Layer 3 interface, you can toggle the IGMP state. **Caution:** Expect traffic loss until IGMP records are built after toggling the IGMP state.
wi01145099	IP multicast packets with a time-to-live (TTL) equal to 1 are not switched across the SPB cloud over a Layer 2 VSN. They are dropped by the ingress BEB.	To prevent IP multicast packets from being dropped, configure multicast senders to send traffic with TTL greater than 1.
wi01159075	VSP 4450GTX-HT-PWR+: Mirroring functionality is not working for RSTP BPDUs.	None.
wi01171670	Telnet packets get encrypted on MACsec enabled ports.	None.
wi01198872	On VSP 4000 Series, a loss of learned MAC addresses occurs in a vIST setup beyond 10k addresses. In a SPB setup the MAC learning is limited to 13k MAC addresses, due to the limitation of the internal architecture when using SPB. Moreover, as vIST uses SPB and due to the way vIST synchronizes MAC addresses with a vIST pair, the MAC learning in a vIST setup is limited to 10K Mac addresses.	None.
wi01210217	The command **`show eapol auth-stats`** displays `LAST-SRC-MAC` for NEAP sessions incorrectly.	n/a
wi01211415	In addition to the fan modules, each power supply also has a fan. The power supply stops working if a power supply fan fails, but there is no LED or software warning that indicates this failure.	Try to recover the power supply fan by resetting the switch. If the fan does not recover, then replace the faulty power supply.

Table continues…

Issue number	Description	Workaround
wi01212034	When you disable EAPoL globally: • Traffic is allowed for static MAC configured on EAPoL enabled port without authentication. • Static MAC config added for authenticated NEAP client is lost.	n/a
wi01212247	BGP tends to have many routes. Frequent additions or deletions impact network connectivity. To prevent frequent additions or deletions, reflected routes are not withdrawn from client 2 even though they are withdrawn from client 1. Disabling route-reflection can create a black hole in the network.	Bounce the BGP protocol globally.
wi01212585	LED blinking in EDM is representative of, but not identical to, the actual LED blinking rates on the switch.	n/a
wi01213040	When you disable auto-negotiation on both sides, the 10 Gbps copper link does not come up.	n/a
wi01213066 wi01213374	EAP and NEAP are not supported on brouter ports.	n/a
wi01213336	When you configure tx mode port mirroring on T-UNI and SPBM NNI ports, unknown unicast, broadcast and multicast traffic packets that ingress these ports appear on the mirror destination port, although they do not egress the mirror source port. This is because tx mode port mirroring happens on the mirror source port before the source port squelching logic drops the packets at the egress port.	n/a
wi01219658	The command **`show khi port-statistics`** does not display the count for NNI ingress control packets going to the CP.	n/a
wi01219295	SPBM QOS: Egress UNI port does not follow port QOS with ingress NNI port and Mac-in-Mac incoming packets.	n/a
wi01223526	ISIS logs duplicate system ID only when the device is a direct neighbor.	n/a

Table continues…

Issue number	Description	Workaround
wi01223557	Multicast outage occurs on LACP MLT when simplified vIST peer is rebooted.	You can perform one of the following work arounds: • Enable PIM on the edge. • Ensure that IST peers are either RP or DR but not both.
wi01224683 wi01224689	Additional link bounce may occur on 10 Gbps ports when toggling links or during cable re-insertion. Additional link bounce may occur with 40 Gbps optical cables and 40 Gbps break-out cables, when toggling links or during cable re-insertion.	n/a
wi01229417	Origination and termination of IPv6 6-in-4 tunnel is not supported on a node with vIST enabled.	None.
wi01232578	When SSH keyboard-interactive-auth mode is enabled, the server generates the password prompt to be displayed and sends it to the SSH client. The server always sends an expanded format of the IPv6 address. When SSH keyboard-interactive-auth mode is disabled and password-auth is enabled, the client itself generates the password prompt, and it displays the IPv6 address format used in the **`ssh`** command.	None.
wi01234289	HTTP management of the ONA is not supported when it is deployed with a VSP 4000 Series device.	None.

VSP 4450GTX-HT-PWR+ Restrictions

⚠ Caution:

The VSP 4450GTX-HT-PWR+ has operating temperature and power restrictions. For safety and optimal operation of the device, ensure that the prescribed thresholds are strictly adhered to.

The following table provides a description of the restriction or behavior and the work around, if one exists.

Behavior	Description	Workaround
For high-temperature threshold	The VSP 4450GTX-HT-PWR+ supports a temperature range of 0°C to 70°C. In the alpha release, power supply does not shut down	To prevent equipment damage, ensure that the operating temperature is within the

Table continues…

Behavior	Description	Workaround
	at an intended over-temperature threshold of 79°C.	supported temperature range of 0°C to 70°C.
For power supply wattage threshold	Software functionality to reduce the POE power budget based on the number of operational power supplies and operating temperature is not available in the Alpha SW image.	Ensure that the POE device power draw is maintained at the following when the device is at temperatures between 61°C and 70°C: • 400W — with 1 operational power supply • 832W — with 2 operational power supplies
For inoperable external USB receptacle	The VSP 4450GTX-HT-PWR+ has an empty external USB receptacle that was not available in GTS models. Software to support the use of the external USB receptacle is not yet available in the Alpha SW image. Therefore the USB port is inoperable.	No workarounds are provided with the alpha image.

SSH Connections

VOSS 4.1.0.0 and VOSS 4.2.0.0 SSH server and SSH client support password authentication mode.

VOSS 4.2.1.0 changed the SSH server from password authentication to keyboard-interactive. VOSS 4.2.1.0 changed the SSH client to automatically support either password authentication or keyboard-interactive mode.

In VOSS 4.2.1.0, you cannot configure the SSH server to support password authentication. This limitation creates a backward compatibility issue for SSH clients that do not support keyboard-interactive mode, including SSH clients that are part of pre-VOSS 4.2.1.0 software releases. For example, VOSS 4.1.0.0 SSH clients, VOSS 4.2.0.0 SSH clients, and external SSH clients that only support password authentication cannot connect to VOSS 4.2.1.0 SSH servers.

This issue is addressed in software release VOSS 4.2.1.1 and later. The default mode of the SSH server starting from VOSS 4.2.1.1 is changed back to password authentication. Beginning with VOSS 5.0, you can use a CLI command to change the SSH server mode to keyboard-interactive.

For more information about how to configure the SSH server authentication mode, see Administering VOSS.

See the following table to understand SSH connections between specific client and server software releases.

Client software release	Server software release	Support
VOSS 4.1.0.0	VOSS 4.2.0.0	Supported
VOSS 4.1.0.0	VOSS 4.2.1.0	Not supported

Table continues…

Client software release	Server software release	Support
VOSS 4.2.0.0	VOSS 4.2.1.0	Not supported
VOSS 4.1.0.0	VOSS 4.2.1.1	Supported
VOSS 4.2.0.0	VOSS 4.2.1.1	Supported

Fabric Extend IP over ELAN/VPLS

This feature allows multiple switches running Fabric Extend IP to be directly connected over a Layer 2 broadcast domain without the need for loopback VRFs in Release 6.0 or later.

Releases earlier than 6.0 have a single next hop/ARP restriction that require the use of loopback VRFs to deploy Fabric Extend IP over ELAN/VPLS.

For more information, see Configuring Fabric Basics and Layer 2 Services for VOSS.

Redirect Next-hop Filter Restrictions

This feature does not behave the same way on all platforms:

- VSP 4000 Series and VSP 7400 Series

  The redirect next-hop filter redirects packets with a time-to-live (TTL) of 1 rather than sending them to the CPU where the CPU would generate ICMP TTL expired messages. IP Traceroute does not correctly report the hop. For more information, see Configuring QoS and ACL-Based Traffic Filtering for VOSS.

- VSP 7200 Series and VSP 8000 Series

  The redirect next-hop filter does not redirect packets with a time-to-live (TTL) of 1 nor does it send them to the CPU where the CPU would generate ICMP TTL expired messages. IP Traceroute reports a timeout for the hop. For more information, see Configuring QoS and ACL-Based Traffic Filtering for VOSS.

IP Source Guard Restrictions

If you enable Application Telemetry, IPv6 Source Guard commands and configurations are blocked and not available on VSP 4000 Series, VSP 7200 Series, and VSP 8000 Series switches.

Filter Restrictions

The following table identifies known restrictions.

Applies To	Restriction
ACL restrictions	
All platforms	Only port-based ACLs are supported on egress. VLAN-based ACLs are not supported.
All platforms	IPv6 ingress and egress QoS ACL/filters are not supported.
All platforms	Control packet action is not supported on InVSN Filter or IPv6 filters generally.
All platforms	IPv4/IPv6 VLAN based ACL filters will be applied on traffic received on all the ports if it matches VLAN ID associated with the ACL.

Table continues…

Applies To	Restriction
VSP 7200 Series VSP 7400 Series VSP 8000 Series	VLAN ID and VLAN_DOT1p attributes for untagged traffic are not supported for ingress/egress filters.
All platforms	Scaling numbers are reduced for IPv6 filters.
All platforms	The InVSN Filter does supports IP Shortcut traffic only on both UNI and NNI ports, but does not support IP Shortcut traffic on UNI ports only and NNI ports only.
All platforms	The InVSN Filter does not filter packets that arrive on NNI ingress ports but are bridged to other NNI ports or are for transit traffic.
All platforms	You can insert an inVsn ACL type for a Switched UNI only if the Switched UNI I-SID is associated with a platform VLAN.
ACE restrictions	
All platforms	When an ACE with action count is disabled, the statistics associated with the ACE are reset.
All platforms	Only security ACEs are supported on egress. QoS ACEs are not supported.
All platforms	ICMP type code qualifier is supported only on ingress filters.
All platforms	For port-based ACLs, you can configure VLAN qualifiers. Configuring port qualifiers are not permitted.
All platforms	For VLAN-based ACLs, you can configure port qualifiers. Configuring VLAN qualifiers are not permitted.
All platforms	Egress QoS filters are not supported for IPv6 filters.
All platforms	Ingress QoS filters are not supported for IPv6 filters.
All platforms	Source/Destination MAC addresses cannot be added as attributes for IPv6 filters ACEs.
VSP 4000 Series VSP 7200 Series VSP 8000 Series	If more than 256 IPv6 filters are configured, the number of IPv4 filters is reduced.
VSP 4000 Series VSP 7200 Series VSP 8000 Series	If you enable Application Telemetry, IPv6 security filter commands and configurations are blocked and not available.

