Platform Migration Reference for Avaya Virtual Services Platform 9000

Release 4.0
NN46250-107
Issue 05.02
July 2015

Notice

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

Documentation disclaimer

“Documentation” means information published by Avaya in varying mediums which may include product information, operating instructions and performance specifications that Avaya may generally make available to users of its products and Hosted Services. Documentation does not include marketing materials. Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of documentation unless such modifications, additions, or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

Link disclaimer

Avaya is not responsible for the contents or reliability of any linked websites referenced within this site or documentation provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

Warranty

Avaya provides a limited warranty on Avaya hardware and software. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya’s standard warranty language, as well as information regarding support for this product while under warranty is available to Avaya customers and other parties through the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Please note that if You acquired the product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to You by said Avaya Channel Partner and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, HTTP://SUPPORT.AVAYA.COM/LICENSEINFO OR SUCH SUCCESSOR SITE AS DESIGNATED BY AVAYA, ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AVAYA CHANNEL PARTNER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AVAYA CHANNEL PARTNER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA CHANNEL PARTNER; AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS “YOU” AND “END USER”), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE (“AVAYA”).

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. “Software” means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. “Designated Processor” means a single stand-alone computing device. “Server” means a Designated Processor that hosts a software application to be accessed by multiple users. “Instance” means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine (“VM”) or similar deployment.

Licence types

Designated System(s) License (DS). End User may install and use each copy or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

Heritage Nortel Software

“Heritage Nortel Software” means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at http://support.avaya.com/LicenseInfo under the link “Heritage Nortel Products” or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants You a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or for communication with Avaya equipment. Charges for Heritage Nortel Software may be based on extent of activation or use authorized as specified in an order or invoice.

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, and the product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components

“Third Party Components” mean certain software programs or portions thereof included in the Software or Hosted Service may contain software (including open source software) distributed under third party agreements (“Third Party Components”), which contain terms regarding the rights to use certain portions of the Software (“Third Party Terms”). As required, information regarding distributed Linux OS source code (for those products that have distributed Linux OS source code) and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply is available in the products, Documentation or on Avaya’s website at: http://support.avaya.com/Copyright or such successor site as designated by Avaya. You agree to the Third Party Terms for any such Third Party Components.

Preventing Toll Fraud

“Toll Fraud” is the unauthorized use of your telecommunications system by an unauthorized party (for example, a person who is not a corporate employee, agent, subcontractor, or is not working on your company's behalf). Be aware that there can be a risk of Toll Fraud associated with your system and that, if Toll Fraud occurs, it can result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

If You suspect that You are being victimized by Toll Fraud and You need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com or such successor site as designated by Avaya. Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com, or such successor site as designated by Avaya.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product or Hosted Service notices and articles, or to report a problem with your Avaya product or Hosted Service. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com (or such successor site as designated by Avaya), scroll to the bottom of the page, and select Contact Avaya Support.

Trademarks

The trademarks, logos and service marks (“Marks”) displayed in this site, the Documentation, Hosted Service(s), and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the Documentation, Hosted Service(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party.

Avaya is a registered trademark of Avaya Inc.

All non-Avaya trademarks are the property of their respective owners. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Contents

Chapter 1: Introduction

Purpose

This document describes the differences between the Avaya Ethernet Routing Switch 8000 series and the Avaya Virtual Services Platform 9000. You must read this document to understand the differences between the two products before you begin to migrate to the Virtual Services Platform 9000.

Intended audience

This document is intended for advanced administrators making the transition from the Avaya Ethernet Routing Switch 8000 series environment to the Avaya Virtual Services Platform 9000 environment.

Related resources

Documentation

See *Documentation Reference for Avaya Virtual Services Platform 9000,* NN46250-100 for a list of the documentation for this product.

Training

Ongoing product training is available. For more information or to register, you can access the website at http://avaya-learning.com/.

Course code	Course title
4D00010E	Knowledge Access: ACIS - Avaya ERS 8000 and VSP 9000 Implementation

Table continues…

Course code	Course title
5D00040E	Knowledge Access: ACSS - Avaya VSP 9000 Support

Viewing Avaya Mentor videos

Avaya Mentor videos provide technical content on how to install, configure, and troubleshoot Avaya products.

About this task

Videos are available on the Avaya Support website, listed under the video document type, and on the Avaya-run channel on YouTube.

Procedure

- To find videos on the Avaya Support website, go to http://support.avaya.com and perform one of the following actions:
  - In **Search**, type `Avaya Mentor Videos` to see a list of the available videos.
  - In **Search**, type the product name. On the Search Results page, select **Video** in the **Content Type** column on the left.
- To find the Avaya Mentor videos on YouTube, go to www.youtube.com/AvayaMentor and perform one of the following actions:
  - Enter a key word or key words in the **Search Channel** to search for a specific product or topic.
  - Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the website.

Note:

Videos are not available for all products.

Support

Go to the Avaya Support website at http://support.avaya.com for the most up-to-date documentation, product notices, and knowledge articles. You can also search for release notes, downloads, and resolutions to issues. Use the online service request system to create a service request. Chat with live agents to get answers to questions, or request an agent to connect you to a support team if an issue requires additional expertise.

Searching a documentation collection

On the Avaya Support website, you can download the documentation library for a specific product and software release to perform searches across an entire document collection. For example, you can perform a single, simultaneous search across the collection to quickly find all occurrences of a particular feature. Use this procedure to perform an index search of your documentation collection.

Before you begin

- Download the documentation collection zip file to your local computer.
- You must have Adobe Acrobat or Adobe Reader installed on your computer.

Procedure

1. Extract the document collection zip file into a folder.
2. Navigate to the folder that contains the extracted files and open the file named <*product_name_release*>.pdx.
3. In the Search dialog box, select the option **In the index named <*product_name_release*>.pdx**.
4. Enter a search word or phrase.
5. Select any of the following to narrow your search:
   - Whole Words Only
   - Case-Sensitive
   - Include Bookmarks
   - Include Comments
6. Click **Search**.

   The search results show the number of documents and instances found. You can sort the search results by Relevance Ranking, Date Modified, Filename, or Location. The default is Relevance Ranking.

Navigation

Use the following table to navigate this document and identify the differences that apply to your current Ethernet Routing Switch 8000 series configuration.

Table 1: Document navigation

Difference	Document location
Hardware	

Table continues…

Difference	Document location
VSP 9000 uses different power supplies.	Power requirements on page 13.
The VSP 9012 and the VSP 9010AC are both larger and weigh more than the ERS 8000 series.	Lift the chassis on page 15. Chassis measurements on page 17.
VSP 9000 has specific module installation and handling requirements.	Module requirements on page 18. Protecting modules on page 21.
VSP 9000 supports different SFP transceivers than ERS 8000 series. VSP 9000 uses SFP+ transceivers rather than XFP.	Optical transceivers on page 21.
Platform functionality and system handling	
VSP 9000 introduces ACLI differences, software patching, new features to monitor health, alarms, and failure detection and recovery.	Quick reference on page 33.
Software	
Specific protocols and features are not supported on both products.	Key software differences on page 55.
Support for spanning tree protocols is different.	Spanning tree on page 56.
VLAN classification and port membership in VLANs is different.	VLANs on page 56.
SMLT and SLT configuration is different.	Link aggregation and loop prevention on page 60.
VSP 9000 uses ACLs to assign MAC or VLAN QoS levels, and does not use egress queue sets.	QoS and traffic filters on page 63.
Attributes and operators for traffic filtering are different.	QoS and traffic filters on page 63.
VSP 9000 supports Layer 3 remote mirroring.	Remote mirroring on page 77.
VSP 9000 supports mixed-AS peer communication for BGP.	IP routing on page 77.
IPv6 routing support is different.	IPv6 routing on page 78.
VSP 9000 provides full IGMPv3 support.	IP multicast on page 78.
VSP 9000 and the ERS 8000 series both support Shortest Path Bridging MAC (SPBM); however, small configuration differences exist. VSP 9000 does not support IP VPN Lite over SPBM.	Shortest Path Bridging MAC on page 79.

Chapter 2: New in this release

The following sections describe what is new in *Platform Migration Reference for Avaya Virtual Services Platform 9000,* NN46250-107, for Release 4.0.

Features

See the following sections for information about feature-related changes.

9012QQ-2 I/O module

Release 4.0.1 introduces a second generation 9012QQ-2 Input/Output (I/O) module. The 9012QQ-2 module is a 12-port 40- gigabit-per-second (Gbps) module that supports the 40GBASE-R QSFP+ transceivers. You can use second generation I/O modules in first generation mode or second generation mode.

For more information about the 9012QQ-2 I/O module, see:

- 9012QQ-2 I/O module on page 31.

For more information about the 9012QQ-2 module specifications, see *Installing Modules in Avaya Virtual Services Platform 9000,* NN46250-301.

9048XS-2 I/O module

Release 4.0 introduces a second generation 9048XS-2 Input/Output (I/O) module.

The 9048XS-2 module is a 48-port 10-Gigabit-per-second (Gbps) module that supports both 10GBASE-R small form factor pluggable plus (SFP+) and 1000BASE-X SFP transceivers. You can use second generation I/O modules in first generation mode or second generation mode. For more information, see:

- 9048XS-2 I/O module on page 21.
- Shortest Path Bridging MAC on page 79.
- Module requirements on page 18.

IPv4 routes

Release 4.0 improves scaling of FIB IPv4 routes to one million, if you use second generation modules in second generation module mode with a premier license. For more software scaling information, see Software scaling comparison on page 86.

New parameter for the reset command

Release 4.0 adds the **`-coredump`** parameter to the **`reset`** command, which allows you to create a coredump for the main process before the switch resets.

⚠ **Caution:**

Only use the -coredump parameter if an issue causes you to reset the switch, and you need to contact customer service for analysis of the problem.

For more information, see Resetting the platform on page 41.

New parameter for the sys action cpu-switch-over command

Release 4.0 adds the **`-coredump`** parameter to the **`sys action cpu-switch-over`** command, which allows you to create a coredump for the main process before the switch changes to the backup CPU.

⚠ **Caution:**

Only use the -coredump parameter if an issue causes you to switchover the switch, and you need to contact customer service for analysis of the problem.

For more information, see Restarting the platform on page 43.

Software scaling update

Release 4.0.1 updates software scaling information in relation to IPv6 support. For more information, see: Software scaling comparison on page 86.

Other changes

See the following sections for other changes to the documentation.

Document title

Release 4.0 updates the document title to *Platform Migration Reference for Avaya Virtual Services Platform 9000*, NN46250–107, from *Avaya Virtual Services Platform 9000 Platform Migration*, NN46250–107.

Chapter 3: Hardware considerations

This section contains information on the migration considerations that pertain to the Avaya Virtual Services Platform 9000 hardware.

Quick reference

The following table provides a quick reference for the hardware differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series. For more detail about the Virtual Services Platform 9000 implementation, see the sections that follow.

Table 2: Hardware quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
ERS 8000 uses AC and DC power supplies	The VSP 9000 platform has the VSP 9012 chassis and the VSP 9010AC chassis. The VSP 9000 only supports AC power in this release.
–	VSP 9000 uses a VSP 9010AC chassis, a VSP 9012 chassis, and new modules.
ERS 8000 uses combined SF/CPU modules	VSP 9000 uses separate SF and CP modules with specific requirements for each. You must use the 9095SF module in the VSP 9010AC chassis and the 9090SF module in the VSP 9012 chassis. The LED lights in the 9080CP module map differently for the VSP 9012 chassis and the VSP 9010AC chassis.
No requirement exists for access to the back of the chassis.	While you insert the power supplies at the front of the chassis, the power cords connect to the back of the chassis. SF modules install at the back of the chassis.
Airflow for the chassis is front-to-back.	Airflow for the VSP 9010AC chassis is front-to-back. For the VSP 9010AC chassis, Avaya recommends 36 inches (in.) (91 centimeters [cm]) of free space in both the front and the back of the machine.

Table continues…

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
	Airflow for the VSP 9012 chassis is both front-to-back and side-to-side. For the VSP 9012 chassis, Avaya recommends 36 inches (in.) (91 centimeters [cm]) of free space in both the front and the back of the machine, and 6 in. (15.2 cm) on each side.
The out-of-band (OOB) Ethernet ports are 10/100 for the 8692 SF/CPU and 10/100/1000 for the 8895 SF/CPU.	The out-of-band (OOB) Ethernet ports are 10/100/1000.
ERS 8000 uses SFPs, XFPs, and GBICs.	VSP 9000 supports different SFPs. VSP 9000 does not support older non-DDI capable 1000BaseX SFPs and does not support GBICs. VSP 9000 uses SFP+ instead of XFP
A port functions even if the SFP or XFP are unsupported.	VSP 9000 cannot use unsupported SFP or SFP+. The port does not come up.
Autonegotiation is enabled by default on SFP modules, for example, 8630GBR, 8648GBRS, 8634XGRS, 8834XG, or 8848GB.	The 9024XL module supports 1000BaseX SFPs with autonegotiation disabled by default. You must enable auto-negotiation manually. The 9024XL ports are preconfigured for SFP+, where autonegotiation does not apply.

Virtual Services Platform 9012 power requirements

The Virtual Services Platform 9012 supports up to six 1200–2000 Watt AC power supplies. The Virtual Services Platform 9012 supports the 9006AC power supply.

The Virtual Services Platform 9012 does not support DC power supplies at this time.

Note:

All the power supplies must run the same voltage. Do not mix 120 and 220 voltages.

Power supply redundancy

You can operate the AC power supplies separately, or in parallel, or parallel redundant configurations. You can use the AC power supplies in one of the following redundant configurations:

- N+1 redundancy

  A single power supply failure or circuit breaker shutdown is backed up by the remaining supplies.

- N+N redundancy

  An n + n redundant configuration provides power in the event of a loss of a single power phase in the building. Balance the line side voltage source between building phases. Use n + n redundancy to ensure redundancy in the event that an external failure occurs; for example, an entire power feed within the building fails. To ensure n + n redundancy, you must install power supplies to provide twice the power requirements of your hardware configuration.

Important:

Avaya recommends that you install each power supply on its own dedicated branch circuit for electrical installation reasons

For more information on AC power supplies, see *Installing AC Power Supplies in Avaya Virtual Services Platform 9000,* NN46250-303.

To determine how many power supplies you need, you can download *ERS 8000 / VSP 9000 Power Supply Calculator*, NN48500–519 from the **System Management & Planning** section of the Virtual Services Platform 9000 product documentation at https://support.avaya.com.

For information about hardware and software compatibility, see *Release Notes for Avaya Virtual Services Platform 9000,* NN46250-401.

Virtual Services Platform 9010 power requirements

The Virtual Services Platform 9010AC supports up to eight 1200–2000 Watt AC power supplies. The Virtual Services Platform 9010AC supports the 9006AC power supply. The Virtual Services Platform 9010AC does not support DC power supplies at this time.

Important:

Avaya recommends that you install each power supply on its own dedicated branch circuit for electrical installation reasons.

Virtual Services Platform 9010AC

The Virtual Services Platform 9010AC supports up to eight 1200–2000 Watt AC power supplies. The nominal input voltage range is 100–120 VAC and 200–240 VAC; however, the output power is limited to 1200 W maximum at 100–120 VAC nominal input voltage conditions.

Power supply redundancy

You can operate the AC power supplies separately, or in parallel, or in parallel redundant configurations. You can use the AC power supplies in one of the following redundant configurations, where n is the number of required power supplies to power the chassis and modules:

- N+1 redundancy

  A single power supply failure or circuit breaker shutdown is backed up by the remaining supplies.

- N+N redundancy

  An n + n redundant configuration provides power in the event of a loss of a single power phase in the building. Balance the line side voltage source between building phases. Use n + n redundancy to ensure redundancy in the event that an external failure occurs; for example, an entire power feed within the building fails. To ensure n + n redundancy, you must install power supplies to provide twice the power requirements of your hardware configuration.