Chapter 8: Resolved Issues

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

Fixes from Previous Releases

VOSS 8.1. incorporates all fixes from prior releases, up to and including VOSS 7.1.4 and VOSS 8.0.6.1.

Resolved Issues in VOSS 8.1

Issue number	Description
VOSS-1402	You cannot use EDM to configure SSH rekey, or to enable or disable SFTP.
VOSS-1404	You cannot use EDM to view the IPv6 DHCP relay counters.
VOSS-2415	The EDM tab name V3 Interface, available under IP is not available under IPv6 VRRP.
VOSS-9642	If you add more ports to an existing MLT used by an IPv6 tunnel to send traffic, the datapath records do not update to support the new port.
VOSS-9670	When rebooting the chassis, the following message can appear: `1 2018-03-05T11:16:36.168-05:00 AVL-156 CP1 - 0x002bc608 - 00000000 GlobalRouter VSPTALK WARNING cppTap unexpected IO error fd 137 errno 100.`
VOSS-11084	In highly scaled environments the command **`show vnid mac-address-entry`** can be slow in printing the expected output.
VOSS-12229	In a vIST/SMLT scenario with IGMP SPB snooping receivers on an SMLT link using IGMPv3, and the sender is local to one of the vIST peers, if you shut down the SMLT ports to change the MLT group membership or change the MLT type from SMLT admin to normal MLT, when you bring the ports back up, traffic from the stream may no longer flow to receivers on the SMLT link.
VOSS-13050	Deleting a VLAN configured for VRRP with no VLAN IP address causes switch reset.
VOSS-13070	Using pipe with **`show running-config`** to include patterns does not work when IP name server configured.
VOSS-13127	Chassis reset due to receiving a DHCP packet with a malformed value in the UDP length field.

Table continues…

Issue number	Description
VOSS-13158	Certificate enhancements: • relax consistency checks for CSR generation by introducing a new 'relaxed' mode for installing them via **`certificate generate-csr relaxed`** command; • relaxed mode also allows for adding SANs (Subject Alternative Names) to CSR • support for PKCS12
VOSS-13185	vIST staying up on VSP 8600 even when vIST peer is rebooted or powered off.
VOSS-13193	MPLS packets with EtherType 0x8847 are not passing over a T-UNI, other EtherTypes work.
VOSS-13194	IS-IS adjacency not coming up for manual-area “00.0000”.
VOSS-13198	Unable to communicate between VRF CLIP IP.
VOSS-13537	Attempting to redistribute a BGP route into IS-IS as an external route with a route-map containing a **`match community xyz`** clause, then that route is getting redistributed into IS-IS to the peer device even if the match fails.
VOSS-13585	**`ip forward-protocol udp`** fails to forward packet received on Layer 2 VSN when no "up" port exists in Platform VLAN.
VOSS-13638	After disabling MSTP on a port, re-enabling it causes the following message: `Error: port 1/40, Invalid value given to MSTP.`
VOSS-13713	**`show isis spbm ip-multicast detail`** for long egress port list may cause chassis to reset. If list exceeds the length of the buffer, "..." is appended to the output display.
VOSS-13731	LLDP neighbor content is not displayed for back-to-back port connections.
VOSS-13754	Chassis may reset during ARP cleanup.
VOSS-13768	Fabric Extend tunnel flapping can occur when IPSec is enabled with F&R with IMIX traffic. Issue is seen only on XA1440 device.
VOSS-13792	If you change a DvR leaf node Virtual IST configuration, vIST may not come up again after the change.
VOSS-13797	Changed traces that were using VERY_TERSE level to TERSE level in PLSB FIB.
VOSS-13783	DvR leaf reports DVR ERROR L3_ENTRY table limit reached. Error counting resources.
VOSS-13825	**`show isis logical interface`** output repeats forever.
VOSS-13835	GlobalRouter IPMC ERROR Insufficient VFI/VPN resources to create McoSpb source.

VOSS-13850	Switch crashed with reboot after memory usage reached 95% because of ARP flap condition. Add log when ARP Flap is detected for up to 100 ARPs.
	`EventCode: 0x000385ff`
	`AlarmId: <0>`

Table continues…

Issue number	Description
	AlarmStatus: <ALARM_NONE> ModuleName: <MOD_P2IP> Severity: <S_ERROR> TerseMsg: <"MAC Flapping detected on Host ARP %d.%d.%d.%d MAC %s"> ProbableCause: <"This message indicates that the Host ARP MAC is flapping. Customer needs to find the offending device to see why the MAC is flapping."> Remedy: <"Check the network for this IP/MAC combination."> # EventCode: 0x00038600 AlarmId: <0x00d00005> AlarmStatus: <ALARM_SET> ModuleName: <MOD_P2IP> Severity: <S_ERROR> TerseMsg: <"MAC Flapping detected on Host ARP %d.%d.%d.%d MAC %s"> ProbableCause: <"This message indicates that the Host ARP MAC is flapping. Customer needs to find the offending device to see why the MAC is flapping."> Remedy: <"Check the network for this IP/MAC combination."> # EventCode: 0x00038601 AlarmId: <0x00d00005> AlarmStatus: <ALARM_CLEAR> ModuleName: <MOD_P2IP> Severity: <S_ERROR> TerseMsg: <"Cleared the condition for MAC Flapping detected: Host ARP %d.%d.%d.%d MAC %s"> ProbableCause: <"This message indicates the condition was cleared."> Remedy: <"No action required"> #
VOSS-13860	Chassis reset during ARP age out.
VOSS-13893	VLAN I-SID mapping can be overwritten without warning message.
VOSS-13924	You may 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. 1 2019-05-30T13:46:03.941-04:00 wolfboro-1 CP1 - 0x00064724 - 00000000 GlobalRouter MLT WARNING SMLT buffer usage over 200M. Low memory warning condition 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 istRxLearnMacCnt 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

Table continues…

Issue number	Description
VOSS-13944	After you enable MSTP-Fabric Connect Multi Homing (**spbm 1 stp-multi-homing enable**), the following error messages appear in the log: • On VSP 8200: `2019-05-16T13:03:55.596Z VSP8200-2 IO1 - 0x0012852b - 00000000 GlobalRouter COP-SW ERROR ercdProcEgressVlanMsg: EGRESS VLAN CTRL Message>> VLAN=4093/port=255 combination is invalid` • `2019-05-16T12:16:10.928Z VSP7400-1 CP1 - 0x000145f2 - 00000000 GlobalRouter BRIDGE WARNING Spanning Tree: Port unknown - Received 2 TCs in 1 minute(s)`
VOSS-13945	**stp-multi-homing** shows unknown port in **show spanning-tree config**.
VOSS-13946	After you enable MSTP-Fabric Connect Multi Homing (**spbm 1 stp-multi-homing enable**) and a BEB acts as root bridge, you cannot tell from another BEB which BEB is the root bridge. The Spanning Tree show commands show the port towards the root bridge. If the root is across the fabric, the **show spanning-tree mstp status** command output shows `fabric` for root port.
VOSS-13977	Chassis reset during ARP deletion.
VOSS-14029	Limit SFTP access to same as FTP according to access permissions.
VOSS-14030	Reports incorrect port num in the log when MAC is corrected to point to vIST after learned on non IST port due to loop on the edge.
VOSS-14042	LLDP PDU with a TLV length greater than 32 bytes caused chassis reset. System truncates string to 32 bytes if length exceeded.
VOSS-14056	Resource manager leak when VRF is deleted. Seen when enabling/disabling DvR controller Following logs may be seen: `CP1 [05/24/19 02:57:19.973:EDT] 0x001087d7 00000000 GlobalRouter RCIP6 INFO 85% of route limit reached for combined ipv4/v6 routes: total(13383), ipv4(3), ipv6 <=64 prefix length(6690)`
VOSS-14064	**clear dvr host-entries** command may cause chassis reset.
VOSS-14068	Allow **untag-port-default-vlan** for MLT/LACP trunks.
VOSS-14089	vIST Peer MAC movement log support for MLT. Log `COP-SW INFO VIST peer mac f4:6e:95:9e:c0:81 on VID 3000 is learnt on non-IST port 1/18, Pointing record back to IST port` reports incorrect non-vIST port.
VOSS-14095	Change log for ICMP Needfrag Reply `MTU size exceeded, next hop MTU size: x` to include the source IP and the required MTU.
VOSS-14104	Packets are dropped on Fabric Extend logical interface if the packet size is larger than the tunnel MTU. This scenario can occur if there are devices such

Table continues…

Issue number	Description
	as firewalls placed in front of XA1400 Series devices, which performs bracket reassembly to detect possible anomalies.
VOSS-14107	**`no ssh encryption`** configuration truncated in saved config (partial configuration loss may occur).
VOSS-14148	Missing GRT default route when DvR Controller Leaf Link bounce.
VOSS-14157	High CPU utilization, after software upgrade, caused by SNMP task(tSnmpTmr) Exception in the SNMP retransmit logic can cause loop.
VOSS-14200	After rebooting VSP7254XSQ, a QSFP+ card won't come up without re-plugging.
VOSS-14210	BGP log messages missing when BGP session goes down.
VOSS-14211	For configurations using brouter interfaces as the FE tunnel endpoints: if the node learns about a non-direct route (say a default route) that encompasses the tunnel endpoint prior to the direct interface coming up the logical ISIS adjacency will fail to establish.
VOSS-14216	Unexpected error when trying to install PKCS12 file but the file does not exist in flash.
VOSS-14217	Cannot install PKCS12 file if a public/private key does not already exist on the switch.
VOSS-14218	A CSR can be generated with relax option when CN or SAN is configured, the error shown when neither CN nor SAN are configured should be more clear (just one of the parameters are required).
VOSS-14220	User should not be allowed to configure invalid values for the certificate SAN.

VOSS-14267	Two DvR controllers reset following 95% memory utilization. Excessive control messages cause memory leak. Log messages added to indicate state of DBsync Message queue:
	`EventCode: 0x00390606`
	`AlarmId: <0x00000000>`
	`AlarmStatus: <ALARM_NONE>`
	`ModuleName: <MOD_DBSYNC>`
	`Severity: <S_WARNING>`
	`TerseMsg: <"Message queue length from DB Sync to tMain reached warning threshold">`
	`ProbableCause: <"CPU utilization is high">`
	`Remedy: <"Check alarm status and network configuration">`
	`EventCode: 0x00390607`
	`AlarmId: <0x00000000>`
	`AlarmStatus: <ALARM_NONE>`
	`ModuleName: <MOD_DBSYNC>`
	`Severity: <S_WARNING>`
	`TerseMsg: <"Message queue length from DB Sync to tMain threshold cleared ">`
	`ProbableCause: <"This message indicates that the queue length has returned to normal operating range">`
	`Remedy: <"No action required">`

VOSS-14276	Chassis reset for no reason. Fixed thread death handling.
VOSS-14278	Use of single quote in regular expression caused chassis reset.

Table continues…

Issue number	Description
VOSS-14315	Chassis reset after entering **`clear telnet 0`** while logged in via telnet.
VOSS-14366	Sunon fans generate many fault on VSP 7432CQ and VSP 7400-48Y platforms.
VOSS-14373	DvR dbsync queue full messages when flapping same IP with 2 MACs.
VOSS-14396	Unable to provision more than 24 VRFs with Genlic license.
VOSS-14444	100gigER4-Lite optic support.
VOSS-14478	Configuration loss after reboot when MSTP-Fabric Connect Multi Homing is enabled on the SPBM instance.
VOSS-14484	Not all DvR host-entries are relearned by all BEBs when **`clear dvr host-entries`** command is issued.
VOSS-14496	SSH is disabled after reboot.
VOSS-14508	**`ipv6 dhcp-relay fwd-path`** configuration lost on reboot.
VOSS-14585	VSP 7200 Series and VSP 7400 Series platforms may not automatically reset after a core dump.
VOSS-14587	Starting multiple VMs results in CPU usage overlap with VOSS and other VMs. VOSS processes affected by VMs.
VOSS-14607	When inserting an optical transciever into an XA1400 the destination connection flaps once before remaining up.
VOSS-14608	VSP 7432CQ and VSP 7400-48Y Fan tray unit2 operational message and alarm message missing.
VOSS-14672	Recovery of Fabric Extend tunnel adjacency can be 4 to 5 minutes after changing a tunnel source IP address.
VOSS-14679	Add one alarm per DvR arp/mac flapping detected.
VOSS-14712	SMLT high/low memory wrong alarm detection.
VOSS-14719	When port is configured as tagged and Egress-VLANID is used alongside Tunnel-Private-Group-ID attribute, the tagging is taken from port level. Egress-VLANID is overwriting port level tagging only to tagged value not otherwise.
VOSS-14765	Add certificate enhancement MIB support for VSP 7400 Series.
VOSS-14794	Forwarding records point to vIST on DvR leaf node instead of client port after reboot of leaf node.
VOSS-14880	**`untag-port-default-vlan`** port command not effective upon enabling EAPoL on a port. Added consistency check preventing EAPol from overriding **`untag-port-default-vlan`** port command.
VOSS-14943	Core dump after IGMP configuration changes when more than one IGMPv3 members joining same group on same port and explicit host tracking is enabled.
VOSS-14991	`The QOS Egress shaper rate can not exceed the port capability` when trying to configure 10G port to shape to 8G.

Table continues…

Issue number	Description
VOSS-15016	SFP+ ports do not power down if 'shut' command is issued without inserting a supported optical transceiver. If you insert a non-DAC optical transceiver on an admin down port 1/5 or 1/6, the peer end link still comes up and port status and activity LED on the XA1400 is on. The root cause is port 1/5,1/6 does not power down if the 'shut' command is issued without first inserting a supported optical transceiver on that port. When an optical transceiver is inserted on a admin down port, the peer end link still comes up in the following two scenarios: • Bootup - inserting optics with default config, when port is in admin shut state. • Run time - in link up state remove existing optics, shut the port and re-insert the optic. When a optics is inserted in above scenario, the peer end link still comes up and local side link LED glows. With 10G optics the issue is seen in both scenarios. With 1G optics the issue is seen only during run time insertion of optics and not at bootup
VOSS-15073	Traffic loss at DvR Leafs after all DvR controllers are rebooted and come back online.
VOSS-15091	VSP is not assigning a value to the IP identification field on IS-IS packets.
VOSS-15119	Mirrored packets are not seen on the second mirroring destination port.
VOSS-15129	**`l2 tracetree-vclan`** command should be disabled.
VOSS-15154	ARP entries stop forming over time. Resources leaked when host route add fails because an inclusive local route exists.
VOSS-15219	Node may hang after core dump occurs that require power cycle to recover.
VOSS-15220	Node is not assigning a value to the IP identification field on IS-IS packets causing issues with downstream paths that may use fragmentation and reassembly.
VOSS-15226	DvR leaf vIST pair are not removing DvR hosts learned via LACP/FA/SMLT attached devices when SMLT goes down.
VOSS-15232	DvR leaf vIST pairs do not port lock SMLT ports during bootup. Multiple seconds of packet loss when leaf comes back online.
VOSS-15292	QSFP28-SR4-100G connected to firewall, not coming up when shut/no-shut or when the box is rebooted until FEC is set manually again.
VOSS-15308	Node becomes hung when **`snmp-server engine-id`** command executed with empty password.
VSP4000-247	Core Dump Generated with memory leak "GlobalRouter SW ERROR Memory reached critical level of 95% utilization from SLA Mon activity.
VSP4000-248	IS-IS adjacency over FE Tunnel not coming up.
VSP7200-78	vIST peers reset when short ARP packet destined to the SLA Mon agent IP address.

Appendix A: Related Information

MIB Changes

Deprecated MIBs

Table 33: Common

Object Name	Object OID	Deprecated in VOSS Release
rcChasFanTable	1.3.6.1.4.1.2272.1.4.7.1	7.0
rcChasFanEntry	1.3.6.1.4.1.2272.1.4.7.1.1	7.0
rcChasFanId	1.3.6.1.4.1.2272.1.4.7.1.1.1	7.0
rcChasFanOperStatus	1.3.6.1.4.1.2272.1.4.7.1.1.2	7.0
rcChasFanAmbientTemperature	1.3.6.1.4.1.2272.1.4.7.1.1.3	7.0
rcChasFanType	1.3.6.1.4.1.2272.1.4.7.1.1.4	7.0
rcChasLedTable	1.3.6.1.4.1.2272.1.4.65.1	7.0
rcChasLedEntry	1.3.6.1.4.1.2272.1.4.65.1.1	7.0
rcChasLedId	1.3.6.1.4.1.2272.1.4.65.1.1.1	7.0
rcChasLedLabel	1.3.6.1.4.1.2272.1.4.65.1.1.2	7.0
rcChasLedStatus	1.3.6.1.4.1.2272.1.4.65.1.1.3	7.0
rcIpConfBfdVrfId	1.3.6.1.4.1.2272.1.8.1.12.1.8	8.1
rcIpConfBfdVrfName	1.3.6.1.4.1.2272.1.8.1.12.1.9	8.1
rcIpv6InterfaceBfdVrfId	1.3.6.14.1.2272.1.62.1.1.14.1.8	8.1
rcIpv6InterfaceBfdVrfName	1.3.6.14.1.2272.1.62.1.1.14.1.9	8.1

Modified MIBs

Table 34: Common

Object Name	Object OID	Modified in VOSS Release	Modification
rcVossSystemMgmtPortLedStatus	1.3.6.1.4.1.2272.1.101.1.1.1.1	8.0.5	ADD_NEW_VALUES: 4 and 5 for first 4 bits (left led)
rcPortType	1.3.6.1.4.1.2272.1.4.10.1.1.2	8.1	ADD_ENUM: 195-212
rcVossSystemVimAdminSpeed	1.3.6.1.4.1.2272.1.101.1.1.1.3	8.1	ADD_ENUM: unsupported(3)
rcVossSystemCardLedId	1.3.6.1.4.1.2272.1.101.1.1.5.1.2	8.1	CHANGE_RANGE: Changed the range from 1..5 to 1..9

Table 35: VSP 4000 Series

Object Name	Object OID	Modified in VOSS Release	Modification
rcPortType	1.3.6.1.4.1.2272.1.4.10.1.1.2	8.0.5	Corrected or added missing values for rc10GbAOC(191), rc100GbSWDM4(192), rc100GbSWDM4Channelized(193), and rc10GbInsight(194).