For more information on AC power supplies, see *Installing AC Power Supplies in Avaya Virtual Services Platform 9000,* NN46250-303.

To determine how many power supplies you need, you can download *ERS 8000 / VSP 9000 Power Supply Calculator*, NN48500–519 from the **System Management & Planning** section of the Virtual Services Platform 9000 product documentation at https://support.avaya.com.

For information about hardware and software compatibility, see *Release Notes for Avaya Virtual Services Platform 9000,* NN46250-401.

Lifting the Virtual Services Platform 9012

The Virtual Services Platform 9012 weighs in excess of 160 lb (73 kg). Each chassis requires a minimum of three people to lift. Always use a mechanical lift when one is available.

Before you begin

Important:

Reduce the weight of the chassis as much as possible before you lift it. Always use a mechanical lift when one is available. Ensure you have at least three people to lift the chassis. Use a third person to support the chassis from behind the rack, as you position the chassis on the shelf and hold it in place. Take care to lift the chassis from the bottom.

Procedure

Use the recessed handles at the top and bottom of the Virtual Services Platform 9012 sides to lift the chassis. From the rear of the Virtual Services Platform 9012, lift the chassis from the bottom only.

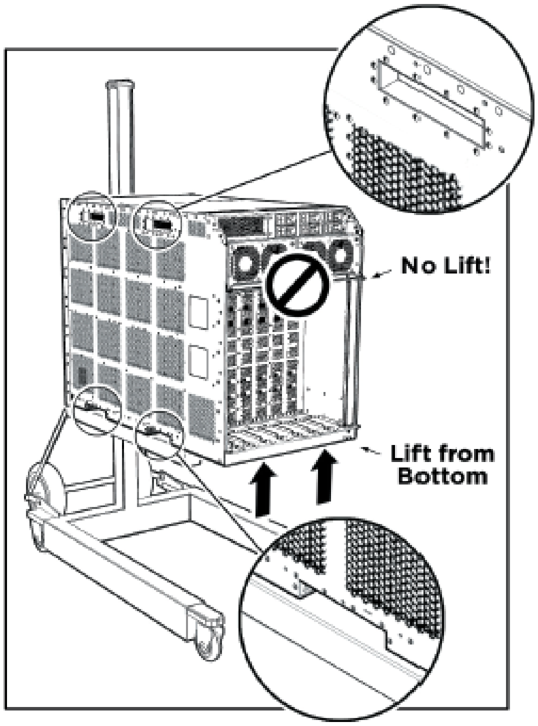

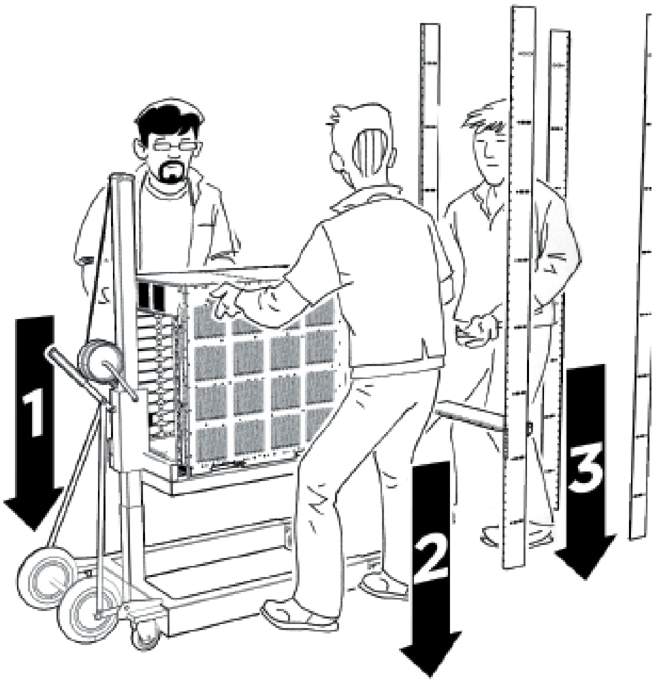

Lifting the chassis

The Virtual Services Platform 9010 weighs in excess of 141 lb (64 kg). Each chassis requires a minimum of three people to lift. Always use a mechanical lift when one is available.

Use the handles that swing out from the top and bottom of the chassis sides to lift the chassis. To use the handles, swing the handle up and out from the chassis. From the rear, lift the chassis from the bottom only.

Before you begin

Reduce the weight of the chassis as much as possible before you lift it. Always use a mechanical lift when one is available. Ensure you have at least three people to lift the chassis. Use two people to lift the chassis at the sides and a third person to support the chassis from behind the rack, as you position the chassis on the shelf and hold it in place. Take care to lift the chassis from the bottom.

Procedure

1. Use the recessed handles at the top and bottom of the chassis sides to lift the chassis. To use the handles, swing the handle up and out from the chassis.

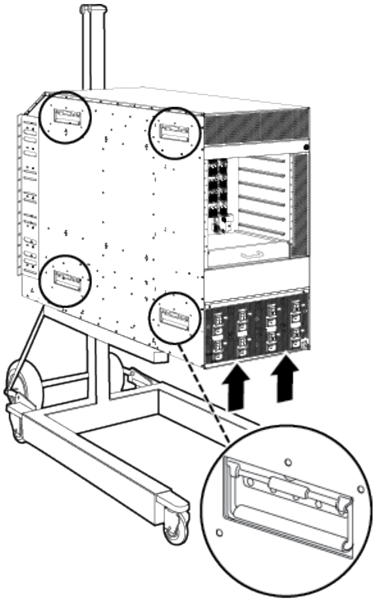

2. From the rear of the chassis, lift the chassis from the bottom only.

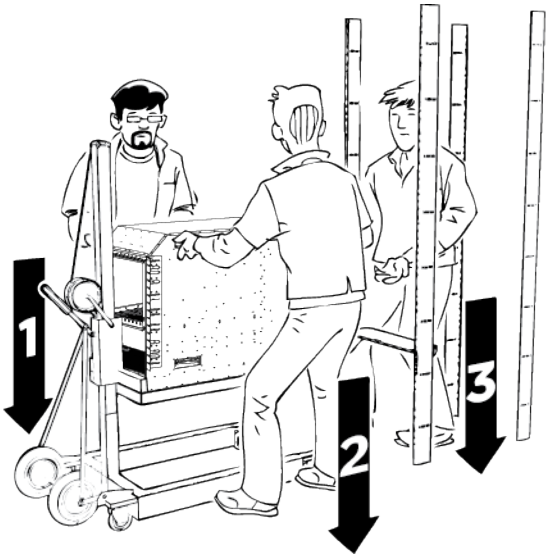

Chassis measurements

The Virtual Services Platform 9012 and Virtual Services Platform 9010 are different sizes than the Ethernet Routing Switch 8000 series chassis, particularly the depth of each chassis. The Virtual Services Platform 9012 and Virtual Services Platform 9010 are each deeper than the Ethernet Routing Switch 8000 series. You must install the chassis in a rack that meets the minimum depth requirements.

The following table provides a comparison of the chassis measurements.

Table 3: Chassis measurements

Measurement	Height	Width	Depth
VSP 9012	24.375 in. (61.91 cm)	17.5 in. (44.45 cm)	32.5 in. (82.55 cm)
VSP 9010	36 in. (91.4 cm)	17.5 in. (44.5 cm)	33.1 in. (84.1)
ERS 8010co	35 in. (88.9 cm)	19 in (48.26 cm)	23.7 in. (60.19 cm)
ERS 8010	22.9 in. (58.2 cm)	22.9 in. (58.2 cm)	19.9 in. (50.5 cm)
ERS 8006	15.8 in.(40.1 cm)	17.5 in. (44.5 cm)	19.9 in. (50.5 cm)
ERS 8803–R	12.25 in. (31.1 cm)	19 in. (48.3 cm)	21 in. (53.3 cm)

Module requirements

Two modules, the Control Processor (CP) module and the Switch Fabric (SF) module, have specific requirements that you must follow.

9080CP module

The 9080CP module requires an external Compact Flash card. This card is not optional.

You can hot swap the external storage devices but you must follow a specific procedure to prevent data loss or hardware damage. For more information, see Removing external storage devices from the CP module on page 19.

SF modules

The Virtual Services Platform 9010AC supports the 9095SF module. The Virtual Services Platform 9012 supports the 9090SF module.

If you install a second generation module in Virtual Services Platform 9010 or Virtual Services Platform 9012, you must have a minimum of five SF modules installed. Populate slots SF1 and SF4, and use any other slots for the remaining three SF modules.

If you install a first generation module in Virtual Services Platform 9010 or Virtual Services Platform 9012, you must have a minimum of three SF modules installed. Populate slots SF1 and SF4, and use any other slot for the remaining SF module.

Releases prior to 3.1 require a minimum of four SF modules.

Note:

Do not leave slots open. Fill all slots with modules or filler modules to maintain safety compliance, proper cooling, and EMI containment.

Take care when you insert the SF modules into the chassis because you can misalign the modules, and potentially damage the connectors in the chassis.

Tip:

It is easier to install the SF modules in the Virtual Services Platform 9012 from left to right because it is easier to compress each vertical seal on the face of the module as you install one after the other.

Removing external storage devices from the CP module

Perform this procedure to safely remove the USB and the external Compact Flash devices from the CP module. You must perform this procedure to prevent data loss or hardware damage.

Important:

Do not unplug the storage device without first performing this procedure.

You must use the appropriate stop command to unmount the device before you physically remove it from the CP module.

Before you begin

Several system tools use the external Compact Flash as the default storage location. Check the following features before you remove the card:

- Packet Capture (PCAP)
- logging
- debug or trace

The Virtual Services Platform 9000 stop command does not succeed if the specified device is in use. Common uses that impede the proper execution of the stop command are:

- USB or external Compact Flash file access is in progress (move, copy, read, or write) to or from the USB, or the external Compact Flash.

  Discontinue operations or wait for access completion before you use the stop command.
- The ACLI session current working directory is configured for the device you need to remove.

  Change the current working directory to internal Compact Flash, which is the default.
- Logging is enabled to the external Compact Flash, which is the default.

  Use the **show logging config** command to verify the current storage location. If the location is the external Compact Flash card that you need to remove, use the **no logging logToExtFlash** command to log to the internal Compact Flash.
- PCAP is enabled.

  Disable PCAP, which requires the external Compact Flash. Use the **show pcap** command to verify if PCAP is enabled. To disable PCAP, use the **no pcap enable** command.
- Debugging features are enabled.

  The debug-config file and trace-logging flags must be disabled, which is the default. Use the **show boot config flags** command to verify the status. Use the **no boot config**

flags debug-config file or the **no boot config flags trace-logging** command to disable these flags.

About this task

Note:

Use the Avaya Compact Flash device (EC1411010-E6) with the Virtual Services Platform 9000 because the Avaya Compact Flash is validated for proper operation on the Virtual Services Platform 9000. Do not use other Compact Flash devices because they are not verified for Virtual Services Platform 9000 compatibility, and can result in loss of access to the Compact Flash device.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Remove a USB device:

   a. Unmount the USB device:

      ```
      usb-stop
      ```

   b. Wait for the response that indicates it is safe to remove the device.

   c. Physically remove the device.

3. Remove an external Compact Flash device:

   a. Unmount the external flash device:

      ```
      extflash-stop
      ```

   b. Wait for the response that indicates it is safe to remove the device.

   c. Physically remove the device.

Example

Unmount and remove the USB:

```
VSP-9012:1>enable
VSP-9012:1#usb-stop
It is now safe to remove the USB device.
VSP-9012:1#extflash-stop
It is now safe to remove the external Compact Flash device.
```

Next steps

No restrictions or requirements exist before you can reinsert a USB or external Compact Flash device. You can insert these devices at any time and Virtual Services Platform 9000 automatically recognizes them. The devices are accessible within seconds after insertion.

After you insert the external Compact Flash, enable logging to the external Compact Flash with the **logging logToExtFlash** command.

Additionally, you can enable the following features as required:

- PCAP
- debug-config file or trace-logging flags

Protecting modules

Virtual Services Platform 9000 modules are larger and heavier than Ethernet Routing Switch 8000 series modules.

Handle the modules used in Virtual Services Platform 9000 with care. Take the following items into consideration when you handle modules:

- To prevent damage from electrostatic discharge, always wear an antistatic wrist strap connected to an ESD jack when you connect cables or you perform maintenance on this device.
- Always place the modules on appropriate antistatic material.
- Support the module from underneath with two hands. Do not touch components or connector pins with your hand, or damage can result.
- Damage to a module can occur if you bump the module into another object, including other modules installed in a chassis. Be careful not to bump module connectors against the action levers of an adjacent module. Damage to connectors can result. Use both hands to support modules.
- Visually inspect the connectors for damage before you insert the module. If you insert a module with damaged connectors you will damage the midplane.
- Check the clearance between the insertion lever and the gasket on adjacent modules during insertion or extraction.
- Do not stack modules one on top of the other when you move them.
- Do not leave slots open. Fill all slots with modules or filler modules to maintain safety compliance, proper cooling, and EMI containment.
- Do not over tighten screws. Tighten until snug. Do not use a power tool to tighten screws.

Optical transceivers

Virtual Services Platform 9000 interface modules support 1 Gb and 10 Gb optical transceivers. The 1 Gb transceiver is a small form factor pluggable (SFP) transceiver. The 10 Gb transceiver is an SFP+ transceiver. Virtual Services Platform 9000 does not support XFPs or gigabit interface converters (GBIC).

9048XS-2 I/O module

The second generation 9048XS-2 Input/Output (I/O) module is a 48 port 10 Gigabit per second (Gbps) module. The 9048XS-2 module supports the 10GBASE-R small form-factor pluggable plus (SFP+) transceivers and the 1000BASE-X SFP transceivers. The Virtual Services Platform 9000 supports the 9048XS-2 module in first generation mode and second generation mode. The Virtual Services Platform 9012 requires the 9012FCHS I/O cooling module to be installed before you install the 9048XS-2 module. You must also have a minimum of five Switch Fabric modules installed, if you install the 9048XS-2 module on the Virtual Services Platform 9012. Populate slots SF1 and SF4, and you can use any other slots for the remaining three SF modules.

This module supports standard management information base (MIB).

Note:

The 9048XS-2 module does not support Lossless Ethernet.

The 9048XS-2 module is oversubscribed 2:1, with full QoS awareness, with regards to line rate over 48 ports of 10 Gbps Ethernet traffic using standard SFP+ fiber transceivers. This module supports a maximum throughput of 357 Million packets per second (Mpps) over 48 ports of 10 Gbps Ethernet traffic using standard SFP+ fiber transceivers. The module supports SR, LR, LRM, ER, and ZR SFP + format.

The following tables provide the multimode fiber (MMF), single-mode fiber (SMF), and copper SFP and SFP+ fiber transceivers that the 9048XS-2 module supports.

Important:

Virtual Services Platform 9000 supports only Avaya-qualified transceivers. Other vendor transceivers will not work and Avaya does not support them.

Table 4: Supported SFP transceivers

Model number	Part number	Description
1000BASE-SX DDI SFP	AA1419048-E6	850 nm, Gigabit Ethernet, duplex LC connector
1000BASE-LX DDI	AA1419049-E6	1310 nm, up to 10 km
1000BASE-XD DDI	AA1419050-E6	1310 nm. The range is up to 40 km over SMF pair. This transceiver has been discontinued but remains supported by the software.
	AA1419051-E6	1550 nm (non-CWDM). The range is up to 40 km over SMF pair. This transceiver has been discontinued but remains supported by the software. Avaya recommends AA1419057-E6 as a replacement.
1000BASE-ZX DDI	AA1419052-E6	1550 nm (non-CWDM). The range is up to 70 km over SMF pair. This transceiver has been discontinued but remains supported by the software. Avaya recommends AA1419065-E6 as a replacement.