Table 36: VSP 7400 Series

Object Name	Object OID	Modified in VOSS Release	Modification
rcVossSystemTemperatureSensorIndex	1.3.6.1.4.1.2272.1.101.1.1.2.1.1	8.0	Changed the range from 1..4 to 1..7
rcVossSystemCardLedId	1.3.6.1.4.1.2272.1.101.1.1.5.1.2	8.0	Changed the range from 1..4 to 1..5
rcPortType	1.3.6.1.4.1.2272.1.4.10.1.1.2	8.0.5	Corrected or added missing values for rc10GbAOC(191), rc100GbSWDM4(192), rc100GbSWDM4Channelized(193), and rc10GbInsight(194).
rcVirtualServiceName	1.3.6.1.4.1.2272.1.101.1.1.8.1.1	8.0.5	Changed the range from 1..128 to 1..80

Table continues…

Object Name	Object OID	Modified in VOSS Release	Modification
rcVirtualServiceDiskVirtServName	1.3.6.1.4.1.2272.1.101.1.1.9.1.1	8.0.5	Changed the range from 1..128 to 1..80
rcVirtualServiceVPortsVirtServName	1.3.6.1.4.1.2272.1.101.1.1.10.1.1	8.0.5	Changed the range from 1..128 to 1..80
rcVirtualServiceVPortsName	1.3.6.1.4.1.2272.1.101.1.1.10.1.2	8.0.5	Changed the range from 1..128 to 1..32
rcVirtualServiceApplicationName	1.3.6.1.4.1.2272.1.101.1.1.11.1.1	8.0.5	Changed the range from 1..128 to 1..80

Table 37: VSP 8000 Series

Object Name	Object OID	Modified in VOSS Release	Modification
rcPortType	1.3.6.1.4.1.2272.1.4.10.1.1.2	8.0.5	Corrected or added missing values for rc10GbAOC(191), rc100GbSWDM4(192), rc100GbSWDM4Channelized(193), and rc10GbInsight(194).

Table 38: XA1400 Series

Object Name	Object OID	Modified in VOSS Release	Modification
rcIsisCircuitPlsbState	1.3.6.1.4.1.2272.1.63.5.1.4	8.0.50	Changed default value from disable to enable
rcIsisGlobalIpTunnelMtu	1.3.6.1.4.1.2272.1.63.1.20	8.0.50	Changed range from 750..1950 to 0 | 750..1950
rcSysMTUSize	1.3.6.1.4.1.2272.1.1.55	8.0.50	Supports MTU frame size 9000 value 2 (instead of 9600 MTU)
rcVlanType	1.3.6.1.4.1.2272.1.3.2.1.10	8.0.50	Supports only byPort(1) and spbm-bvlan(11) values
rcVlanProtocolId	1.3.6.1.4.1.2272.1.3.2.1.15	8.0.50	Supports only none(0) value

Table continues…

Object Name	Object OID	Modified in VOSS Release	Modification
rcPlsbGlobalEtherType	1.3.6.1.4.1.2272.1.78.1.4	8.0.50	Supports only default value (0x8100)
rcIsidServiceType	1.3.6.1.4.1.2272.1.87.2.1.2	8.0.50	Supports only l2vsn(4) value
rcPrFilterAclType	1.3.6.1.4.1.2272.1.202.1.1.2.3.1.1.4	8.0.50	Supports only inVlan(1), inPort(3) and outPort(4) values
rcIsisLogicalInterfaceType	1.3.6.1.4.1.2272.1.63.26.1.3	8.0.50	Supports only ip(2) value
rc2kCpuSerialPortBaudRate	1.3.6.1.4.1.2272.1.100.3.1.6	8.0.50	Supports only baud115200 value
rcIsisGlobalIpTunnelMtu	. 1.3.6.1.4.1.2272.1.63.1.20.0	8.1	CHANGE_RANGE: Changed the range from 750..1950 to 750..9000
rcIsisLogicalInterfaceShapingRate	1.3.6.1.4.1.2272.1.63.26.1.16	8.1	CHANGE_RANGE: Changed the range from 0..5000 to 0..1000. Changed the type from Integer32 to INTEGER.
rcCfmTransmitL2IpPingIpAddrType	1.3.6.1.4.1.2272.1.69.37.1.1	8.1.1	Supports only Ipv4 address type.
rcCfmTransmitL2IpPingIpAddr	1.3.6.1.4.1.2272.1.69.37.1.2	8.1.1	Supports only ipv4 address.
rcIpsecPolicyDstAddressType	1.3.6.1.4.1.2272.1.213.1.1.2	8.1.1	Supports only Ipv4 address type.
rcIpsecPolicyDstAddress	1.3.6.1.4.1.2272.1.213.1.1.3	8.1.1	Supports only ipv4 address.
rcIpsecPolicySrcAddressType	1.3.6.1.4.1.2272.1.213.1.1.4	8.1.1	Supports only Ipv4 address type.

Table continues…

Object Name	Object OID	Modified in VOSS Release	Modification
rcIpsecPolicySrcAddress	1.3.6.1.4.1.2272.1.213.1.1.5	8.1.1	Supports only ipv4 address.
rcIpsecPolicyL4Protocol	1.3.6.1.4.1.2272.1.213.1.1.9	8.1.1	Supports only tcp(6), udp(17) and icmp(1) values.
rcIpRedistributeInterVrfProtocol	1.3.6.1.4.1.2272.1.8.100.22.1.2	8.1.1	Supports only ospf(1), bgp(2), isis(3), vrf-ext(4) and rip(6) values.
rcIpRedistributeInterVrfRouteSource	1.3.6.1.4.1.2272.1.8.100.22.1.4	8.1.1	Supports only direct(1), static(2), rip(3), ospf(4), bgp(5) and isis(13) values.
rcCfmTransmitL2IpTraceIpAddrType	1.3.6.1.4.1.2272.1.69.39.1.1	8.1.1	Supports only ipv4 address type.
rcCfmTransmitL2IpTraceIpAddr	1.3.6.1.4.1.2272.1.69.39.1.2	8.1.1	Supports only ipv4 address.

New MIBs

Table 39: Common

Object Name	Object OID	New in VOSS Release
rcPortOperAutoNegotiate	1.3.6.1.4.1.2272.1.4.10.1.1.128	8.0.5
rcPortOperForwardErrorCorrection	1.3.6.1.4.1.2272.1.4.10.1.1.129	8.0.5
rcnChasPowerSupplyDifferentInputTypesDetectedTrap	1.3.6.1.4.1.2272.1.21.0.355	8.0.5
rcnChasPowerSupplyDifferentInputTypesClearTrap	1.3.6.1.4.1.2272.1.21.0.356	8.0.5
rcPortVLacpFlapDetectEnable	1.3.6.1.4.1.2272.1.4.10.5.1.11	8.1
rcPortVLacpFlapFrequency	1.3.6.1.4.1.2272.1.4.10.5.1.12	8.1
rcPortVLacpFlapInterval	1.3.6.1.4.1.2272.1.4.10.5.1.13	8.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcPortVLacpTotalFlapCount	1.3.6.1.4.1.2272.1.4.10.5.1.14	8.1
rcPortVLacpFirstFlapTimeStamp	1.3.6.1.4.1.2272.1.4.10.5.1.15	8.1
rcPortVLacpLastFlapTimeStamp	1.3.6.1.4.1.2272.1.4.10.5.1.16	8.1
rcPortVLacpFlapClearStats	1.3.6.1.4.1.2272.1.4.10.5.1.17	8.1
rcIpConfBfdEnable	1.3.6.1.4.1.2272.1.8.1.1.1.25	8.1
rcIpConfOspfBfdEnable	1.3.6.1.4.1.2272.1.8.1.2.1.12	8.1
rcIpConfBfdTable	1.3.6.1.4.1.2272.1.8.1.12	8.1
rcIpConfBfdEntry	1.3.6.1.4.1.2272.1.8.1.12.1	8.1
rcIpStaticRouteBfdTableSize	1.3.6.1.4.1.2272.1.8.15.3.0	8.1
rcIpStaticRouteBfdTable	1.3.6.1.4.1.2272.1.8.15.4	8.1
rcIpStaticRouteBfdEntry	1.3.6.1.4.1.2272.1.8.15.4.1	8.1
rcIpStaticRouteBfdNextHop	1.3.6.1.4.1.2272.1.8.15.4.1.1	8.1
rcIpStaticRouteBfdRowStatus	1.3.6.1.4.1.2272.1.8.15.4.1.2	8.1
rcIpStaticRouteBfdVrfId	1.3.6.1.4.1.2272.1.8.15.4.1.3	8.1
rcIpStaticRouteBfdVrfName	1.3.6.1.4.1.2272.1.8.15.4.1.4	8.1
rcIpBgpPeerGroupBfdEnable	1.3.6.1.4.1.2272.1.8.101.11.1.39	8.1
rcIpBgpExtPeerAfBfdEnable	1.3.6.1.4.1.2272.1.8.101.16.6.1.39	8.1
rcIpBfdTrapEnabled	1.3.6.1.4.1.2272.1.8.104.1	8.1
rcIpv6InterfaceBfdEnable	1.3.6.1.4.1.2272.1.62.1.1.2.1.28	8.1
rcIpv6InterfaceBfdTable	1.3.6.1.4.1.2272.1.62.1.1.14	8.1
rcIpv6InterfaceBfdEntry	1.3.6.1.4.1.2272.1.62.1.1.14.1	8.1
rcIpv6InterfaceBfdIfIndex	1.3.6.1.4.1.2272.1.62.1.1.14.1.1	8.1
rcIpv6InterfaceBfdMinRxInterval	1.3.6.1.4.1.2272.1.62.1.1.14.1.2	8.1
rcIpv6InterfaceBfdTxInterval	1.3.6.1.4.1.2272.1.62.1.1.14.1.3	8.1
rcIpv6InterfaceBfdMultiplier	1.3.6.1.4.1.2272.1.62.1.1.14.1.4	8.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcIpv6InterfaceBfdVrfId	1.3.6.1.4.1.2272.1.62.1.1.14.1.8	8.1
rcIpv6InterfaceBfdVrfName	1.3.6.1.4.1.2272.1.62.1.1.14.1.9	8.1
rcIpv6StaticRouteBfdTable	1.3.6.1.4.1.2272.1.62.1.1.24	8.1
rcIpv6StaticRouteBfdEntry	1.3.6.1.4.1.2272.1.62.1.1.24.1	8.1
rcIpv6StaticRouteBfdIfIndex	1.3.6.1.4.1.2272.1.62.1.1.24.1.1	8.1
rcIpv6StaticRouteBfdNextHop	1.3.6.1.4.1.2272.1.62.1.1.24.1.2	8.1
rcIpv6StaticRouteBfdRowStatus	1.3.6.1.4.1.2272.1.62.1.1.24.1.3	8.1
rcIpv6StaticRouteBfdVrfId	1.3.6.1.4.1.2272.1.62.1.1.24.1.4	8.1
rcIpv6StaticRouteBfdVrfName	1.3.6.1.4.1.2272.1.62.1.1.24.1.5	8.1
rcOspfv3IfBfdEnable	1.3.6.1.4.1.2272.1.67.1.1.7.1.27	8.1
rcPlugOptModExtremeSku	1.3.6.1.4.1.2272.1.71.1.1.83	8.1
rcPlugOptModPowerOnCounter	1.3.6.1.4.1.2272.1.71.1.1.84	8.1
rcPlugOptModTxDdmInitial	1.3.6.1.4.1.2272.1.71.1.1.85	8.1
rcPlugOptModTxDdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.86	8.1
rcPlugOptModRxDdmInitial	1.3.6.1.4.1.2272.1.71.1.1.87	8.1
rcPlugOptModRxDdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.88	8.1
rcPlugOptModQSFPTx1DdmInitial	1.3.6.1.4.1.2272.1.71.1.1.89	8.1
rcPlugOptModQSFPTx1DdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.90	8.1
rcPlugOptModQSFPRx1DdmInitial	1.3.6.1.4.1.2272.1.71.1.1.91	8.1
rcPlugOptModQSFPRx1DdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.92	8.1
rcPlugOptModQSFPTx2DdmInitial	1.3.6.1.4.1.2272.1.71.1.1.93	8.1
rcPlugOptModQSFPTx2DdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.94	8.1
rcPlugOptModQSFPRx2DdmInitial	1.3.6.1.4.1.2272.1.71.1.1.95	8.1
rcPlugOptModQSFPRx2DdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.96	8.1
rcPlugOptModQSFPTx3DdmInitial	1.3.6.1.4.1.2272.1.71.1.1.97	8.1
rcPlugOptModQSFPTx3DdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.98	8.1
rcPlugOptModQSFPRx3DdmInitial	1.3.6.1.4.1.2272.1.71.1.1.99	8.1
rcPlugOptModQSFPRx3DdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.100	8.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcPlugOptModQSFPTx4DdmInitial	1.3.6.1.4.1.2272.1.71.1.1.101	8.1
rcPlugOptModQSFPTx4DdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.102	8.1
rcPlugOptModQSFPRx4DdmInitial	1.3.6.1.4.1.2272.1.71.1.1.103	8.1
rcPlugOptModQSFPRx4DdmLastGasp	1.3.6.1.4.1.2272.1.71.1.1.104	8.1
rcPlugOptModSupportExtremeExtraFeatures	1.3.6.1.4.1.2272.1.71.1.1.105	8.1
rcBfd	1.3.6.1.4.1.2272.1.81	8.1
rcBfdTmpMib	1.3.6.1.4.1.2272.1.81.1	8.1
rcBfdTmpObjects	1.3.6.1.4.1.2272.1.81.1.1	8.1
rcBfdTmpScalarObjects	1.3.6.1.4.1.2272.1.81.1.1.1	8.1
rcBfdTmpAdminStatus	1.3.6.1.4.1.2272.1.81.1.1.1.1	8.1
rcBfdTmpVersionNumber	1.3.6.1.4.1.2272.1.81.1.1.1.3	8.1
rcBfdTmpSessTable	1.3.6.1.4.1.2272.1.81.1.1.2	8.1
rcBfdTmpSessEntry	1.3.6.1.4.1.2272.1.81.1.1.2.1	8.1
rcBfdTmpSessPerfTable	1.3.6.1.4.1.2272.1.81.1.1.3	8.1
rcBfdTmpSessPerfEntry	1.3.6.1.4.1.2272.1.81.1.1.3.1	8.1
rcBfdTmpSessPerfPktIn	1.3.6.1.4.1.2272.1.81.1.1.3.1.1	8.1
rcBfdTmpSessPerfClearStats	1.3.6.1.4.1.2272.1.81.1.1.3.1.10	8.1
rcBfdTmpSessPerfVrfId	1.3.6.1.4.1.2272.1.81.1.1.3.1.11	8.1
rcBfdTmpSessPerfVrfName	1.3.6.1.4.1.2272.1.81.1.1.3.1.12	8.1
rcBfdTmpSessPerfPktOut	1.3.6.1.4.1.2272.1.81.1.1.3.1.2	8.1
rcBfdTmpSessUpTime	1.3.6.1.4.1.2272.1.81.1.1.3.1.3	8.1
rcBfdTmpSessPerfLastSessDownTime	1.3.6.1.4.1.2272.1.81.1.1.3.1.4	8.1
rcBfdExtMib	1.3.6.1.4.1.2272.1.81.2	8.1
rcBfdExtObjects	1.3.6.1.4.1.2272.1.81.2.1	8.1
rcBfdExtSessTable	1.3.6.1.4.1.2272.1.81.2.1.1	8.1
rcBfdExtSessEntry	1.3.6.1.4.1.2272.1.81.2.1.1.1	8.1
rcBfdExtSessPeerState	1.3.6.1.4.1.2272.1.81.2.1.1.1.1	8.1
rcBfdExtSessPeerAddrType	1.3.6.1.4.1.2272.1.81.2.1.1.1.2	8.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcBfdExtSessPeerAddr	1.3.6.1.4.1.2272.1.81.2.1.1.1.3	8.1
rcBfdExtSessApp	1.3.6.1.4.1.2272.1.81.2.1.1.1.4	8.1
rcBfdExtSessAppRun	1.3.6.1.4.1.2272.1.81.2.1.1.1.5	8.1
rcNlsMgmtAddressTable	1.3.6.1.4.1.2272.1.223.2	8.1.1
rcNlsMgmtIpArpTable	1.3.6.1.4.1.2272.1.223.3	8.1.1
rcNlsMgmtIpStaticRouteTable	1.3.6.1.4.1.2272.1.223.5	8.1.1
rcNlsMgmtStatsTable	1.3.6.1.4.1.2272.1.223.7	8.1.1
rcNlsMgmtIpRouteTable	1.3.6.1.4.1.2272.1.223.8	8.1.1
rcCloudIq	1.3.6.1.4.1.2272.1.230	8.1.1
rcCloudIqObjects	1.3.6.1.4.1.2272.1.230.1	8.1.1
rcCloudIqScalars	1.3.6.1.4.1.2272.1.230.1.1	8.1.1
rcCloudIqAgentEnable	1.3.6.1.4.1.2272.1.230.1.1.1	8.1.1
rcCloudIqAgentVersion	1.3.6.1.4.1.2272.1.230.1.1.2	8.1.1
rcCloudIqServerAddressType	1.3.6.1.4.1.2272.1.230.1.1.3	8.1.1
rcCloudIqServerAddress	1.3.6.1.4.1.2272.1.230.1.1.4	8.1.1
rcCloudIqProxyAddressType	1.3.6.1.4.1.2272.1.230.1.1.5	8.1.1
rcCloudIqProxyAddress	1.3.6.1.4.1.2272.1.230.1.1.6	8.1.1
rcCloudIqProxyTcpPort	1.3.6.1.4.1.2272.1.230.1.1.7	8.1.1
rcCloudIqProxyUserName	1.3.6.1.4.1.2272.1.230.1.1.8	8.1.1
rcCloudIqProxyPassword	1.3.6.1.4.1.2272.1.230.1.1.9	8.1.1
rcCloudIqNotificationEnable	1.3.6.1.4.1.2272.1.230.1.1.10	8.1.1
rcCloudIqOperStatus	1.3.6.1.4.1.2272.1.230.1.1.11	8.1.1
rcCloudIqAssociationUrl	1.3.6.1.4.1.2272.1.230.1.1.12	8.1.1
rcCloudIqPollUrl	1.3.6.1.4.1.2272.1.230.1.1.13	8.1.1
rcCloudIqMonitorFreq	1.3.6.1.4.1.2272.1.230.1.1.14	8.1.1
rcCloudIqPollFreq	1.3.6.1.4.1.2272.1.230.1.1.15	8.1.1
rcCloudIqLastOnboardTime	1.3.6.1.4.1.2272.1.230.1.1.16	8.1.1
rcCloudIqLastPollStatus	1.3.6.1.4.1.2272.1.230.1.1.17	8.1.1
rcCloudIqLastPollTime	1.3.6.1.4.1.2272.1.230.1.1.18	8.1.1
rcCloudIqLastMonitorStatus	1.3.6.1.4.1.2272.1.230.1.1.19	8.1.1
rcCloudIqLastMonitorTime	1.3.6.1.4.1.2272.1.230.1.1.20	8.1.1
rcCloudIqLastHealthStatus	1.3.6.1.4.1.2272.1.230.1.1.21	8.1.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcCloudIqLastHealthTime	1.3.6.1.4.1.2272.1.230.1.1.22	8.1.1
rcnCloudIqUpTrap	1.3.6.1.4.1.2272.1.21.0.357	8.1.1
rcnCloudIqDownTrap	1.3.6.1.4.1.2272.1.21.0.358	8.1.1

Table 40: VSP 4450 Series

Object Name	Object OID	New in VOSS Release
rcDigitalCertRelaxedMode	1.3.6.1.4.1.2272.1.222.1.1.1.13	7.1.3
rcDigitalCertPkcs12Password	1.3.6.1.4.1.2272.1.222.1.1.1.14	7.1.3
rcDigitalCertSanTable	1.3.6.1.4.1.2272.1.222.1.1.6	7.1.3
rcDigitalCertSanEntry	1.3.6.1.4.1.2272.1.222.1.1.6.1	7.1.3
rcDigitalCertSanType	1.3.6.1.4.1.2272.1.222.1.1.6.1.1	7.1.3
rcDigitalCertSanName	1.3.6.1.4.1.2272.1.222.1.1.6.1.2	7.1.3
rcDigitalCertSanRowStatus	1.3.6.1.4.1.2272.1.222.1.1.6.1.3	7.1.3