Table continues…

Model number	Part number	Description
1000BASE-BX-U-10	AA1419069-E6	Transmits at 1310 nm. The range is up to 10km upstream.
1000BASE-BX-D-10	AA1419070-E6	Transmits at 1490 nm. The range is up to 10 km downstream.
1000BASE-BX-U-40	AA1419076-E6	Transmits at 1310 nm. The range is up to 40 km upstream.
1000BASE-BX-D-40	AA1419077-E6	Transmits at 1490 nm. The range is up to 40 km downstream.
1000BASE-EX DDI	AA1419071-E6	1550 nm, up to 120 km (non-CWDM)
1000BASE DDI CWDM	AA1419053-E6	1470 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419054-E6	1490 (CWDM). The range is up to 40km over SMF pair.
	AA1419055-E6	1510 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419056-E6	1530 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419057-E6	1550 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419058-E6	1570 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419059-E6	1590 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419060-E6	1610 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419061-E6	1470 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419062-E6	1490 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419063-E6	1510 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419064-E6	1530 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419065-E6	1550 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419066-E6	1570 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419067-E6	1590 nm (CWDM). The range is up to 70km over SMF pair.

Table continues…

Model number	Part number	Description
	AA1419068-E6	1610 nm (CWDM). The range is up to 70km over SMF pair.
1000BASE-T	AA1419043-E6	CAT5 UTP, up to 100 m. Because the 1000BASE-T device is all electrical, it does not need DDI support.

Table 5: Supported SFP+ transceivers and cables

Model number	Part number	Description
10GBASE-CX	AA1403018-E6 to AA1403021-E6	4-pair twinaxial copper cable to connect 10 Gb ports. The maximum range is 15 m.
10GBASE-ER/EW	AA1403013-E6	1550 nm SMF. The range is up to 40 km.
10GBASE-ER CWDM DDI	AA1403153-E6	1471 nm SMF. The range is up to 40 km.
	AA1403154-E6	1491 nm SMF. The range is up to 40 km.
	AA1403155-E6	1511 nm SMF. The range is up to 40 km.
	AA1403156-E6	1531 nm SMF. The range is up to 40 km.
	AA1403157-E6	1551 nm SMF. The range is up to 40 km.
	AA1403158-E6	1571 nm SMF. The range is up to 40 km.
	AA1403159-E6	1591 nm SMF. The range is up to 40 km.
	AA1403160-E6	1611 nm SMF. The range is up to 40 km.
10GBASE-LR/LW	AA1403011-E6	1310 nm SMF. The range is up to 10 km.
10GBASE-LRM	AA1403017-E6	1310 nm. Up to 220 m reach over Fiber Distributed Data Interface (FDDI)-grade 62.5 µm multimode fiber. Suited for campus LANs.
10GBASE-SR/SW	AA1403015-E6	850 nanometers (nm). The range is up to the following: • 26 m using 62.5 micrometer (µm), 160 megaHertz times km (MHz-km) MMF

Table continues…

Model number	Part number	Description
		• 33 m using 62.5 µm, 200 MHz-km MMF • 66 m using 62.5 µm, 400 MHz-km MMF • 82 m using 50 µm, 500 MHz-km MMF • 300 m using 50 µm, 2000 MHz-km MMF • 400 m using 50 µm, 4700 MHz-km MMF (OM4)
10GBASE-ZR/ZW	AA1403016-E6	1550 nm SMF. The range is up to 70 km.
10GBASE-ZR CWDM DDI	AA1403161-E6	1471 nm SMF. The range is up to 70 km.
	AA1403162-E6	1491 nm SMF. The range is up to 70 km.
	AA1403163-E6	1511 nm SMF. The range is up to 70 km.
	AA1403164-E6	1531 nm SMF. The range is up to 70 km.
	AA1403165-E6	1551 nm SMF. The range is up to 70 km.
	AA1403166-E6	1571 nm SMF. The range is up to 70 km.
	AA1403167-E6	1591 nm SMF. The range is up to 70 km.
	AA1403168-E6	1611 nm SMF. The range is up to 70 km.

The 9048XS-2 I/O module has a dual core P2020 processor and 2 GB onboard DDR3 memory. You can use the 9048XS-2 module in both the Virtual Services Platform 9010 and Virtual Services Platform 9012 chassis.

The 9048XS-2 has the following characteristics:

- compliant with IEEE 802.3ae standards
- 802.3 Ethernet frame format, MAC layer functionality
- 64B/66B line encoding
- asynchronous Ethernet interface

9024XL I/O module

The 9024XL I/O module is a 24 port 10 gigabits per second (Gbps) small form-factor pluggable plus (SFP+) I/O module.

The module has approximately a 3.5:1 oversubscribed line rate over 24 ports of 10 Gbps Ethernet traffic using standard SFP+ fiber transceivers. Each continuous physical group of 4 ports supports a combined bandwidth of 11.3GE. Use only a single port for each grouping to ensure no oversubscription. As a helpful guide the last port in each group has a black mark on the faceplate.

The module supports a maximum throughput of 105 Mpps over 24 ports of 10 Gbps Ethernet traffic using standard SFP+ fiber transceivers. The module supports SR, LR, LRM, ER, and ZR SFP+ transceivers.

The following tables provide the multimode fiber (MMF), single-mode fiber (SMF), and copper SFP and SFP+ fiber transceivers that the 9024XL module supports.

Important:

Virtual Services Platform 9000 supports only Avaya-qualified transceivers. Other vendor transceivers will not work and Avaya does not support them.

Table 6: Supported SFP transceivers

Model	Part number	Description
1000BASE-XD DDI	AA1419050-E6	1310 nm. The range is up to 40 km over SMF pair. This transceiver has been discontinued but remains supported by the software.
	AA1419051-E6	1550 nm (non-CWDM). The range is up to 40 km over SMF pair. This transceiver has been discontinued but remains supported by the software. Avaya recommends AA1419057-E6 as a replacement.
1000BASE-ZX DDI	AA1419052-E6	1550 nm (non-CWDM). The range is up to 70 km over SMF pair. This transceiver has been discontinued but remains supported by the software. Avaya recommends AA1419065-E6 as a replacement.
1000BASE-BX-U-10	AA1419069-E6	Transmits at 1310 nm. The range is up to 10km upstream.

Table continues…

Model	Part number	Description
1000BASE-BX-D-10	AA1419070-E6	Transmits at 1490 nm. The range is up to 10 km downstream.
1000BASE-BX-U-40	AA1419076-E6	Transmits at 1310 nm. The range is up to 40 km upstream.
1000BASE-BX-D-40	AA1419077-E6	Transmits at 1490 nm. The range is up to 40 km downstream.
1000BASE-EX DDI	AA1419071-E6	1550 nm, up to 120 km (non-CWDM)
1000BASE DDI CWDM	AA1419053-E6	1470 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419054-E6	1490 (CWDM). The range is up to 40km over SMF pair.
	AA1419055-E6	1510 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419056-E6	1530 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419057-E6	1550 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419058-E6	1570 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419059-E6	1590 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419060-E6	1610 nm (CWDM). The range is up to 40km over SMF pair.
	AA1419061-E6	1470 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419062-E6	1490 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419063-E6	1510 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419064-E6	1530 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419065-E6	1550 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419066-E6	1570 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419067-E6	1590 nm (CWDM). The range is up to 70km over SMF pair.
	AA1419068-E6	1610 nm (CWDM). The range is up to 70km over SMF pair.

Table continues…

Model	Part number	Description
1000BASE-T	AA1419043-E6	CAT5 UTP, up to 100 m. Because the 1000BASE-T device is all electrical, it does not need DDI support.

Table 7: Supported SFP+ transceivers and cables

Model number	Part number	Description
10GBASE-CX	AA1403018-E6 to AA1403021-E6	4-pair twinaxial copper cable to connect 10 Gb ports. The maximum range is 15 m.
10GBASE-ER/EW	AA1403013-E6	1550 nm SMF. The range is up to 40 km.
10GBASE-ER CWDM DDI	AA1403153-E6	1471 nm SMF. The range is up to 40 km.
	AA1403154-E6	1491 nm SMF. The range is up to 40 km.
	AA1403155-E6	1511 nm SMF. The range is up to 40 km.
	AA1403156-E6	1531 nm SMF. The range is up to 40 km.
	AA1403157-E6	1551 nm SMF. The range is up to 40 km.
	AA1403158-E6	1571 nm SMF. The range is up to 40 km.
	AA1403159-E6	1591 nm SMF. The range is up to 40 km.
	AA1403160-E6	1611 nm SMF. The range is up to 40 km.
10GBASE-LR/LW	AA1403011-E6	1310 nm SMF. The range is up to 10 km.
10GBASE-LRM	AA1403017-E6	1310 nm. Up to 220 m reach over Fiber Distributed Data Interface (FDDI)-grade 62.5 µm multimode fiber. Suited for campus LANs.
10GBASE-SR/SW	AA1403015-E6	850 nanometers (nm). The range is up to the following: • 26 m using 62.5 micrometer (µm), 160 megaHertz times km (MHz-km) MMF • 33 m using 62.5 µm, 200 MHz-km MMF

Table continues…

Model number	Part number	Description
		• 66 m using 62.5 µm, 400 MHz-km MMF • 82 m using 50 µm, 500 MHz-km MMF • 300 m using 50 µm, 2000 MHz-km MMF • 400 m using 50 µm, 4700 MHz-km MMF (OM4)
10GBASE-ZR/ZW	AA1403016-E6	1550 nm SMF. The range is up to 70 km.
10GBASE-ZR CWDM DDI	AA1403161-E6	1471 nm SMF. The range is up to 70 km.
	AA1403162-E6	1491 nm SMF. The range is up to 70 km.
	AA1403163-E6	1511 nm SMF. The range is up to 70 km.
	AA1403164-E6	1531 nm SMF. The range is up to 70 km.
	AA1403165-E6	1551 nm SMF. The range is up to 70 km.
	AA1403166-E6	1571 nm SMF. The range is up to 70 km.
	AA1403167-E6	1591 nm SMF. The range is up to 70 km.
	AA1403168-E6	1611 nm SMF. The range is up to 70 km.

The 9024XL I/O module has a 1 GHz 8584E processor and 1 GB onboard DDR2 memory.

You can use the 9024XL module in both the Virtual Services Platform 9010 and Virtual Services Platform 9012 chassis.

The 9024XL has the following characteristics:

- compliant with IEEE 802.3ae standards
- 802.3 Ethernet frame format, MAC layer functionality
- 64B/66B line encoding
- asynchronous Ethernet interface

9048GB interface module

The 9048GB interface module is a 48 port 1 Gbps small form-factor pluggable (SFP) interface module that supports multimode fiber (MMF), single-mode fiber (SMF), and copper connections.

The following table details the SFP transceivers supported by the 9048GB module.

Important:

Virtual Services Platform 9000 supports only Avaya-qualified transceivers. Other vendor transceivers will not work and Avaya does not support them.

Table 8: Supported SFP transceivers

Model	ROHS product number	Description
1000BASE-T	AA1419043-E6	CAT5 UTP, up to 100 m. Because the 1000BASE-T device is all electrical, it does not need DDI support.
1000BASE-SX DDI	AA1419048-E6	850 nm up to 275 m using 62.5 m MMF optic cable up to 550 m using 50 µm MMF optic cable
1000BASE-LX DDI	AA1419049-E6	1310 nm, up to 10 km
1000BASE-XD DDI	AA1419050-E6	1310 nm, up to 40 km This transceiver has been discontinued but remains supported by the software.
	AA1419051-E6	1550 nm, up to 40km (non-CWDM) This transceiver has been discontinued but remains supported by the software. Avaya recommends AA1419057-E6 as a replacement.
1000BASE-ZX DDI	AA1419052-E6	1550 nm, up to 70 km (non-CWDM) This transceiver has been discontinued but remains supported by the software. Avaya recommends AA1419065-E6 as a replacement.
1000BASE-BX-U-10	AA1419069-E6	Transmits at 1310 nm. The range is up to 10km upstream.
1000BASE-BX-D-10	AA1419070-E6	Transmits at 1490 nm. The range is up to 10 km downstream.
1000BASE-BX-U-40	AA1419076-E6	Transmits at 1310 nm. The range is up to 40 km upstream.
1000BASE-BX-D-40	AA1419077-E6	Transmits at 1490 nm. The range is up to 40 km downstream.
1000BASE-EX DDI	AA1419071-E6	1550 nm, up to 120 km (non-CWDM)
1000BASE DDI CWDM	AA1419053-E6	1470 nm, up to 40 km
	AA1419054-E6	1490 nm, up to 40 km
	AA1419055-E6	1510 nm, up to 40 km

Table continues…

Model	ROHS product number	Description
	AA1419056-E6	1530 nm, up to 40 km
	AA1419057-E6	1550 nm, up to 40 km
	AA1419058-E6	1570 nm, up to 40 km
	AA1419059-E6	1590 nm, up to 40 km
	AA1419060-E6	1610 nm, up to 40 km
	AA1419061-E6	1470 nm, up to 70 km
	AA1419062-E6	1490 nm, up to 70 km
	AA1419063-E6	1510 nm, up to 70 km
	AA1419064-E6	1530 nm, up to 70 km
	AA1419065-E6	1550 nm, up to 70 km
	AA1419066-E6	1570 nm, up to 70 km
	AA1419067-E6	1590 nm, up to 70 km
	AA1419068-E6	1610 nm, up to 70 km
100BASE-FX	AA1419074-E6	1310 nm, up to 2km

The 9048GB is 100/1000 Mbps capable.

The 9048GB has a 1 GHz 8584E processor and 1 GB onboard DDR2 memory. This module has a maximum throughput of 70 Mpps.

The 9048GB has the following characteristics:

- compliant with IEEE 802.3z standards
- 802.3 Ethernet frame format, MAC layer functionality
- asynchronous Ethernet interface

You can use the 9048GB module in the Virtual Services Platform 9010 and Virtual Services Platform 9012.

9012QQ-2 I/O module

The second generation 9012QQ-2 Input/Output (I/O) module is a 12-port 40 Gigabits per second (Gbps) module. The 9012QQ-2 module supports the 40GBASE-R QSFP+ transceivers. You must also have a minimum of five Switch Fabric modules installed, if you install the 9012QQ-2 module on the Virtual Services Platform 9010 or Virtual Services Platform 9012. The Virtual Services Platform 9012 requires the High-Speed Front Cooling Modules be installed before you install the 9012QQ-2 module.

This module supports standard management information base (MIB), 802.3ba.

Note:

The 9012QQ-2 module does not support Lossless Ethernet.

The 9012QQ-2 module is oversubscribed 2:1 with regards to line rate over 12 ports of 40 Gbps Ethernet traffic using standard QSFP+ fiber transceivers. QSFP+ fiber transceivers, also support directly-attached cables (DACs).

The module supports SR4, LR4 optical modules, and CR4 DACs format.

Important:

Virtual Services Platform 9000 supports only Avaya-qualified transceivers. Other vendor transceivers will not work and Avaya does not support them.

QSFP+ transceivers are hot-swappable input and output enhancement components that allow 40 Gigabit Ethernet ports to link with other 40 Gigabit Ethernet ports. All Avaya QSFP+ transceivers use Lucent connectors (LC) and MTO/MTR connectors to provide precision keying and low interface losses.

The following table lists and describes the Avaya QSFP+ models.

Table 9: Compatible 40 Gigabit QSFP+ transceivers

Hardware	Description	Part number
QSFP+ to QSFP+ DAC	1 meter Passive DAC	AA1404029-E6
QSFP+ to QSFP+ DAC	2 meter Passive DAC	AA1404030-E6
QSFP+ to QSFP+ DAC	3 meter Passive DAC	AA1404031-E6
QSFP+ to QSFP+ DAC	5 meter Passive DAC	AA1404032-E6
QSFP+ to QSFP+ DAC	0.5 meter Passive DAC	AA1404037-E6
QSFP+ to QSFP+ DAC	0.5 meter Passive flexi-DAC (TAA)	AA1404037-E6GS
QSFP+ to QSFP+ DAC	1 meter Passive flexi-DAC (TAA)	AA1404038-E6GS
QSFP+ to QSFP+ DAC	3 meter Passive flexi-DAC (TAA)	AA1404039-E6GS
40GBASE-LR4 QSFP+	10 km	AA1404001-E6
40GBASE-SR4 4x10GBASE-SR QSFP+	100 meters with OM3 fiber cable 150 meters with OM4 fiber cable	AA1404005-E6

The 9012QQ-2 I/O module has a dual core P2020 processor and 2 GB onboard DDR3 memory.