Table 41: VSP 4900 Series

Object Name	Object OID	New in VOSS Release
pethFastPoeEnable	1.3.6.1.2.1.105.1.3.1.1.6	8.1
pethPerpetualPoeEnable	1.3.6.1.2.1.105.1.3.1.1.7	8.1
bsnesEEEPortStatsEntry	1.3.6.1.4.1.45.5.34.1.5.1	8.1
bsnesEEEPortStatsPort	1.3.6.1.4.1.45.5.34.1.5.1.1	8.1
bsnesEEEPortStatsState	1.3.6.1.4.1.45.5.34.1.5.1.2	8.1
bsnesEEEPortStatsTxEvents	1.3.6.1.4.1.45.5.34.1.5.1.3	8.1
bsnesEEEPortStatsTxDuration	1.3.6.1.4.1.45.5.34.1.5.1.4	8.1
bsnesEEEPortStatsRxEvents	1.3.6.1.4.1.45.5.34.1.5.1.5	8.1
bsnesEEEPortStatsRxDuration	1.3.6.1.4.1.45.5.34.1.5.1.6	8.1
rcSysLocatorLED	1.3.6.1.4.1.2272.1.1.125	8.1
rcVossSystemVimAdminSpeed	1.3.6.1.4.1.2272.1.101.1.1.1.3	8.1
rcCloudIq	1.3.6.1.4.1.2272.1.230	8.1.1
rcCloudIqObjects	1.3.6.1.4.1.2272.1.230.1	8.1.1
rcCloudIqScalars	1.3.6.1.4.1.2272.1.230.1.1	8.1.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcCloudIqAgentEnable	1.3.6.1.4.1.2272.1.230.1.1.1	8.1.1
rcCloudIqAgentVersion	1.3.6.1.4.1.2272.1.230.1.1.2	8.1.1
rcCloudIqServerAddressType	1.3.6.1.4.1.2272.1.230.1.1.3	8.1.1
rcCloudIqServerAddress	1.3.6.1.4.1.2272.1.230.1.1.4	8.1.1
rcCloudIqProxyAddressType	1.3.6.1.4.1.2272.1.230.1.1.5	8.1.1
rcCloudIqProxyAddress	1.3.6.1.4.1.2272.1.230.1.1.6	8.1.1
rcCloudIqProxyTcpPort	1.3.6.1.4.1.2272.1.230.1.1.7	8.1.1
rcCloudIqProxyUserName	1.3.6.1.4.1.2272.1.230.1.1.8	8.1.1
rcCloudIqProxyPassword	1.3.6.1.4.1.2272.1.230.1.1.9	8.1.1
rcCloudIqNotificationEnable	1.3.6.1.4.1.2272.1.230.1.1.10	8.1.1
rcCloudIqOperStatus	1.3.6.1.4.1.2272.1.230.1.1.11	8.1.1
rcCloudIqAssociationUrl	1.3.6.1.4.1.2272.1.230.1.1.12	8.1.1
rcCloudIqPollUrl	1.3.6.1.4.1.2272.1.230.1.1.13	8.1.1
rcCloudIqMonitorFreq	1.3.6.1.4.1.2272.1.230.1.1.14	8.1.1
rcCloudIqPollFreq	1.3.6.1.4.1.2272.1.230.1.1.15	8.1.1
rcCloudIqLastOnboardTime	1.3.6.1.4.1.2272.1.230.1.1.16	8.1.1
rcCloudIqLastPollStatus	1.3.6.1.4.1.2272.1.230.1.1.17	8.1.1
rcCloudIqLastPollTime	1.3.6.1.4.1.2272.1.230.1.1.18	8.1.1
rcCloudIqLastMonitorStatus	1.3.6.1.4.1.2272.1.230.1.1.19	8.1.1
rcCloudIqLastMonitorTime	1.3.6.1.4.1.2272.1.230.1.1.20	8.1.1
rcCloudIqLastHealthStatus	1.3.6.1.4.1.2272.1.230.1.1.21	8.1.1
rcCloudIqLastHealthTime	1.3.6.1.4.1.2272.1.230.1.1.22	8.1.1
rcnCloudIqUpTrap	1.3.6.1.4.1.2272.1.21.0.357	8.1.1
rcnCloudIqDownTrap	1.3.6.1.4.1.2272.1.21.0.358	8.1.1

Table 42: VSP 7200 Series

Object Name	Object OID	New in VOSS Release
rcDigitalCertRelaxedMode	1.3.6.1.4.1.2272.1.222.1.1.1.13	7.1.3
rcDigitalCertPkcs12Password	1.3.6.1.4.1.2272.1.222.1.1.1.14	7.1.3
rcDigitalCertSanTable	1.3.6.1.4.1.2272.1.222.1.1.6	7.1.3
rcDigitalCertSanEntry	1.3.6.1.4.1.2272.1.222.1.1.6.1	7.1.3
rcDigitalCertSanType	1.3.6.1.4.1.2272.1.222.1.1.6.1.1	7.1.3

Table continues…

Object Name	Object OID	New in VOSS Release
rcDigitalCertSanName	1.3.6.1.4.1.2272.1.222.1.1.6.1.2	7.1.3
rcDigitalCertSanRowStatus	1.3.6.1.4.1.2272.1.222.1.1.6.1.3	7.1.3
rcEndpointTrackingAutoIsidOffset	1.3.6.1.4.1.2272.1.228.1.1.1.1	8.1
rcEndpointTrackingAutoIsidOffsetEnable	1.3.6.1.4.1.2272.1.228.1.1.1.2	8.1
rcEndpointTrackingGlobalEnable	1.3.6.1.4.1.2272.1.228.1.1.1.3	8.1
rcEndpointTrackingInterfaceTable	1.3.6.1.4.1.2272.1.228.1.1.2	8.1
rcEndpointTrackingInterfaceEntry	1.3.6.1.4.1.2272.1.228.1.1.2.1	8.1
rcEndpointTrackingInterfaceIndex	1.3.6.1.4.1.2272.1.228.1.1.2.1.1	8.1
rcEndpointTrackingInterfaceEnable	1.3.6.1.4.1.2272.1.228.1.1.2.1.2	8.1
rcEndpointTrackingInterfaceRowStatus	1.3.6.1.4.1.2272.1.228.1.1.2.1.3	8.1
rcEndpointTrackingBindingTable	1.3.6.1.4.1.2272.1.228.1.1.3	8.1
rcEndpointTrackingBindingEntry	1.3.6.1.4.1.2272.1.228.1.1.3.1	8.1
rcEndpointTrackingBindingIfIndex	1.3.6.1.4.1.2272.1.228.1.1.3.1.1	8.1
rcEndpointTrackingBindingMacAddr	1.3.6.1.4.1.2272.1.228.1.1.3.1.2	8.1
rcEndpointTrackingBindingStatus	1.3.6.1.4.1.2272.1.228.1.1.3.1.3	8.1
rcEndpointTrackingBindingVlanId	1.3.6.1.4.1.2272.1.228.1.1.3.1.4	8.1
rcEndpointTrackingBindingIsid	1.3.6.1.4.1.2272.1.228.1.1.3.1.5	8.1
rcEndpointTrackingBindingIsidSource	1.3.6.1.4.1.2272.1.228.1.1.3.1.6	8.1
rcEndpointTrackingBindingTimeout	1.3.6.1.4.1.2272.1.228.1.1.3.1.7	8.1
rcEndpointTrackingBindingTimeRemaining	1.3.6.1.4.1.2272.1.228.1.1.3.1.8	8.1