The 9012QQ-2 has the following characteristics:

- compliant with IEEE 802.3ba standards
- 802.3 Ethernet frame format, MAC layer functionality
- asynchronous Ethernet interface

Chapter 4: Platform functionality and system handling

This section contains information and procedures on platform functionality and system handling.

Quick reference

The following table provides a quick reference to the platform functionality and system handling differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series. For more detail about the Virtual Services Platform 9000 implementation, see the sections that follow.

Table 10: Platform functionality and system handling quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
ERS 8000 supports CLI and ACLI.	VSP 9000 supports ACLI.
ERS 8000 ports are enabled by default.	Ports are shutdown by default, which saves energy and is a security factor, if starting the switch without a configuration.
ERS 8000 uses a boot.cfg file If you reset the system to factory defaults, the boot flags are not returned to default values because they are stored in boot.cfg.	VSP 9000 does not use a boot.cfg file. Out of band (OOB) IP addresses are stored in config.cfg. Boot flags, except the factorydefault flag, are stored in config.cfg. If you reset the VSP 9000 to factory defaults, the boot flags are returned to default values. The configuration file boot choices are stored in release/version.cfg.
In a dual CPU configuration, the `reset` and `boot` commands restart only the Master CPU, and the Standby CPU takes over.	In a dual CPU configuration, both the `reset` and `boot` commands produce a full chassis restart, where both CPUs (and all IO and SF modules) restart.
The boot flags ha-cpu and savetostandby are disabled by default.	The boot flags ha-cpu and savetostandby are enabled by default.
ERS 8000 cannot use PCAP in HA mode.	VSP 9000 can use PCAP in HA mode.

Table continues…

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
-	OOB Ethernet ports belong to a separate predefined MgmtRouter VRF, which changes the configuration steps.
ERS 8000 does not support software patching.	VSP 9000 supports software patching.
-	With a dual-redundant system configuration, you can perform upgrades without a network outage. For example, in an SMLT configuration, you can take one system offline and upgrade it. The redundant system manages all network activity. After you bring the system back online, you can upgrade the second system.
-	VSP 9000 supports Key Health Indicators to monitor system health.
-	VSP 9000 includes an alarm database to view local alarms for troubleshooting.
ERS 8000 creates one core file	VSP 9000 packages core file and flight recorder archives and stores them on the external flash. For more information see *Troubleshooting Avaya Virtual Services Platform 9000,* NN46250-700.
-	VSP 9000 includes hardware support for hardware failure detection (Rapid Failure Detection and Recovery) with in-service health checks.

Upgrading the software

Perform this procedure to upgrade the software on the Avaya Virtual Services Platform 9000. This procedure shows how to upgrade the software using the internal flash memory as the file storage location; you can use other storage locations.

Before you begin

- Back up the configuration files.
- Download the upgrade file to the Virtual Services Platform 9000.
- Avaya Virtual Services Platform 9010 supports two upgrade paths:
  - Release 3.3.3.x to Release 4.0.
  - Release 3.4.x to Release 4.0.
- Avaya Virtual Services Platform 9012: No restrictions exist for the upgrade. For Avaya Virtual Services Platform 9012, you can upgrade directly to Release 4.0. You do not need to install Release 3.3.3.0 or later prior to upgrading to Release 4.0.

Procedure

1. Enter Privileged EXEC mode:

```
enable
```

2. Extract the release distribution files to the `/intflash/release/` directory:

   ```
   software add WORD<1-99>
   ```

3. Install the image:

   ```
   software activate WORD<1-99>
   ```

4. Restart the Virtual Services Platform 9000:

   ```
   reset
   ```

   Important:

   After you restart the system, you have the amount of time configured for the commit timer to verify the upgrade and commit the software to gold. If you do not commit the software to gold and you did not enable auto-commit, the system restarts with the last known working version after the commit timer expires. This feature ensures you can regain control of the system if an upgrade fails.

5. Confirm the software upgrade:

   ```
   show software [verbose]
   ```

Example

The following figures show detailed examples of the **show software** and **show software verbose** commands throughout the various stages of the upgrade process, such as **software add**, **software activate**, and **software commit** (manual and auto). The following example shows a software upgrade from Release 3.3.1.1 to Release 3.4.

Note:

The figures in this example demonstrate a 3.4.0.0 software build. These examples are valid for all post-3.4.0.0 software upgrades.

Software add:

In the following two examples, note that Release 3.4.0.0.GA now appears in the output. The **show software verbose** output also shows the Committed Type as Not Committed.

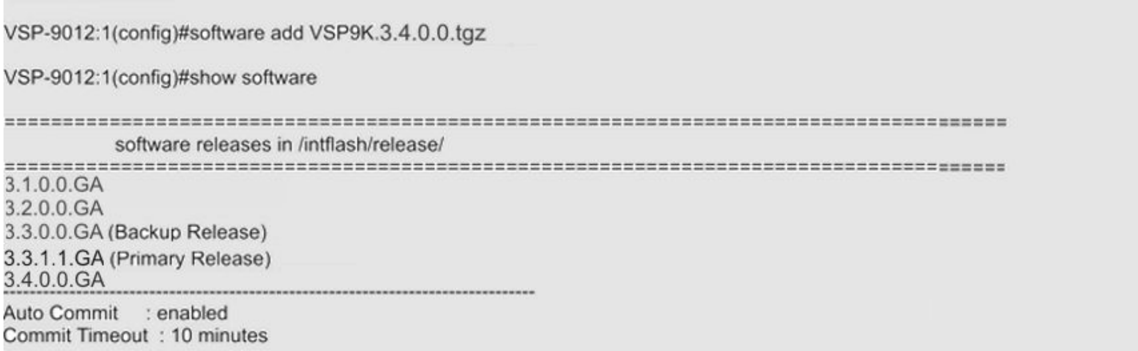

```
VSP-9012:1(config)#software add VSP9K.3.4.0.0.tgz

VSP-9012:1(config)#show software

===================================================================================
              software releases in /intflash/release/
===================================================================================
3.1.0.0.GA
3.2.0.0.GA
3.3.0.0.GA (Backup Release)
3.3.1.1.GA (Primary Release)
3.4.0.0.GA--------------------------------------------------------------------
Auto Commit     : enabled
Commit Timeout  : 10 minutes
```

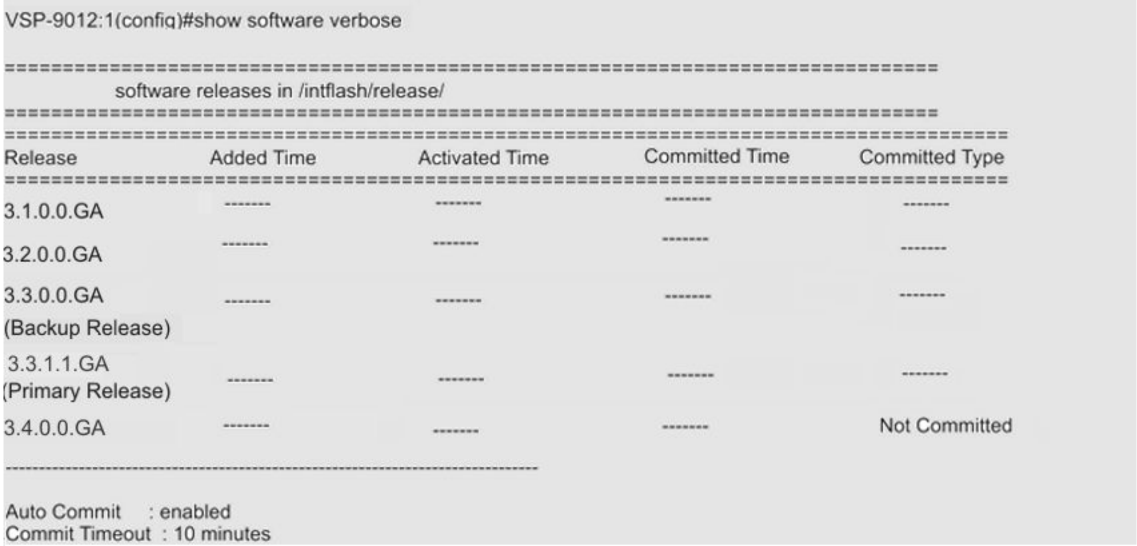

```
VSP-9012:1(config)#show software verbose

==============================================================================
              software releases in /intflash/release/
==============================================================================
===================================================================================
Release           Added Time        Activated Time     Committed Time     Committed Type
===================================================================================
3.1.0.0.GA          -------           -------            -------             -------

3.2.0.0.GA          -------           -------            -------             -------

3.3.0.0.GA          -------           -------            -------             -------
(Backup Release)
 3.3.1.1.GA         -------           -------            -------             -------
(Primary Release)
3.4.0.0.GA          -------           -------            -------          Not Committed

--------------------------------------------------------------------------

Auto Commit     : enabled
Commit Timeout  : 10 minutes
```

Software activate:

In the following two examples, note that Release 3.4.0.0.GA is now shown as the Next Boot Release.

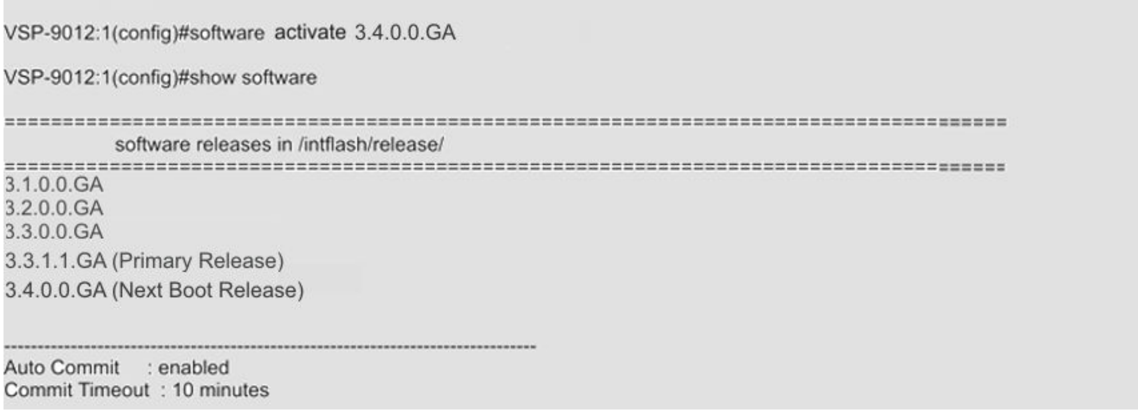

```
VSP-9012:1(config)#software activate 3.4.0.0.GA

VSP-9012:1(config)#show software

==========================================================================================
              software releases in /intflash/release/
==========================================================================================
3.1.0.0.GA
3.2.0.0.GA
3.3.0.0.GA
3.3.1.1.GA (Primary Release)
3.4.0.0.GA (Next Boot Release)

--------------------------------------------------------------------------------
Auto Commit     : enabled
Commit Timeout  : 10 minutes
```

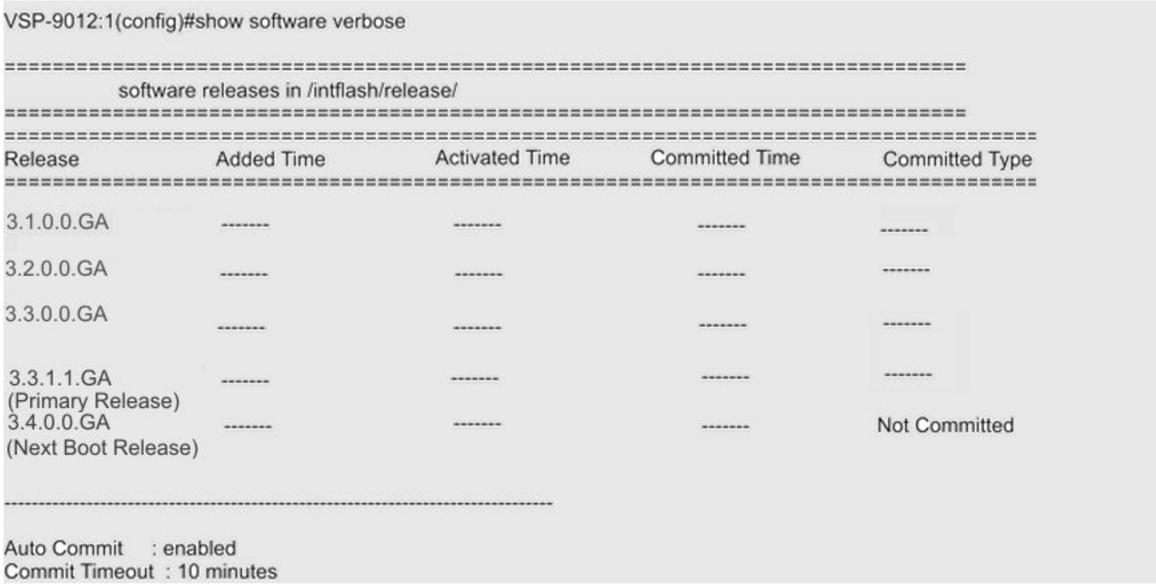

```
VSP-9012:1(config)#show software verbose

=====================================================================================
              software releases in /intflash/release/
=====================================================================================
===========================================================================================
Release            Added Time        Activated Time      Committed Time      Committed Type
===========================================================================================

3.1.0.0.GA          -------           -------             -------             -------

3.2.0.0.GA          -------           -------             -------             -------

3.3.0.0.GA          -------           -------             -------             -------

3.3.1.1.GA          -------           -------             -------             -------
(Primary Release)
3.4.0.0.GA          -------           -------             -------             Not Committed
(Next Boot Release)

--------------------------------------------------------------------------------

Auto Commit     : enabled
Commit Timeout  : 10 minutes
```

After a system restart:

In the following two examples, note that Release 3.4.0.0.GA is now shown as the Primary Release. The auto-commit timer shows the remaining time before an auto-commit. The **`show software verbose`** output shows the timestamps for the Added Time and Activated Time.

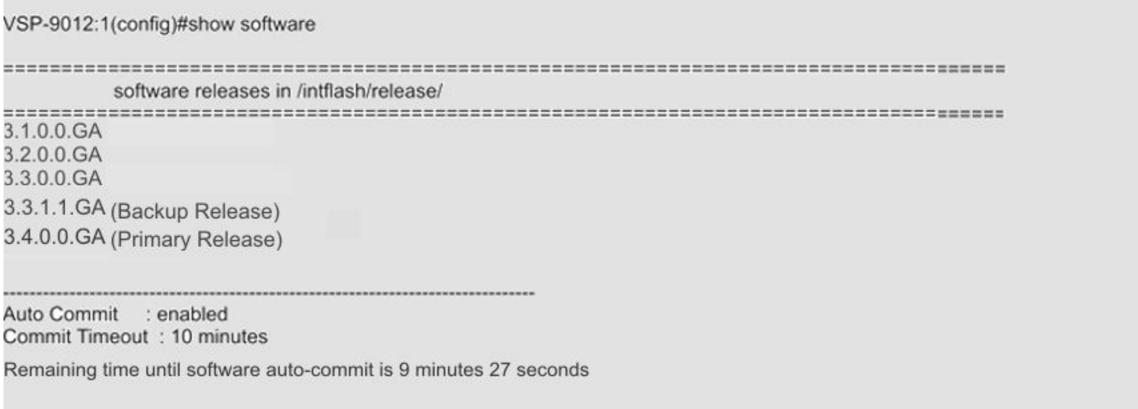

```
VSP-9012:1(config)#show software

========================================================================================
              software releases in /intflash/release/
========================================================================================
3.1.0.0.GA
3.2.0.0.GA
3.3.0.0.GA
3.3.1.1.GA (Backup Release)
3.4.0.0.GA (Primary Release)

--------------------------------------------------------------------------------
Auto Commit     : enabled
Commit Timeout  : 10 minutes
Remaining time until software auto-commit is 9 minutes 27 seconds
```

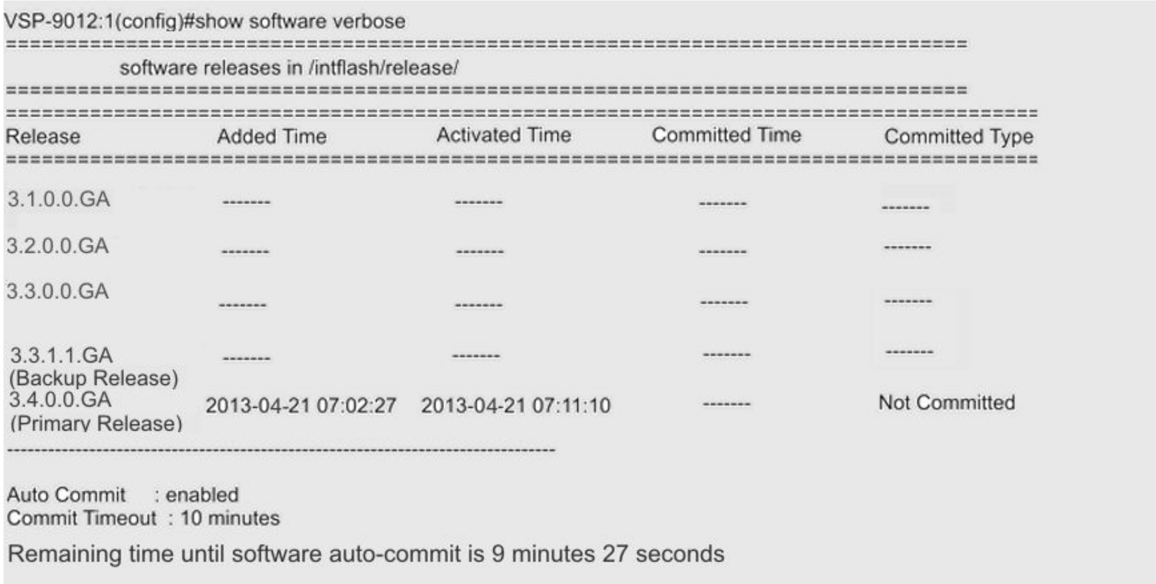

```
VSP-9012:1(config)#show software verbose
========================================================================================
              software releases in /intflash/release/
========================================================================================
==============================================================================================
Release              Added Time            Activated Time        Committed Time      Committed Type
==============================================================================================

3.1.0.0.GA           -------               -------               -------             -------

3.2.0.0.GA           -------               -------               -------             -------

3.3.0.0.GA           -------               -------               -------             -------

3.3.1.1.GA           -------               -------               -------             -------
(Backup Release)
3.4.0.0.GA           2013-04-21 07:02:27   2013-04-21 07:11:10   -------             Not Committed
(Primary Release)
--------------------------------------------------------------------------------

Auto Commit     : enabled
Commit Timeout  : 10 minutes
Remaining time until software auto-commit is 9 minutes 27 seconds
```

Manual and automatic software commit:

The following example shows the **show software verbose** output for a manual software commit. The Committed Time timestamp appears and the Committed Type shows as Manual.

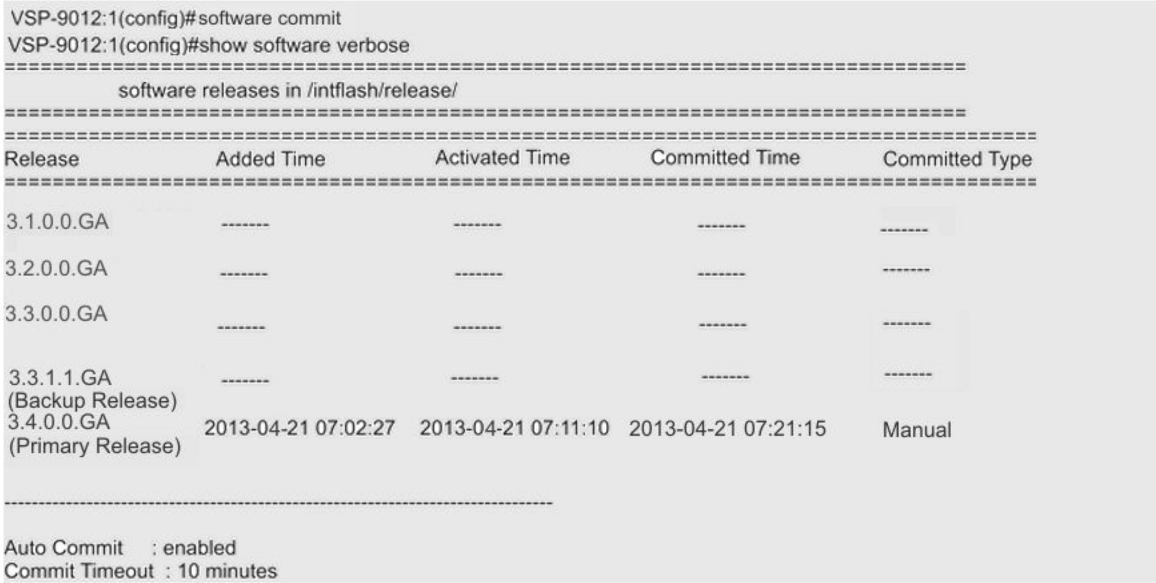

```
VSP-9012:1(config)#software commit
VSP-9012:1(config)#show software verbose
================================================================================
          software releases in /intflash/release/
================================================================================
======================================================================================
Release            Added Time            Activated Time         Committed Time         Committed Type
======================================================================================
3.1.0.0.GA         -------               -------                -------                -------
3.2.0.0.GA         -------               -------                -------                -------
3.3.0.0.GA         -------               -------                -------                -------

3.3.1.1.GA         -------               -------                -------                -------
(Backup Release)
3.4.0.0.GA         2013-04-21 07:02:27   2013-04-21 07:11:10   2013-04-21 07:21:15    Manual
(Primary Release)

--------------------------------------------------------------------------------

Auto Commit     : enabled
Commit Timeout  : 10 minutes
```

The following example shows the **`show software verbose`** output for an automatic software commit after the 10 minute commit timer expired. The Committed Time timestamp appears and the Committed Type shows as Auto.

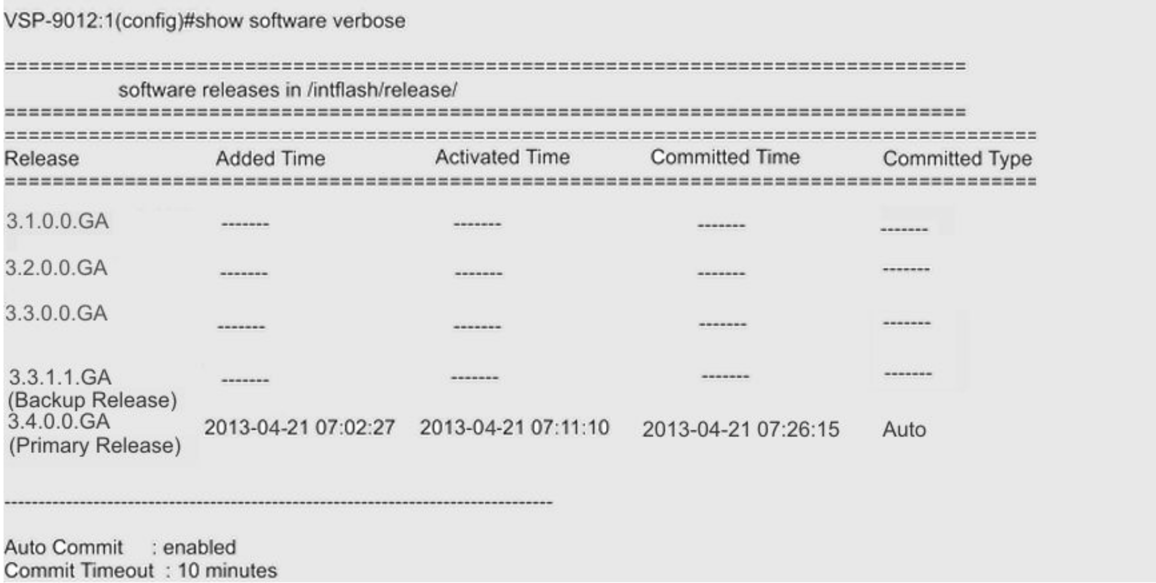

```
VSP-9012:1(config)#show software verbose

================================================================================
          software releases in /intflash/release/
================================================================================
======================================================================================
Release            Added Time            Activated Time         Committed Time         Committed Type
======================================================================================
3.1.0.0.GA         -------               -------                -------                -------
3.2.0.0.GA         -------               -------                -------                -------
3.3.0.0.GA         -------               -------                -------                -------

3.3.1.1.GA         -------               -------                -------                -------
(Backup Release)
3.4.0.0.GA         2013-04-21 07:02:27   2013-04-21 07:11:10    2013-04-21 07:26:15   Auto
(Primary Release)

--------------------------------------------------------------------------------

Auto Commit     : enabled
Commit Timeout  : 10 minutes
```

All releases prior to Release 3.4 do not show a timestamp in the **`show software verbose`** command output regardless of whether you add, activate, or commit. When upgrading from releases prior to Release 3.4, the Added Time and Activated Time timestamps appear after a system restart and you only see the timestamps for Release 3.4 and later.

Note:

For upgrades post Release 3.4, you do not need to restart the system to see the updated timestamps for Added Time and Activated Time. The timestamps immediately appear in **`show software verbose`** output after you issue the commands **`software add`** and **`software activate`**.

Variable definitions

Use the data in the following table to use the **`software`** command.

Variable	Value
activate *WORD<1-99>*	Specifies the name of the software release image.
add *WORD<1-99>*	Specifies the path and version of the compressed software release archive file.
remove *WORD<1-99>*	Specifies the path and version of the compressed software release archive file.

Committing an upgrade

Perform the following procedure to commit an upgrade.

About this task

The commit function for software upgrades allows maximum time set by the commit timer (the default is 10 minutes) to ensure that the upgrade is successful. If you enable the auto-commit option, the system automatically commits to the new software version after the commit timer expires. If you disable the auto-commit option, you must issue the software commit command before the commit timer expires to commit the new software version, otherwise the system restarts automatically to the previous (committed) version.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. **(Optional)** Extend the time to commit the software:

   ```
   software reset-commit-time [<1-60>]
   ```

3. Commit an upgrade:

   ```
   software commit
   ```

Adding an encryption module

You can add encryption modules to a software release to use encryption features.

Before you begin

- Download the module file to the Avaya Virtual Services Platform 9000.
- You must log on to at least the Privileged EXEC mode in ACLI.

Procedure

1. Add an encryption module to a release version:

   ```
   software add-modules WORD<1-99> WORD<1-99>
   ```

2. Enter Global Configuration mode:

   ```
   enable
   ```

   ```
   configure terminal
   ```

3. Load the encryption module:

   ```
   load-encryption-module {3DES|AES|DES}
   ```

Example

Add the Advanced Encryption Standard (AES) module to the Release 3.4:

```
VSP-9012:1#software add-modules 3.4.0.0.GA VSP9K.3.4.0.0_modules.tgz
```

Log on to Global Configuration mode:

```
VSP-9012:1#configure terminal
```

Load the AES encryption module:

```
VSP-9012:1(config)#load-encryption-module AES
```

Variable definitions

Use the data in the following table to use the **software add-modules** command.

Variable	Value
The first *WORD<1-99>*.	Specifies the release version to which you want to add modules to.
The second *WORD<1-99>*.	Specifies the module archive you want to add.

Use the data in the following table to use the **load-encryption-module** command.

Variable	Value
<3DES|AES|DES>	Specifies the encryption module to load.

Resetting the platform

About this task

Reset the platform to reload system parameters from the most recently saved configuration file.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Use one of the following commands to reset the switch:

   a. Reset the switch and receive a prompt to confirm the reset:

      ```
      reset
      ```

   b. Reset the switch and do not receive a prompt to confirm the reset:

      ```
      reset -y
      ```

   c. Reset the switch, receive a prompt to confirm the reset, and create a core dump file:

      ```
      reset -coredump
      ```

   d. Reset the switch, do not receive a prompt to confirm the reset, and create a core dump file:

      ```
      reset -coredump -y
      ```

Example

```
VSP-9012:1>enable
```

Reset the switch:

```
VSP-9012:1#reset

Are you sure you want to reset the switch? (y/n)y
```

Variable definitions

Use the data in the following table to use the **reset** command.

Table 11: Variable definitions

Variable	Value
-coredump	Creates a coredump for the main process before the switch resets. ⚠ **Caution:** Only use the -coredump parameter if an issue causes you to reset the switch, and you need to contact customer service for analysis of the problem.
-y	Suppresses the confirmation message before the switch resets. If you omit this parameter, you must confirm the action before the system resets.

Restarting the platform

About this task

Restart the switch to implement configuration changes or recover from a system failure. When you restart the system, you can specify the boot source (internal flash, external flash, USB, or TFTP server) and file name. If you do not specify a device and file, the run-time ACLI uses the software and configuration files on the primary boot device defined by the `boot config choice` command.

After the switch restarts normally, it sends a cold trap within 45 seconds after a restart. If a CPU (9080CP module) switchover occurs during operation, the switch sends a warm-start management trap within 45 seconds of a restart.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. Restart the switch:

   ```
   boot [config WORD<1-99>] [-y]
   ```

   Important:

   If you enter the `boot` command with no arguments, you cause the switch to start using the current boot choices defined by the `boot config choice` command.

3. Use one of the following commands to switch-over to the backup CP:

   a. Change to the backup CP:

      ```
      sys action cpu-switch-over
      ```

   b. Generate a core dump of the master CP, and then switch to the backup:

      ```
      sys action cpu-switch-over -coredump
      ```

4. Reset system functions:

   ```
   sys action reset {console|counters}
   ```

Variable definitions

Use the data in the following table to use the **`boot`** command.

Table 12: Variable definitions

Variable	Value
config *WORD<1–99>*	Specifies the software configuration device and file name in one of the following formats: • a.b.c.d:<file> • /intflash/ <file> • /extflash/ <file> • /usb/<file> The file name, including the directory structure, can include up to 99 characters.
-y	Suppresses the confirmation message before the switch restarts. If you omit this parameter, you must confirm the action before the system restarts.

Use the data in the following table to use the `sys action` command.

Table 13: Variable definitions

Variable	Value
-coredump	Creates a coredump for the main process before the switch changes to the backup CPU. ⚠ **Caution:** Only use the -coredump parameter if an issue causes you to switchover the switch, and you need to contact customer service for analysis of the problem.
cpu-switch-over	Resets the switch to change over to the backup CPU.
reset {console|counters}	Reinitializes the hardware universal asynchronous receiver transmitter (UART) drivers. Use this command only if the console connection does not respond. Resets all the statistics counters in the switch to zero. Resets the console port.

Configuring system flags

Configure the system flags to enable specific services and functions for the chassis.

About this task

❗ **Important:**

- If you enable the hsecure flag, you cannot enable the flags for the web server or SSH password-authentication.

- After you change certain configuration parameters using the `boot config flags` command, you must save the changes to the configuration file.

File Transfer Protocol (FTP), Trivial File Transfer Protocol (TFTP) and Telnet server support both IPv4 and IPv6 addresses, with no difference in functionality or configuration.