Table 43: VSP 7400 Series

Object Name	Object OID	New in VOSS Release
rcDigitalCertRelaxedMode	1.3.6.1.4.1.2272.1.222.1.1.1.13	8.0.6
rcDigitalCertPkcs12Password	1.3.6.1.4.1.2272.1.222.1.1.1.14	8.0.6
rcDigitalCertSanTable	1.3.6.1.4.1.2272.1.222.1.1.6	8.0.6
rcDigitalCertSanEntry	1.3.6.1.4.1.2272.1.222.1.1.6.1	8.0.6
rcDigitalCertSanType	1.3.6.1.4.1.2272.1.222.1.1.6.1.1	8.0.6
rcDigitalCertSanName	1.3.6.1.4.1.2272.1.222.1.1.6.1.2	8.0.6
rcDigitalCertSanRowStatus	1.3.6.1.4.1.2272.1.222.1.1.6.1.3	8.0.6
rcEndpointTrackingAutoIsidOffset	1.3.6.1.4.1.2272.1.228.1.1.1.1	8.1
rcEndpointTrackingAutoIsidOffsetEnable	1.3.6.1.4.1.2272.1.228.1.1.1.2	8.1
rcEndpointTrackingGlobalEnable	1.3.6.1.4.1.2272.1.228.1.1.1.3	8.1
rcEndpointTrackingInterfaceTable	1.3.6.1.4.1.2272.1.228.1.1.2	8.1
rcEndpointTrackingInterfaceEntry	1.3.6.1.4.1.2272.1.228.1.1.2.1	8.1
rcEndpointTrackingInterfaceIndex	1.3.6.1.4.1.2272.1.228.1.1.2.1.1	8.1
rcEndpointTrackingInterfaceEnable	1.3.6.1.4.1.2272.1.228.1.1.2.1.2	8.1
rcEndpointTrackingInterfaceRowStatus	1.3.6.1.4.1.2272.1.228.1.1.2.1.3	8.1
rcEndpointTrackingBindingTable	1.3.6.1.4.1.2272.1.228.1.1.3	8.1
rcEndpointTrackingBindingEntry	1.3.6.1.4.1.2272.1.228.1.1.3.1	8.1
rcEndpointTrackingBindingIfIndex	1.3.6.1.4.1.2272.1.228.1.1.3.1.1	8.1
rcEndpointTrackingBindingMacAddr	1.3.6.1.4.1.2272.1.228.1.1.3.1.2	8.1
rcEndpointTrackingBindingStatus	1.3.6.1.4.1.2272.1.228.1.1.3.1.3	8.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcEndpointTrackingBindingVlanId	1.3.6.1.4.1.2272.1.228.1.1.3.1.4	8.1
rcEndpointTrackingBindingIsid	1.3.6.1.4.1.2272.1.228.1.1.3.1.5	8.1
rcEndpointTrackingBindingIsidSource	1.3.6.1.4.1.2272.1.228.1.1.3.1.6	8.1
rcEndpointTrackingBindingTimeout	1.3.6.1.4.1.2272.1.228.1.1.3.1.7	8.1
rcEndpointTrackingBindingTimeRemaining	1.3.6.1.4.1.2272.1.228.1.1.3.1.8	8.1
rcCloudIq	1.3.6.1.4.1.2272.1.230	8.1.1
rcCloudIqObjects	1.3.6.1.4.1.2272.1.230.1	8.1.1
rcCloudIqScalars	1.3.6.1.4.1.2272.1.230.1.1	8.1.1
rcCloudIqAgentEnable	1.3.6.1.4.1.2272.1.230.1.1.1	8.1.1
rcCloudIqAgentVersion	1.3.6.1.4.1.2272.1.230.1.1.2	8.1.1
rcCloudIqServerAddressType	1.3.6.1.4.1.2272.1.230.1.1.3	8.1.1
rcCloudIqServerAddress	1.3.6.1.4.1.2272.1.230.1.1.4	8.1.1
rcCloudIqProxyAddressType	1.3.6.1.4.1.2272.1.230.1.1.5	8.1.1
rcCloudIqProxyAddress	1.3.6.1.4.1.2272.1.230.1.1.6	8.1.1
rcCloudIqProxyTcpPort	1.3.6.1.4.1.2272.1.230.1.1.7	8.1.1
rcCloudIqProxyUserName	1.3.6.1.4.1.2272.1.230.1.1.8	8.1.1
rcCloudIqProxyPassword	1.3.6.1.4.1.2272.1.230.1.1.9	8.1.1
rcCloudIqNotificationEnable	1.3.6.1.4.1.2272.1.230.1.1.10	8.1.1
rcCloudIqOperStatus	1.3.6.1.4.1.2272.1.230.1.1.11	8.1.1
rcCloudIqAssociationUrl	1.3.6.1.4.1.2272.1.230.1.1.12	8.1.1
rcCloudIqPollUrl	1.3.6.1.4.1.2272.1.230.1.1.13	8.1.1
rcCloudIqMonitorFreq	1.3.6.1.4.1.2272.1.230.1.1.14	8.1.1
rcCloudIqPollFreq	1.3.6.1.4.1.2272.1.230.1.1.15	8.1.1
rcCloudIqLastOnboardTime	1.3.6.1.4.1.2272.1.230.1.1.16	8.1.1
rcCloudIqLastPollStatus	1.3.6.1.4.1.2272.1.230.1.1.17	8.1.1
rcCloudIqLastPollTime	1.3.6.1.4.1.2272.1.230.1.1.18	8.1.1
rcCloudIqLastMonitorStatus	1.3.6.1.4.1.2272.1.230.1.1.19	8.1.1
rcCloudIqLastMonitorTime	1.3.6.1.4.1.2272.1.230.1.1.20	8.1.1
rcCloudIqLastHealthStatus	1.3.6.1.4.1.2272.1.230.1.1.21	8.1.1
rcCloudIqLastHealthTime	1.3.6.1.4.1.2272.1.230.1.1.22	8.1.1
rcnCloudIqUpTrap	1.3.6.1.4.1.2272.1.21.0.357	8.1.1
rcnCloudIqDownTrap	1.3.6.1.4.1.2272.1.21.0.358	8.1.1

Table 44: VSP 8000 Series

Object Name	Object OID	New in VOSS Release
rcDigitalCertRelaxedMode	1.3.6.1.4.1.2272.1.222.1.1.1.13	7.1.3
rcDigitalCertPkcs12Password	1.3.6.1.4.1.2272.1.222.1.1.1.14	7.1.3
rcDigitalCertSanTable	1.3.6.1.4.1.2272.1.222.1.1.6	7.1.3
rcDigitalCertSanEntry	1.3.6.1.4.1.2272.1.222.1.1.6.1	7.1.3
rcDigitalCertSanType	1.3.6.1.4.1.2272.1.222.1.1.6.1.1	7.1.3
rcDigitalCertSanName	1.3.6.1.4.1.2272.1.222.1.1.6.1.2	7.1.3
rcDigitalCertSanRowStatus	1.3.6.1.4.1.2272.1.222.1.1.6.1.3	7.1.3
rcPortMacsecMKAProfileName	1.3.6.1.4.1.2272.1.4.10.1.1.127	8.1
rcMACSecMKAProfileTable	1.3.6.1.4.1.2272.1.88.3	8.1
rcMACSecMKAProfileEntry	1.3.6.1.4.1.2272.1.88.3.1	8.1
rcMACSecMKAProfileId	1.3.6.1.4.1.2272.1.88.3.1.1	8.1
rcMACSecMKAProfileName	1.3.6.1.4.1.2272.1.88.3.1.2	8.1
rcMACSecMKAProfileReplayProtectEnable	1.3.6.1.4.1.2272.1.88.3.1.3	8.1
rcMACSecMKAProfileReplayProtectWindow	1.3.6.1.4.1.2272.1.88.3.1.4	8.1
rcMACSecMKAProfileOffsetValue	1.3.6.1.4.1.2272.1.88.3.1.5	8.1
rcMACSecMKAProfileRowStatus	1.3.6.1.4.1.2272.1.88.3.1.6	8.1
rcMACSecMKAProfilePortMembers	1.3.6.1.4.1.2272.1.88.3.1.7	8.1
rcMACSecMKAProfileCipherSuite	1.3.6.1.4.1.2272.1.88.3.1.8	8.1
rcMACSecMKAStatsTable	1.3.6.1.4.1.2272.1.88.4	8.1
rcMACSecMKAStatsEntry	1.3.6.1.4.1.2272.1.88.4.1	8.1
rcMACSecMKAMKPDUValidatedPkts	1.3.6.1.4.1.2272.1.88.4.1.1	8.1
rcMACSecMKARxDistributedSAKPkts	1.3.6.1.4.1.2272.1.88.4.1.2	8.1
rcMACSecMKAMKPDUTransmittedPkts	1.3.6.1.4.1.2272.1.88.4.1.3	8.1
rcMACSecMKATxDistributedSAKPkts	1.3.6.1.4.1.2272.1.88.4.1.4	8.1
rcMACSecMKAClearStats	1.3.6.1.4.1.2272.1.88.4.1.5	8.1
rcEndpointTrackingAutoIsidOffset	1.3.6.1.4.1.2272.1.228.1.1.1.1	8.1
rcEndpointTrackingAutoIsidOffsetEnable	1.3.6.1.4.1.2272.1.228.1.1.1.2	8.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcEndpointTrackingGlobalEnable	1.3.6.1.4.1.2272.1.228.1.1.1.3	8.1
rcEndpointTrackingInterfaceTable	1.3.6.1.4.1.2272.1.228.1.1.2	8.1
rcEndpointTrackingInterfaceEntry	1.3.6.1.4.1.2272.1.228.1.1.2.1	8.1
rcEndpointTrackingInterfaceIndex	1.3.6.1.4.1.2272.1.228.1.1.2.1.1	8.1
rcEndpointTrackingInterfaceEnable	1.3.6.1.4.1.2272.1.228.1.1.2.1.2	8.1
rcEndpointTrackingInterfaceRowStatus	1.3.6.1.4.1.2272.1.228.1.1.2.1.3	8.1
rcEndpointTrackingBindingTable	1.3.6.1.4.1.2272.1.228.1.1.3	8.1
rcEndpointTrackingBindingEntry	1.3.6.1.4.1.2272.1.228.1.1.3.1	8.1
rcEndpointTrackingBindingIfIndex	1.3.6.1.4.1.2272.1.228.1.1.3.1.1	8.1
rcEndpointTrackingBindingMacAddr	1.3.6.1.4.1.2272.1.228.1.1.3.1.2	8.1
rcEndpointTrackingBindingStatus	1.3.6.1.4.1.2272.1.228.1.1.3.1.3	8.1
rcEndpointTrackingBindingVlanId	1.3.6.1.4.1.2272.1.228.1.1.3.1.4	8.1
rcEndpointTrackingBindingIsid	1.3.6.1.4.1.2272.1.228.1.1.3.1.5	8.1
rcEndpointTrackingBindingIsidSource	1.3.6.1.4.1.2272.1.228.1.1.3.1.6	8.1
rcEndpointTrackingBindingTimeout	1.3.6.1.4.1.2272.1.228.1.1.3.1.7	8.1
rcEndpointTrackingBindingTimeRemaining	1.3.6.1.4.1.2272.1.228.1.1.3.1.8	8.1
ieee8021XPaeKaY	1.3.111.2.802.1.1.15.1.6	8.1
ieee8021XKayMkaTable	1.3.111.2.802.1.1.15.1.6.1	8.1
ieee8021XKayMkaEntry	1.3.111.2.802.1.1.15.1.6.1.1	8.1
ieee8021XKayMkaActive	1.3.111.2.802.1.1.15.1.6.1.1.1	8.1
ieee8021XKayMkaAuthenticated	1.3.111.2.802.1.1.15.1.6.1.1.2	8.1
ieee8021XKayMkaSecured	1.3.111.2.802.1.1.15.1.6.1.1.3	8.1
ieee8021XKayMkaFailed	1.3.111.2.802.1.1.15.1.6.1.1.4	8.1
ieee8021XKayMkaActorSCI	1.3.111.2.802.1.1.15.1.6.1.1.5	8.1