On IPv6 networks, VSP 9000 supports SSH server, remote login (rlogin) server and Remote Shell (rsh) server only. VSP 9000 does not support outbound SSH client over IPv6, rlogin client over IPv6 or rsh client over IPv6. On IPv4 networks, VSP 9000 supports both server and client for SSH, rlogin and rsh.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   ```

   ```
   configure terminal
   ```

2. Enable system flags:

   ```
   boot config flags <block-snmp|debug-config [console|file]|debugmode|
   fabric-profile <1-3>|factorydefaults|ftpd|ha-cpu|hsecure|logging|
   reboot|rlogind|savetostandby|spanning-tree-mode <mstp|rstp>|sshd|
   telnetd|tftpd|trace-logging|verify-config>
   ```

3. Disable system flags:

   ```
   no boot config flags <block-snmp|debug-config|debugmode|
   factorydefaults|ftpd|ha-cpu|hsecure|logging|reboot|rlogind|
   savetostandby|spanning-tree-mode|sshd|telnetd|tftpd|trace-logging|
   verify-config>
   ```

4. Configure the system flag to the default value:

   ```
   default boot config flags <block-snmp|debug-config|debugmode|fabric-
   profile|factorydefaults|ftpd|ha-cpu|hsecure|logging|reboot|rlogind|
   savetostandby|spanning-tree-mode|sshd|telnetd|tftpd|trace-logging|
   verify-config>
   ```

5. Save the changed configuration.

6. Restart the switch.

Example

Activate High Secure mode:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#boot config flags hsecure
VSP-9012:1(config)#save config
VSP-9012:1(config)#reset
```

To debug loading of the configuration file, enable logging of the line-by-line configuration file processing and the result of the execution to the debug file, while the device loads the configuration file:

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#boot config flags debug-config file
VSP-9012:1(config)#save config
VSP-9012:1(config)#reset
```

Variable definitions

Use the data in the following table to use the **boot config flags** command.

Table 14: Variable definitions

Variable	Value
block-snmp	Activates or disables Simple Network Management Protocol management. The default value is false (disabled), which permits SNMP access.
debug-config [console][file]	Enables you to debug the configuration file during loading configuration at system boot up. The default is disabled. You do not have to restart the switch after you enable debug-config, unless you want to immediately debug the configuration. After you enable debug-config and save the configuration, the debug output either displays on the console or logs to an output file the next time the switch reboots. The two options include: • debug-config [console] – Displays the line-by-line configuration file processing and result of the execution on the console while the device loads the configuration file. • debug-config [file] – Logs the line-by-line configuration file processing and result of the execution to the debug file while the device loads the configuration file. The system logs the debug-config output to the `/extflash/debugconfig_primary.txt` for the primary configuration file. The system logs the debug config output to the `/extflash/debugconfig_backup.txt` for the backup configuration, if the backup configuration file loads.

Table continues…

Variable	Value
debugmode	❗ **Important:** Do not change this parameter unless directed by Avaya. Controls whether the switch stops in debug mode following a fatal error. Debug mode provides information equivalent to the **`trace`** commands. If you enable this flag, the switch does not restart following a fatal error. The default value is disabled. If you change this parameter, you must restart the switch.
fabric-profile <1–3>	Configures the system to give preference to one type of traffic over the other in times of over subscription. The values are: • 1: balanced • 2: unicast optimized • 3: multicast optimized If you change this parameter, you must restart the switch. The default profile is 1, balanced.
factorydefaults	Specifies whether the switch uses the factory default settings at startup. The default value is disabled. This flag is automatically reset to the default setting after the CPU restarts. If you change this parameter, you must restart the switch.
ftpd	Activates or disables the FTP server on the switch. The default value is disabled. To enable FTP, ensure that the tftpd flag is disabled.
ha-cpu	Activates or disables High Availability (HA) mode. Switches with two CPUs use HA mode to recover quickly from a failure of one of the CPUs. If you enable or disable High Availability mode, the secondary CPU resets automatically to load settings from the saved configuration file.
hsecure	Activates or disables High Secure mode. The hsecure command provides the following password behavior: • 10 character enforcement • aging time • failed login attempt limitation The default value is disabled. If you enable High Secure mode, you must restart the switch to enforce secure passwords. If you operate the switch in High

Table continues…

Variable	Value
	Secure mode, the switch prompts a password change if you enter invalid-length passwords.
logging	Activates system logging. The default value is enabled. The system names log files according to the following: • File names appear in 8.3 (log.xxxxxxxx.sss) format. • The first 6 characters of the file name contain the last three bytes of the chassis base MAC address. • The next two characters in the file name specify the slot number of the CPU that generated the logs. • The last three characters in the file name are the sequence number of the log file. The system generates multiple sequence numbers for the same chassis and same slot if the system reaches the maximum log file size.
reboot	❗ **Important:** Do not change this parameter unless directed by Avaya. Activates or disables automatic reboot on a fatal error. The default value is activated. The reboot command is equivalent to the debugmode command. If you change the reboot variable value, you must restart the switch.
rlogind	Activates or disables the rlogin and rsh server. The default value is disabled.
savetostandby	Activates or disables automatic save of the configuration file to the standby CPU. The default value is enabled. If you operate a dual CPU system, Avaya recommends that you enable this flag for ease of operation.
spanning-tree-mode <mstp|rstp>	Specifies the Multiple Spanning Tree Protocol or Rapid Spanning Tree Protocol mode. If you do not specify a protocol, the switch uses the default mode. The default mode is mstp. If you change the spanning tree mode, you must save the current configuration and restart the switch.
sshd	Activates or disables the SSH server service. The default value is enabled.
telnetd	Activates or disables the Telnet server service. The default is disabled.

Table continues…

Variable	Value
	If you disable the Telnet server service in a dual CPU system, the Telnet server prevents a Telnet connection initiated from the other CPU.
tftpd	Activates or disables Trivial File Transfer Protocol server service. The default value is disabled. If you disable the TFTP server, you can still copy files between the CPUs.
trace-logging	❗ **Important:** Do not change this parameter unless directed by Avaya. Activates or disables the creation of trace logs. The default value is disabled.
verify-config	Activates syntax checking of the configuration file. The default is enabled. • Primary config behavior: When the verify-config flag is enabled, the primary config file is pre-checked for syntax errors. If the system finds an error, the primary config file is not loaded, instead the system loads the backup config file. If the verify-config flag is disabled, the system does not pre-check syntax errors. When the verify-config flag is disabled, the system ignores any lines with errors during loading of the primary config file. If the primary config file is not present or cannot be found, the system tries to load the backup file. • Backup config behavior: If the system loads the backup config file, the system does not check the backup file for syntax errors. It does not matter if the verify-config flag is disabled or enabled. With the backup config file, the system ignores any lines with errors during the loading of the backup config file. If no backup config file exists, the system defaults to factory defaults. Avaya recommends that you disable the verify-config flag.

High Availability mode

High Availability (HA) mode, also called HA-CPU, activates two CPUs simultaneously. These CPUs exchange topology data so that, if a failure occurs, either CPU can take over the operations of the other.

In HA-CPU mode, the two CPUs are active and exchange topology data through an internal dedicated bus. This configuration allows for a complete separation of traffic. To guarantee total security, users cannot access this bus.

In HA-CPU mode, also called Hot Standby, the two CPUs synchronize. In non HA-CPU mode, also called Warm Standby, the two CPUs do not synchronize.

The following tables lists feature support and synchronization information for HA-CPU.

Table 15: Feature support for HA-CPU

Feature	Release 3.4
Modules	Yes
Platform	Yes
Layer 2	Yes
Layer 3	Yes; partial-HA for Border Gateway Protocol and IPv6
Multicast	Partial HA
Multicast virtualization (virtualized IGMP only)	Full HA
Security	Yes
Applications	Partial HA

Table 16: Synchronization capabilities in HA-CPU mode

Synchronization of	Release 3.4
Layer 1	
Port configuration parameters	Yes
Layer 2	
Multiple Spanning Tree Protocol parameters	Yes
Quality of Service (QoS) parameters	Yes
Rapid Spanning Tree Protocol parameters	Yes
Shortest Path Bridging MAC (MAC)	Yes
Connectivity Fault Management (CFM)	Yes
SMLT parameters	Yes
VLAN parameters	Yes
Layer 3	
ARP entries	Yes

Table continues…

Synchronization of	Release 3.4
Border Gateway Protocol (BGP)	Partial (configuration only)
Dynamic Host Configuration Protocol (DHCP) Relay	Partial (configuration only)
Internet Group Management Protocol (IGMP)	Yes
IGMP virtualization Snooping SPB	Yes
IGMP Snooping routed-SPB	Yes
IP filters	Yes
IPv6 filters	Yes
IPv6	Partial (configuration only)
Layer 3 Filters: access control entries, access control lists	Yes
Open Shortest Path First (OSPF)	Yes
Packet Capture (PCAP) tool	Yes
PIM	Partial (configuration only)
Prefix lists and route policies	Yes
Routing Information Protocol	Yes
Router Discovery	Yes
Routed Split Multi-Link Trunking (RSMLT)	Yes
RSMLT edge support	Yes
Shortest Path Bridging MAC (SPBM)	Yes
Connectivity Fault Management (CFM)	Yes
Static and default routes	Yes
Virtual IP (VLANs)	Yes
Virtual Router Redundancy Protocol	Yes
VRF Lite	Yes
Transport Layer	
Network Load Balancing (NLB)	Yes
Remote Access Dial-In User Services (RADIUS)	Yes
Terminal Access Controller Access-Control System plus (TACACS+)	Partial (configuration only)
UDP forwarding	Yes
Applications	
VSP Talk	Partial (configuration only). After a CPU switchover, you must re-enable event notification.

For more information about how to configure HA-CPU, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

Partial HA implementation

A few applications in HA-mode, or Hot Standby mode, have partial HA implementation. This means that the system synchronizes user configuration data (including interfaces, IPv6 addresses and static routes) between the master CPU and standby CPU. However, for applications in HA-mode with partial HA implementation, the platform does not synchronize dynamic data learned by protocols. As a result, after failure those applications need to restart and rebuild their tables. This operation causes an interruption to traffic that is dependent on a protocol or application with Partial HA support. The following applications support Partial High Availability:

- Border Gateway Protocol (BGP)
- Dynamic Host Configuration Protocol (DHCP) Relay
- IPv6
- Link Layer Discovery Protocol (LLDP), 802.1AB
- Protocol Independent Multicast-Sparse Mode (PIM-SM)
- Protocol Independent Multicast-Source Specific Mode (PIM-SSM)
- Terminal Access Controller Access-Control System plus (TACACS+)
- VSP Talk

HA-CPU limitations and considerations

You must take the following limitations and considerations into account when you use the HA-CPU feature:

- Activating or deactivating HA-CPU mode causes the standby CP to reset. The active CP continues to operate normally.
- In HA-CPU mode, Avaya recommends that you do not configure the Open Shortest Path First (OSPF) dead router interval for less than 15 seconds.

Out-of-band configuration

On Virtual Services Platform 9000, OOB Ethernet ports belong to a separate, predefined MgmtRouter VRF with VRF ID 512. You cannot delete the MgmtRouter VRF. This VRF membership changes the commands you use to configure the OOB ports.

Ethernet Routing Switch 8000 series

This section shows an example configuration for the OOB ports on the Ethernet Routing Switch 8000 series.

```
config term
boot config net mgmt ip 192.168.10.105/24 cpu-slot 5
boot config net mgmt ip 192.168.10.106/24 cpu-slot 6
boot config net mgmt route 192.168.0.0/16 192.168.10.1
save bootconfig            # Saves preceding configuration to boot.cfg
sys mgmt-virtual-ip 192.168.10.100/24
save config                   # Saves mgmt-virtual-ip to config.cfg
ping 192.168.10.1          # Verifies connectivity to OOB default gateway
```

Virtual Services Platform 9000

This section shows an example configuration for the OOB ports on the Virtual Services Platform 9000.

```
config term
interface mgmtEthernet 1/1
ip address 192.168.10.101/24
exit
interface mgmtEthernet 2/1
ip address 192.168.10.102/24
exit
router vrf MgmtRouter
ip route 192.168.0.0 255.255.0.0 192.168.10.1 weight 10
exit
sys mgmt-virtual-ip 192.168.10.100/24
save config                               # Saves all to config.cfg
ping 192.168.10.1 vrf MgmtRouter          # Verifies connectivity to OOB default gateway
```

Key Health Indicators

The Key Health Indicators (KHI) feature of the Virtual Services Platform 9000 provides a subset of health information that allows for quick assessment of the overall operational state of the device. KHI periodically measures important system information that reflects the state of the system.

The KHI feature does not provide a comprehensive debugging solution. Instead, KHI identifies key information that can lead Avaya support personnel towards discovery of a specific failure. After the technician assesses the KHI information, you must do further debugging to determine the specific reason for the fault.

Avaya recommends that you capture KHI information during normal operations to provide a baseline for Avaya support personnel when detecting fault situations. For example, after you first install and configure the Virtual Services Platform 9000, and verify that it operates as expected, capture KHI information.

For more information about KHI, see *Monitoring Performance on Avaya Virtual Services Platform 9000,* NN46250-701.

Local alarms

The Avaya Virtual Services Platform 9000 contains a local alarms mechanism. Applications that run on the switch raise and clear local alarms. Use the **`show alarm database`** command to view active alarms. Local alarms are an automatic mechanism run by the system that do not require additional user configuration. Check local alarms occasionally to ensure no alarms require additional operator attention. The raising and clearing of local alarms also creates a log entry for each event.

For more information, see *Troubleshooting Avaya Virtual Services Platform 9000,* NN46250-700.

Rapid Failure Detection and Recovery

Virtual Services Platform 9000 provides Rapid Failure Detection and Recovery (RFDR) of less than 20 milliseconds (ms) . RFDR applies to the data path including MultiLink Trunking (MLT), Distributed MLT (DMLT), Split MLT (SMLT), and Equal Cost Multipath (ECMP) configurations.

A link state table contains the link states of all the ports and logical connections in the local system and in the remote peer node in the case of SMLT. The state updates every 3.3 ms. All packets that forward to an MLT port, SMLT port, or ECMP route use the link state table with real-time hashing and intelligent pruning to forward the packet to an active port.

Chapter 5: Software considerations

This section contains information on the migration considerations that pertain to the Avaya Virtual Services Platform 9000 software.

Key software differences

Note key support differences between the VSP 9000 software feature set and those you are familiar with from the Avaya Ethernet Routing Switch 8000 series, up to and including Release 7.2. The following table identifies key support differences for software features and protocols.

Feature	Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Avaya VENA Unified Access	Yes	No
Distance Vector Multicast Routing Protocol (DVMRP)	Yes	No
IGMPv3 (full)	No — partial only	Yes
IP Virtual Private Network (IPVPN)	Yes	IP VPN over SPBM only
IP Virtual Private Network Lite (IPVPN Lite)	Yes	No
Lossless Ethernet	No	Yes
Multicast Source Discovery Protocol (MSDP)	Yes	No
Multiprotocol Label Switching (MPLS)	Yes	No
Nortel Spanning Tree	Yes	No
Per VLAN Spanning Tree Plus (PVST+)	Yes	No
Multicast virtualization support of IGMP, PIM-SM, and PIM-SSM	Yes	IGMP virtualization only
VSP Talk (or similar instant messaging client application)	No	Yes

Spanning tree

The following table provides a quick reference to the spanning tree differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series.

Table 17: Spanning tree quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Default spanning tree mode is STP (NT-STG).	Default spanning tree mode is MSTP. Does not support NT-STG.

Virtual Services Platform 9000 only supports Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP) but does provide STP compatibility for interoperability with switches that run in STP mode. An MSTP port automatically downgrades to RSTP operation if it receives RSTP Bridge Protocol Data Units (BPDU) on the port. An MSTP or RSTP port automatically downgrades to legacy STP if it receives a legacy BPDU on the port.

The default spanning-tree mode is MSTP. The default STG for RSTP and MSTP is 0. In RSTP mode all VLANs run in the default STG. In MSTP mode, you can create additional STGs by using the VLAN create command. The Virtual Services Platform 9000 supports up to 64 STGs.