Table continues…

Object Name	Object OID	New in VOSS Release
ieee8021XKayMkaActorsPriority	1.3.111.2.802.1.1.15.1.6.1.1.6	8.1
ieee8021XKayMkaKeyServerPriority	1.3.111.2.802.1.1.15.1.6.1.1.7	8.1
ieee8021XKayMkaKeyServerSCI	1.3.111.2.802.1.1.15.1.6.1.1.8	8.1
ieee8021XKayAllowedJoinGroup	1.3.111.2.802.1.1.15.1.6.1.1.9	8.1
ieee8021XKayAllowedFormGroup	1.3.111.2.802.1.1.15.1.6.1.1.10	8.1
ieee8021XKayCreateNewGroup	1.3.111.2.802.1.1.15.1.6.1.1.11	8.1
ieee8021XKayMacSecCapability	1.3.111.2.802.1.1.15.1.6.1.1.12	8.1
ieee8021XKayMacSecDesired	1.3.111.2.802.1.1.15.1.6.1.1.13	8.1
ieee8021XKayMacSecProtect	1.3.111.2.802.1.1.15.1.6.1.1.14	8.1
ieee8021XKayMacSecReplayProtect	1.3.111.2.802.1.1.15.1.6.1.1.15	8.1
ieee8021XKayMacSecValidate	1.3.111.2.802.1.1.15.1.6.1.1.16	8.1
ieee8021XKayMacSecConfidentialityOffset	1.3.111.2.802.1.1.15.1.6.1.1.17	8.1
ieee8021XKayMkaTxKN	1.3.111.2.802.1.1.15.1.6.1.1.18	8.1
ieee8021XKayMkaTxAN	1.3.111.2.802.1.1.15.1.6.1.1.19	8.1
ieee8021XKayMkaRxKN	1.3.111.2.802.1.1.15.1.6.1.1.20	8.1
ieee8021XKayMkaRxAN	1.3.111.2.802.1.1.15.1.6.1.1.21	8.1
ieee8021XKayMkaParticipantTable	1.3.111.2.802.1.1.15.1.6.2	8.1
Ieee8021XKayMkaParticipantEntry	1.3.111.2.802.1.1.15.1.6.2.1	8.1
ieee8021XKayMkaPartCKN	1.3.111.2.802.1.1.15.1.6.2.1.1	8.1
ieee8021XKayMkaPartKMD	1.3.111.2.802.1.1.15.1.6.2.1.2	8.1
ieee8021XKayMkaPartNID	1.3.111.2.802.1.1.15.1.6.2.1.3	8.1
ieee8021XKayMkaPartCached	1.3.111.2.802.1.1.15.1.6.2.1.4	8.1
ieee8021XKayMkaPartActive	1.3.111.2.802.1.1.15.1.6.2.1.5	8.1
ieee8021XKayMkaPartRetain	1.3.111.2.802.1.1.15.1.6.2.1.6	8.1
ieee8021XKayMkaPartActivateControl	1.3.111.2.802.1.1.15.1.6.2.1.7	8.1
ieee8021XKayMkaPartPrincipal	1.3.111.2.802.1.1.15.1.6.2.1.8	8.1

Table continues…

Object Name	Object OID	New in VOSS Release
ieee8021XKayMkaPartDistCKN	1.3.111.2.802.1.1.15.1.6.2.1.9	8.1
ieee8021XKayMkaPartRowStatus	1.3.111.2.802.1.1.15.1.6.2.1.10	8.1
ieee8021XKayMkaPeerListTable	1.3.111.2.802.1.1.15.1.6.3	8.1
Ieee8021XKayMkaPeerListEntry	1.3.111.2.802.1.1.15.1.6.3.1	8.1
ieee8021XKayMkaPeerListMI	1.3.111.2.802.1.1.15.1.6.3.1.1	8.1
ieee8021XKayMkaPeerListMN	1.3.111.2.802.1.1.15.1.6.3.1.2	8.1
ieee8021XKayMkaPeerListType	1.3.111.2.802.1.1.15.1.6.3.1.3	8.1
ieee8021XKayMkaPeerListSCI	1.3.111.2.802.1.1.15.1.6.3.1.4	8.1

Table 45: XA1400 Series

Object Name	Object OID	New in VOSS Release
rcLicenseDisplayType	1.3.6.1.4.1.2272.1.56.13	8.0.50
rcIsisLogicalInterfaceIpsecEnable	1.3.6.1.4.1.2272.1.63.26.1.14	8.0.50
rcIsisLogicalInterfaceAuthenticationKey	1.3.6.1.4.1.2272.1.63.26.1.15	8.0.50
rcIsisLogicalInterfaceShapingRate	1.3.6.1.4.1.2272.1.63.26.1.16	8.0.50
rcNlsMgmtInterfaceTable	1.3.6.1.4.1.2272.1.223.1	8.1.1
rcNlsMgmtAddressTable	1.3.6.1.4.1.2272.1.223.2	8.1.1
rcNlsMgmtIpArpTable	1.3.6.1.4.1.2272.1.223.3	8.1.1
rcNlsMgmtIpStaticRouteTable	1.3.6.1.4.1.2272.1.223.5	8.1.1
rcNlsMgmtStatsTable	1.3.6.1.4.1.2272.1.223.7	8.1.1
rcNlsMgmtIpRouteTable	1.3.6.1.4.1.2272.1.223.8	8.1.1
rcCloudIq	1.3.6.1.4.1.2272.1.230	8.1.1
rcCloudIqObjects	1.3.6.1.4.1.2272.1.230.1	8.1.1
rcCloudIqScalars	1.3.6.1.4.1.2272.1.230.1.1	8.1.1
rcCloudIqAgentEnable	1.3.6.1.4.1.2272.1.230.1.1.1	8.1.1
rcCloudIqAgentVersion	1.3.6.1.4.1.2272.1.230.1.1.2	8.1.1
rcCloudIqServerAddressType	1.3.6.1.4.1.2272.1.230.1.1.3	8.1.1
rcCloudIqServerAddress	1.3.6.1.4.1.2272.1.230.1.1.4	8.1.1
rcCloudIqProxyAddressType	1.3.6.1.4.1.2272.1.230.1.1.5	8.1.1
rcCloudIqProxyAddress	1.3.6.1.4.1.2272.1.230.1.1.6	8.1.1
rcCloudIqProxyTcpPort	1.3.6.1.4.1.2272.1.230.1.1.7	8.1.1
rcCloudIqProxyUserName	1.3.6.1.4.1.2272.1.230.1.1.8	8.1.1
rcCloudIqProxyPassword	1.3.6.1.4.1.2272.1.230.1.1.9	8.1.1

Table continues…

Object Name	Object OID	New in VOSS Release
rcCloudIqNotificationEnable	1.3.6.1.4.1.2272.1.230.1.1.10	8.1.1
rcCloudIqOperStatus	1.3.6.1.4.1.2272.1.230.1.1.11	8.1.1
rcCloudIqAssociationUrl	1.3.6.1.4.1.2272.1.230.1.1.12	8.1.1
rcCloudIqPollUrl	1.3.6.1.4.1.2272.1.230.1.1.13	8.1.1
rcCloudIqMonitorFreq	1.3.6.1.4.1.2272.1.230.1.1.14	8.1.1
rcCloudIqPollFreq	1.3.6.1.4.1.2272.1.230.1.1.15	8.1.1
rcCloudIqLastOnboardTime	1.3.6.1.4.1.2272.1.230.1.1.16	8.1.1
rcCloudIqLastPollStatus	1.3.6.1.4.1.2272.1.230.1.1.17	8.1.1
rcCloudIqLastPollTime	1.3.6.1.4.1.2272.1.230.1.1.18	8.1.1
rcCloudIqLastMonitorStatus	1.3.6.1.4.1.2272.1.230.1.1.19	8.1.1
rcCloudIqLastMonitorTime	1.3.6.1.4.1.2272.1.230.1.1.20	8.1.1
rcCloudIqLastHealthStatus	1.3.6.1.4.1.2272.1.230.1.1.21	8.1.1
rcCloudIqLastHealthTime	1.3.6.1.4.1.2272.1.230.1.1.22	8.1.1
rcnCloudIqUpTrap	1.3.6.1.4.1.2272.1.21.0.357	8.1.1
rcnCloudIqDownTrap	1.3.6.1.4.1.2272.1.21.0.358	8.1.1

Obsolete MIBs

Table 46: Common

Object Name	Object OID	Obsolete in VOSS Release
rc2kCpuSerialPortDescr	1.3.6.1.4.1.2272.1.100.3.1.2	7.1
rc2kCpuSerialPortMode	1.3.6.1.4.1.2272.1.100.3.1.3	7.1
rc2kCpuSerialPortAdminStatus	1.3.6.1.4.1.2272.1.100.3.1.4	7.1
rc2kCpuSerialPortOperStatus	1.3.6.1.4.1.2272.1.100.3.1.5	7.1
rc2kCpuSerialPortDataBits	1.3.6.1.4.1.2272.1.100.3.1.7	7.1
rc2kCpuSerialPortMyAddr	1.3.6.1.4.1.2272.1.100.3.1.8	7.1
rc2kCpuSerialPortPeerAddr	1.3.6.1.4.1.2272.1.100.3.1.9	7.1
rc2kCpuSerialPortSlipMtu	1.3.6.1.4.1.2272.1.100.3.1.10	7.1
rc2kCpuSerialPortSlipTxRxCompress	1.3.6.1.4.1.2272.1.100.3.1.11	7.1
rc2kCpuSerialPortSlipRxCompress	1.3.6.1.4.1.2272.1.100.3.1.12	7.1
rc2kCpuSerialPortPppConfigFile	1.3.6.1.4.1.2272.1.100.3.1.13	7.1
rcIpConfBfdStaticFlag	1.3.6.1.4.1.2272.1.8.1.12.1.7	8.1
rcIpv6InterfaceBfdStaticFlag	1.3.6.14.1.2272.1.62.1.1.14.1.7	8.1

Table continues…

Object Name	Object OID	Obsolete in VOSS Release
rcVirtualServicePackageInfoName	1.3.6.1.4.1.2272.1.101.1.1.8.1.5	8.1
rcVirtualServicePackageInfoPath	1.3.6.1.4.1.2272.1.101.1.1.8.1.6	8.1
rcVirtualServicePackageAppName	1.3.6.1.4.1.2272.1.101.1.1.8.1.7	8.1
rcVirtualServicePackageAppVersion	1.3.6.1.4.1.2272.1.101.1.1.8.1.8	8.1