Configuration of NT-STP non-default STGs on the Ethernet Routing Switches is incompatible with MSTP non-default STGs on the Virtual Services Platform 9000. The Virtual Services Platform 9000 drops NT-STP BPDUs for non-default STGs.

For more information about RSTP and MSTP, see *Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000,* NN46250-500.

VLANs

The following table provides a quick reference to the VLAN differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series. For more detail about the Virtual Services Platform 9000 implementation, see the sections that follow.

Table 18: VLANs quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports the following dynamic VLAN classification model: Tagged->Subnet->SrcMAC->Protocol-based->Port Default VLAN If an incoming untagged packet matches a dynamic VLAN, then the port membership is checked (potential, static, active). If the port is not-allowed, classification continues.	Supports the following dynamic VLAN classification model: Tagged->SrcMAC->Subnet->Protocol-based>Port Default VLAN If an incoming untagged packet matches a dynamic VLAN, then the port membership is checked (potential, static, active). If the port is not-allowed, the device drops the packet.

Table continues…

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
	Supports configuration of default model for source MAC or subnet.
If the default VLAN for the port is NULL, the newly assigned VLAN becomes the default for the port.	The default VLAN for the port does not change unless the port belongs to no other port-based VLAN.
If the port belongs to another port-based VLAN, the lowest numbered VLAN becomes the default for the port.	The default VLAN for the port becomes NULL.
After a port changes from tagged (trunk) to untagged (access), the port is removed from all but the lowest port-based VLAN to which it belonged.	You must manually remove the port from all but one port-based VLAN before the port can change to untagged.
After you remove a port from an MLT, the port becomes a member of VLAN 1.	The port is removed from all VLANs.
After you add a port to an MLT with no VLAN assigned, the port is removed from the spanning-tree group (STG) and all VLANs.	The MLT and the port become part of the default STG, and the port is removed from all VLANs with the following warning message: Port's Default VLAN is set to NULL. Untagged packets may be dropped.

Policy-based VLANs

You can base a policy on protocol, IP subnet, or source MAC address.

Protocol-based VLANs

The Virtual Services Platform 9000 supports the following protocol-based VLANs:

- IP version 4 (IP)
- IP version 6 (IPv6)
- AppleTalk on Ethernet Type 2 and Ethernet SNAP frames (AppleTalk)
- Digital Equipment Corporation Local Area Transport (DEC LAT) Protocol (decLat)
- Other DEC protocols (decOther)
- International Business Machines Systems Network Architecture (IBM SNA) on IEEE 802.2 frames (sna802dot2)
- IBM SNA on Ethernet Type 2 frames (snaEthernet2)
- NetBIOS Protocol (netBIOS)
- Xerox Network Systems (XNS)
- Banyan VINES (vines)
- Reverse Address Resolution Protocol (RARP)
- Point-to-Point Protocol over Ethernet (PPPoE)
- ipx802.2
- ipx802.3

- ipxEthernet2
- ipxSnap
- user-defined protocols

Multiple protocol-based VLANs cannot be defined for the same protocol.

The maximum number of protocol-based VLANs that you can configure is 16. This restriction is based on a table of 16 entries. Some protocols create more than one entry in the table. For example, an IP protocol-based VLAN creates two entries; one entry for IP ProtocolId= (0x800) and another for ARP ProtocolId=(0x806). If you configure a IP protocol-based VLAN, you can configure only 14 more protocol-based VLANs.

Configuring a DecOther protocol VLAN uses nine table entries, leaving only seven remaining. The following table provides the standard protocol VLANs supported on the VSP 9000 and the number of records created for each.

Table 19: Records types created for standard protocol VLAN types

Protocol	Protocol ID	Encapsulation	Number of records
IP	800	Ether2	2
	806	Ether2	
IPv6	0x86DD	Ether2	1
Ipx802.2	0xE0E0	LLC	1
Ipx802.3	0xFFFF	SNAP	1
IpxEther2	0x8137	Ether2	2
	0x8138	Ether2	
IpxSnap	0x8137	SNAP	2
	0x8138	SNAP	
AppleTalk	0x809b	Ether2	4
	0x809b	SNAP	
	0x80F3	Ether2	
	0x80F3	SNAP	
DecLat	0x6004	Ether2	1
DecOther	0x6000	Ether2	9
	0x6001	Ether2	
	0x6002	Ether2	
	0x6003	Ether2	
	0x6005	Ether2	
	0x6006	Ether2	
	0x6007	Ether2	
	0x6008	Ether2	

Table continues…

Protocol	Protocol ID	Encapsulation	Number of records
	0x6009	Ether2	
NetBios	0xF0F0	LLC	1
PPPoE	0x8863	Ether2	2
	0x8864	Ether2	
RARP	0x8035	Ether2	1
SnaEther2	0x80D5	Ether2	1
sna802dot2	0x04xx	LLC	2
	xx04	LLC	
Vines	0xBAD	Ether2	1
XNS	0x600	Ether2	2
	0x807	Ether2	

IP routing and VLANs

Virtual Services Platform 9000 modules support IP routing on the following types of VLANs:

- Port-based VLANs
- Source IP subnet-based VLANs
- IP protocol-based VLANs
- Source MAC-based VLANs
- Management VLAN 4092: the VLAN comprising the VSP 9000 Management interface

VLAN implementation

This section describes default VLANs and the unassigned (null) VLAN on VSP 9000.

Default VLAN

Virtual Services Platform 9000 devices are factory-configured so that all ports are in a port-based VLAN called the default VLAN. Because all ports are in the default VLAN, the device behaves like a Layer 2 device. The VLAN ID of this default VLAN is always 1, and the default VLAN is always a port-based VLAN. You cannot delete the default VLAN.

Null VLAN

Internally, a Virtual Services Platform 9000 supports a placeholder for ports that is called a null port-based VLAN or unassigned VLAN. This concept is used for ports that are removed from all port-based VLANs. A port that is not a member of a port-based VLAN is a member of the null VLAN. Ports can belong to policy-based VLANs as well as to the null VLAN. If a frame does not meet the policy criteria and no underlying port-based VLAN exists, the port belongs to the null VLAN and the system drops the frame.

Because the null VLAN is an internal construct, you cannot delete the null VLAN.

VLAN configuration rules

The following list provides the VLAN configuration rules for the Virtual Services Platform 9000, which differ from the Ethernet Routing Switch 8000 series configuration rules:

- The Virtual Services Platform 9000 supports up to 4084 configurable VLANS. VLAN IDs range from 1 to 4084. VSP 9000 reserves VLAN IDs 4085 to 4094 for internal use.
- You can configure only one protocol-based VLAN for a certain protocol. VSP 9000 supports up to 16 protocol-based VLAN.
- VSP 9000 determines the VLAN membership of a frame by the following order of precedence, if applicable:
    1. IEEE 802.1Q tagged VLAN ID
    2. source MAC-based VLAN
    3. IP subnet-based VLAN
    4. protocol-based VLAN
    5. port-based VLAN default VLAN of the receiving port

VLAN MAC security

Use MAC security to control traffic from specific MAC addresses. You can also limit the number of allowed MAC addresses. You can enable this feature at two levels: globally and at the port level.

At the global level this feature is a filter mechanism to filter out (drop) packets that contain certain MAC addresses as the source or destination. You configure a set of MAC addresses. The system drops a packet that contains one of these configured addresses as the source or destination.

Port-level MAC security provides more flexibility over the global configuration. Port-level security applies to traffic for all VLANs received on that port.

For more information about MAC security, see *Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000,* NN46250-500.

Link aggregation and loop prevention

The following table provides a quick reference to the link aggregation and loop prevention differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series. For more detail about the Virtual Services Platform 9000 implementation, see the sections that follow.

Table 20: Link aggregation and loop prevention quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Loop prevention	
You can configure the ethertype for SLPP. In versions prior to 7.1, the default is 0x8104. In versions 7.1 and later, the default is 0x8102.	You cannot configure the ethertype for SLPP. The ethertype is 0x8102.
SMLT, MLT, and IST	
An IST session is not supported between ERS and VSP 9000.	-
Uses SMLT IDs.	Uses MLT IDs. Configure an MLT as an SMLT.
Can configure SLTs.	Configure an SMLT with a single local and remote link.
-	Hardware-based SMLT (clustering) with sub 20 ms failover.

For more information about SLPP fundamentals and configuration, see *Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000,* NN46250-500.

On the Ethernet Routing Switch 8000 series, you can configure the SMLT ID to be independent of the MLT ID; only the SMLT ID must match between the peer nodes. For example:

Node A — MLT ID 10, SMLT ID 100

Node B — MLT ID 12, SMLT ID 100

In this scenario, both nodes treat their MLT ID (10 or 12) as an SMLT. When the nodes exchange information about the SMLT over the IST control channel, they use SMLT 100 and map this ID to the local MLT ID.

Virtual Services Platform 9000 does not use SMLT IDs. Instead, configure an MLT with an MLT ID, and select SMLT as the type. Both nodes must have an MLT 10 to function as an SMLT. After you enable SMLT on both nodes, the link functions as an SMLT. Until you enable SMLT on both nodes, the link functions as a normal MLT. To create a single-link SMLT, enable SMLT on an MLT with a single local and remote link.

The following two figures show an SMLT and IST example for the Ethernet Routing Switch 8600 and the Virtual Services Platform 9000.

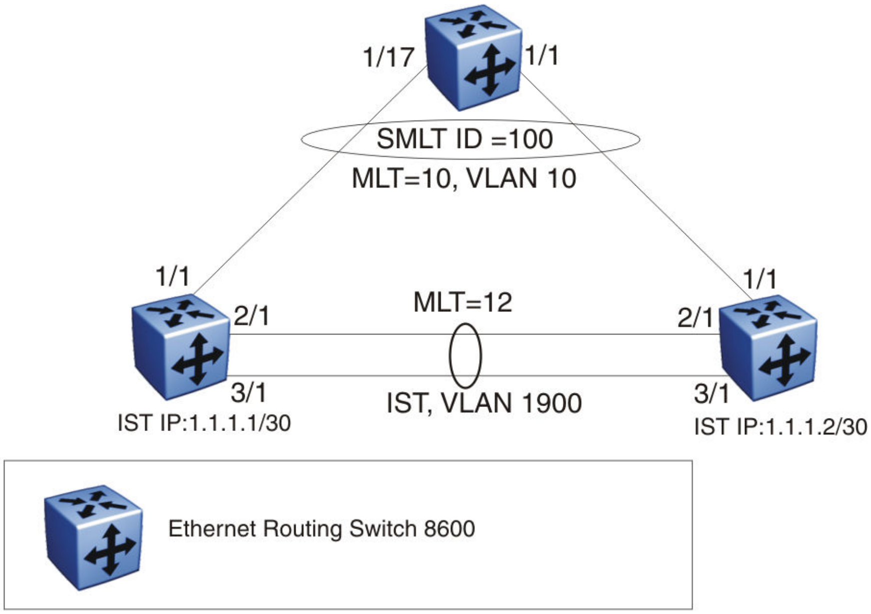

Figure 1: Ethernet Routing Switch 8600 Network topology

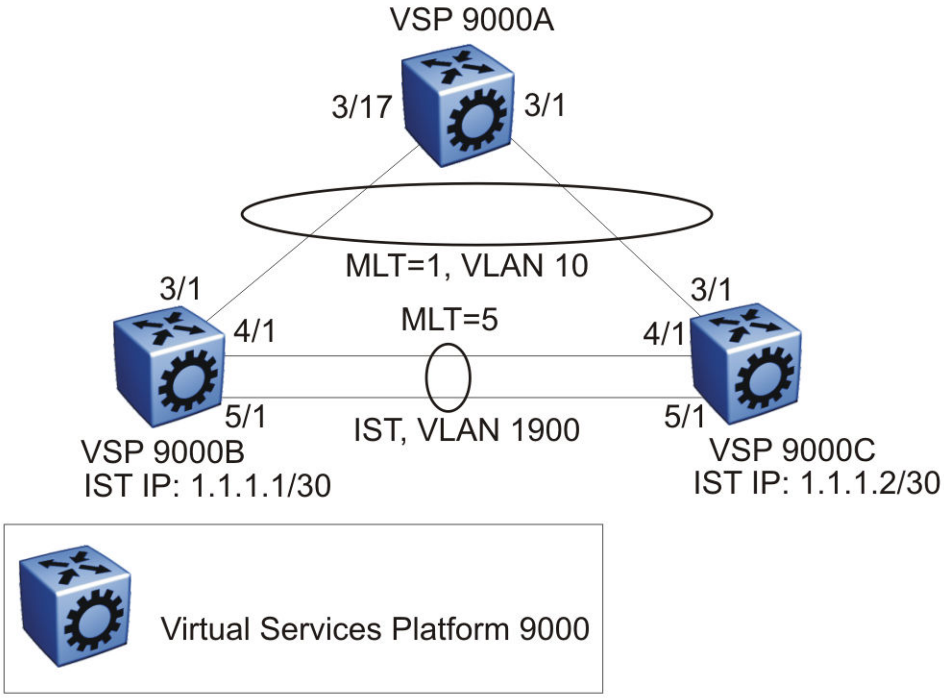

Figure 2: Virtual Services Platform 9000 Network topology

For information about MLT and SMLT, see *Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000,* NN46250-503. For information about Routed SMLT, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

For more information about SMLT configuration, see *Switch Clustering using Split-MultiLink Trunking (SMLT) with VSP 9000, ERS 8600/8800, 8300, and 5000 Technical Configuration Guide* , NN48500-518. For more information about the technical configuration guide, go to the Avaya Web site: http://support.avaya.com.

QoS and traffic filters

The following table provides a quick reference to the QoS and traffic filter differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series. For more detail about the Virtual Services Platform 9000 implementation, see the sections that follow.

Table 21: QoS and filters quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports Extended CP limit.	Does not support Extended CP limit.
You can configure the egress queue sets.	Does not support egress queue sets.
Supports MAC and VLAN QoS level.	Does not support MAC or VLAN QoS level. Uses ACLs for QoS assignments.
–	Uses fabric profiles to give preference to one type of traffic over another in times of over subscription.
Uses Access Control Templates (ACT) to define to list of attributes to filter on.	Uses hardware-based TCAM search, which eliminates the need for ACTs.
-	All packets sourced from the CP destined to an egress port bypass an egress filter.
Does not support a global filter action for control packet protection.	Supports a global filter action for control packet protection for deny mode.
Supports IPv4 and IPv6 Access Control Lists (ACL).	Supports IPv4 ACLs.
Supports the less-than or equal-to operator for Access Control Entries (ACE).	Does not support the less-than or equal-to operator for ACEs.
Supports the greater-than or equal-to operator for ACEs.	Does not support the greater-than or equal-to operator for ACEs.
Supports the not-equal-to operator for ACEs.	Does not support the not-equal-to operator for ACEs.
After a rule matches, the search does not stop. The search continues until all the rules are searched and all non-contradicting actions are applied for all hit ACEs. The only time a search is stopped is if you explicitly define an action stop on match.	After a rule matches, the search stops. The search is performed in the ascending order of ACE IDs.

Table continues…

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports a range value as in the following command: `filter acl ace ethernet 1 1 vlan-id eq 1-127`	Supports a mask operator as in the following command: `filter acl ace ethernet 11 vlan-id mask 1 0x7F` This command matches for VLAN ID equal to 1 and masks the lower 7 bits, which provides the range up to 127. If the given range is not in the bit boundary, you must create multiple ACEs.

Specifies TCP or UDP ports as part of the Layer 4 protocol attribute.
`filter acl ace protocol 770 20 tcp-src-port eq 20-23`
	Specifies the protocol type as part of the IP attribute.
	`filter acl ace ip 770 10 ip-protocol-type eq tcp`
	`filter acl ace protocol 770 10 src-port mask 20 0x3`

Does not use the concept of security and QoS ACEs.	Supports 1000 ACE rules for each ACL, divided between security ACEs (IDs 1-1000) and QoS ACEs (IDs 1001-2000).
Supports the egress-queue ACE action.	Does not support the egress-queue ACE action.
Supports the egress-queue-nnsc ACE action.	Does not support the egress-queue-nnsc ACE action.
Does not support the internal-qos ACE action.	Supports the internal-qos ACE action.
Supports the copy-to-primary ACE action.	Does not support the copy-to-primary ACE action.
Supports the copy-to-secondary ACE action.	Supports a PCAP ACE action.
-	Supports a new action of log , which logs the packet in buffer.

For configuration and conceptual information, and advanced filter examples, see *Configuring QoS and ACL-Based Traffic Filtering on Avaya Virtual Services Platform 9000,* NN46250-502.

Queuing

Virtual Services Platform 9000 supports Weighted Random Early Detection (WRED) on the fabric queues to provide congestion avoidance capabilities.

WRED is enabled for all queues except the highest priority expedited forwarding queue (EF). Expedited Forwarding Per Hop Bevhaviour (PHB) is a forwarding treatment for a DiffServ microflow when the transmission rate ensures that it is the highest priority and it experiences no packet loss for in-profile traffic.

With WRED, early discard starts when the queue reaches 75 percent of its maximum allowed length. If the queue reaches 100 percent of its maximum allowed length packets destined to it are tail dropped. WRED parameters are independent of the fabric profile and are not user configurable.

Internal QoS level

The internal QoS level or effective QoS level is a key element in the Virtual Services Platform 9000 QoS architecture. The internal QoS level specifies the kind of treatment a packet receives. Virtual Services Platform 9000 classifies every packet that enters and assigns it an internal QoS level.

Internal QoS levels map to the queues on a port. For example, for an access port the internal QoS level derives from the port QoS level. For Layer 3 trusted (core) ports, the system honors incoming DSCP or type of service (TOS) bits. The system assigns the internal QoS level using the ingress DSCP to QoS level map.

Ingress mappings

The system uses ingress maps to translate incoming packet QoS markings to the internal QoS level. The system uses the internal QoS level to classify packets.

The following logical table shows how the system performs ingress mappings for data packets and for control packets not destined for the Control Processor (CP).

Table 22: Data packet ingress mapping

DSCP	Layer 2 trusted	Layer 3 trusted and DiffServ enabled	IP packet	Routed packet	Ingress tagged	Internal QoS
x	No	x	No	x	x	Use port QoS
x	Yes	x	No	x	No	Use port QoS
x	Yes	x	No	x	Yes	Use ingress p-bits mapping
0x1B	x	x	Yes	x	x	4
0x23	x	x	Yes	x	x	5
0x29	x	x	Yes	x	x	5
0x2F	x	x	No	x	x	6
x	No	No	x	x	x	Use port QoS
x	No	Yes	Yes	x	x	Use ingress DSCP mapping
x	Yes	No	Yes	x	No	Use port QoS
x	Yes	No	Yes	x	Yes	Use ingress p-bits mapping

Table continues…

DSCP	Layer 2 trusted	Layer 3 trusted and DiffServ enabled	IP packet	Routed packet	Ingress tagged	Internal QoS
x	Yes	Yes	Yes	No	No	Use ingress DSCP mapping
x	Yes	Yes	Yes	No	Yes	Use ingress p-bits mapping
x	Yes	Yes	Yes	Yes	Yes	Use ingress DSCP mapping

The QoS level for control packets destined for the CPU is assigned internally to ensure timely packet processing and scaling numbers. You cannot configure the QoS level for these control packets. The system assigns the highest QoS-level to time-critical protocols.

CPU protection

Avaya Virtual Services Platform 9000 protects the CPU from Denial-of-Service (DOS) attacks through the following methods:

- CPU meters

  CPU meters are another mechanism to protect the CPU on the Control Processor (CP) module from becoming overloaded. The hardware counts every packet destined to each CPU over a specific time period. If the packet count exceeds the packet limit, the system drops the packets. Avaya limits the number of packets to each CPU on the CP module. You cannot configure CPU meters.

  CPU meters also provide packet priority scheduling. CPU meters use eight FIFO queues in FPGA. You cannot configure which packet types go into which queue. Each queue has a meter with packet limits. A scheduler services the eight queues, using a combination of strict priority and round-robin. Queues six and seven drain completely. Queues one through five use round-robin and queue zero uses best effort.
- port and MLT meters (CP Limit)

  Use port and MLT meters to configure the limit on the number of control and data exception packets that can enter on each port or MLT interface. You can configure port and MLT meters to shutdown the port or all ports in the MLT. If the number of packets exceeds the configured limit, the system generates a message in the log file. If enabled, the system shuts down the port or all ports in the MLT and raises an alarm. You can disable the port to clear the alarm. The default value is 8000 packets per second with no shutdown.
- protocol meters

  Protocol meters configure the limit on the number of control packets of specific packet types that can reach the CPU on the CP module. The system classifies every packet and assigns it an internal packet type. Protocol meters use the internal packet type to limit the number of each type of packet. You cannot configure protocol meters.

For more information about how to protect the CPU from DOS attacks, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

Traffic management profiles

The Avaya Virtual Services Platform 9000 provides different paths through the switch fabric for unicast and multicast traffic. You can configure the system to give preference to one type of traffic over the other in times of over subscription.

Over subscription occurs if the incoming traffic on a port or interface is more than the system can switch or route through the system. In these situations of high traffic flow, the Virtual Services Platform 9000 needs to drop traffic. You can control what traffic the system drops and what traffic the system switches or routes by using the fabric-profile boot configuration flag.

Use the boot configuration flag fabric-profile to configure preferences. You can select one of the following three profiles:

- balanced

  In the balanced profile, if the egress port is over subscribed and the unicast traffic is greater than 80% of line rate, the system limits multicast traffic to 20%.

- unicast optimized

  In the unicast optimized profile, if the egress port is over subscribed and the unicast traffic is greater than 90% of line rate, the system limits multicast traffic to 10%.

- multicast optimized

  In the multicast optimized profile, if the egress port is over subscribed and the unicast traffic is greater than 70% of line rate, the system limits multicast traffic to 30%.

Note:

9012QQ-2 traffic profiles depend on particular traffic flows, distribution, and may deviate from the exact numbers due to the hashing architecture.

After you make this configuration change, you need to restart the system. After the restart, the device applies the correct fabric-profile configuration.

Oversubscription Behavior

There are eight unicast queues per egress port. There are one EF queue and seven Weighted RED queues. The EF traffic limits at 10% of line rate. The other queues have a minimum guaranteed bandwidth as shown in the table below. Multicast traffic uses the remaining bandwidth. If there is no multicast traffic, the unicast traffic is distributed across the other WRED queues.

COS #	Behavior	Queue Type	Minimum bandwidth guarantee Balanced	Minimum bandwidth guarantee Unicast	Minimum bandwidth guarantee Multicast
7	EF	Strict Priority	10%	10%	10%
6	AF	WRED	25%	27%	20%
5	AF	WRED	15%	20%	15%
4	AF	WRED	12%	12%	10%
3	AF	WRED	8%	9%	7%
2	AF	WRED	6%	7%	5%
1	AF	WRED	4%	5%	3%
0	Best Effort	WRED	0%	0%	0%

ACTs and ACLs

Access Control Templates (ACT) define the list of attributes to filter on. Hardware-based TCAM search in VSP 9000 eliminate the need for ACTs. You can configure ACEs with a number of defined attributes in a search.

ACLs define global actions and default actions. The following table compares the actions in both implementations.

Table 23: Global action comparison

	Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Mirror	Port, port list, MLT, or VLAN	Port, port list, MLT, or VLAN
IP FIX	Supported	Supported
Count	Supported with a CLI parameter	Supported. Statistics collected by default.
Drop mode	Permit/Deny	Permit/Deny

Table 24: Default action comparison

	Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Drop mode	Permit/Deny	Permit/Deny
Control packet protection	Not supported	Supported with an ACLI configuration for Deny mode
Count	Supported with a CLI parameter	Supported. Statistics collected by default.

Access control entries

The system supports a maximum of 16,000 ACEs globally and a maximum of 1,000 ACEs for each individual ACL.

Operators

The Virtual Services Platform 9000 supports the following operators:

- Equal-to

  This rule operator looks for an exact match with the field defined. If the field matches exactly with the rule, the system returns a match (hit). If the rule does not match, the search continues and at the end of the search the system returns a miss.

- Mask

  ACL-based filters provide the mask operator to match on Layer 2, Layer 3, and Layer 4 packet fields. Use the mask operator to mask bits in packet fields during a search or to match on a partial value of a packet field. This section provides examples of the mask operator.

  If a mask bit is set to 1, it means the mask bit is not part of the match criteria (treated as do not care), and a mask bit of 0 means that the value represented is part of the match criteria. You can use the mask operator for the following attributes:

  - source MAC address
  - destination MAC address
  - VLAN ID
  - Dot1p
  - source IP address
  - destination IP address
  - DSCP
  - Layer 4 source port
  - Layer 4 destination port
  - TCP flags

  The ACL and ACE configuration syntax for a mask is similar to how you use the equal operator except that you must provide the mask value. As part of the configuration you can specify a mask value (number) to represent the bits to mask in the attribute. You can define a mask in different ways depending on the attribute you need to mask. If you use a decimal number for the mask, the mask value applies to the least significant bits on that attribute. For example, a mask of 24 used with an IP address is the same as a mask of 0.255.255.255, and a mask of 24 used with a MAC address is the same as 0x000000ffffff. A mask of 16 used with an IP address is the same as a mask of 0.0.255.255, and a mask of 32 used with a MAC address is the same as 0x0000ffffffff.

The following table explains the mask operator for MAC addresses.

Table 25: Mask operator for MAC address

Rule	Result

`filter acl ace ethernet 10 10 dst-mac mask 0x01:00:5e:00:00:01 24`
which is the same as
`filter acl ace ethernet 10 10 dst-mac mask 01:00:5e:00:00:01 0x000000ffffff`
	The rule matches only on the most significant 24 bits as they are not masked, for example, 01:00:5e, and does not care about the least significant 24 bits because they are masked. The least significant 24 bits can have a value of 00:00:00 - FF:FF:FF.

`filter acl ace ethernet 10 10 dst-mac mask 0x01:00:5e:00:00:01 0xFFFFFFFF0000`	The rule matches only on the least significant 16 bits because they are not masked, for example, 00:01, and does not care about the most significant 32 bits because they are masked. The most significant 32 bits can have a value of 00:00:00:00 – FF:FF:FF:FF.
`filter acl ace ethernet 10 10 dst-mac mask 0x01:00:5e:00:00:01 0xFF00FF0000FF`	The rule matches only on the unmasked bits, for example, 0xXX:00:XX:00:00:XX. The rule matches only on the bits not masked, for example, all the zeroes and the x represents a do not care (0xXX: 00:XX:00:00:XX)

The following table explains the mask operator for IP addresses.

Table 26: Mask operator for IP address

Rule	Result

`filter acl ace ip 10 10 src-ip mask 2.10.10.12 24`
which is the same as
`filter acl ace ip 10 10 src-ip mask 2.10.10.12 0.255.255.255`
	The rule matches only the most significant 8 bits, for example, 2, and does not care about the value of the remaining 16 bits because they are masked, for example, 10.10.12. Packets with a source IP address of 2.15.16.122 or 2.3.4.5 match on the filter rule while packets with a source IP address of 3.10.10.12 and 4.10.10.12 do not match on the filter rule.
	The mask appears as 0.255.255.255 in the command output for **`show filter acl config`**.

`filter acl ace ip 10 10 src-ip mask 3.4.5.6 255.255.255.0`	The rule matches only the least significant 8 bits, for example, 6, and does not case about the most significant 24 bits, 3.4.5. Packets with a source IP address of 17.16.5.6 or 192.168.1.6 match on the filter rule while packets with a source IP address of 3.4.5.4 or 3.4.5.7 do not match on the filter rule. If you use the mask value in decimal format <0–31> with an IP address, do not interpreted the mask as or associated with the subnet-mask in IP configuration. This mask value represents the least significant bits to treat as "do not care" and are ignored when matching filter attributes.

The following table explains the mask operator for Layer 4 source port.

Table 27: Mask operator for Layer 4 source port

Rule	Result
`filter acl ace protocol 10 10 src-port mask 80 0xF`	The filter rule matches on Layer 4 source port 80 (1010000). The mask value 0xF (1111) masks the least significant 4 bits, which means source port 81 (1010001) through 95 (1011111) also match this filter rule. This means the range 80–95 is a match on this rule.

The following table demonstrates the resulting action based on mask configuration and example packets.

Table 28: Mask operator configuration examples

Filter configuration	Address examples that match the filter	Address examples that do not match the filter

Ethernet mask:
`filter acl 1000 type inport`
`filter acl port 1000 6/5,9/11`
`filter acl ace 1000 12`
`filter acl ace ethernet 1000 12 src-mac mask 00:00:11:11:16:00 0x00ff000000f0`
`filter acl ace action 1000 12 permit count`
`filter acl ace 1000 12 enable`
	Source MAC:
	00:01:11:11:16:10
	00:10:11:11:16:f0 00:1f:11:11:16:10 00:ff:11:11:16:f0
	00:00:11:11:16:60
	00:e6:11:11:16:e0
		Source MAC:
		00:00:11:11:16:01
		00:ff:11:11:16:f1

`filter acl ace 1000 1000`
`filter acl ace ethernet 1000 1000 dst-mac mask 00:00:00:64:16:00 0x00000060001f`
`filter acl ace action 1000 1000 deny log count`
`filter acl ace 1000 1000 enable`
	Destination MAC:
	00:00:00:64:16:01
	00:00:00:04:16:01
	00:00:00:24:16:1f
	00:00:00:64:16:1f
	00:00:00:44:16:10
	00:00:00:04:16:05
		Destination MAC:
		00:00:00:24:16:20
		00:00:00:64:16:20
		00:00:00:63:16:01
		00:00:00:65:16:01

IP mask (dotted decimal notation):
`filter acl 10 type outport`
`filter acl port 10 5/13`
`filter acl ace 10 11`
`filter acl ace ethernet 10 11 ether-type eq ip`
`filter acl ace ip 10 11 src-ip mask 192.168.4.0 0.0.0.31`
`filter acl ace action 10 11 permit count`
`filter acl ace 10 11 enable`
	Source IP: 192.168.4.1
	192.168.4.10
	192.168.4.30
	192.168.4.31
		Source IP:
		192.168.3.1
		192.168.4.32

`filter acl ace 10 12`
`filter acl ace ethernet 10 12 ether-type eq ip`
`filter acl ace ip 10 12 dst-ip mask 192.168.7.0 0.0.0.3`
`filter acl ace action 10 12 deny count`
`filter acl ace 10 12 enable`
	Destination IP:
	192.168.7.1 192.168.7.3
		Destination IP:
		192.168.7.4
		192.168.7.5

Table continues…

Filter configuration	Address examples that match the filter	Address examples that do not match the filter

IP mask (decimal notation):
`filter acl 10 type outport`
`filter acl port 10 5/13`
`filter acl ace 10 11`
`filter acl ace ethernet 10 11 ether-type eq ip`
`filter acl ace ip 10 11 src-ip mask 192.168.4.0 5`
`filter acl ace action 10 11 permit count`
`filter acl ace 10 11 enable`
	Source IP: 192.168.4.1
	192.168.4.10
	192.168.4.30
	192.168.4.31
		Source IP:
		192.168.3.1
		192.168.4.32

`filter acl ace 10 12`
`filter acl ace ethernet 10 12 ether-type eq ip`
`filter acl ace ip 10 12 dst-ip mask 192.168.7.0 2`
`filter acl ace action 10 12 deny count`
`filter acl ace 10 12 enable`
	Destination IP:
	192.168.7.1 192.168.7.3
		Destination IP:
		192.168.7.4
		192.168.7.5

Protocol mask:
`filter acl 901 type inport`
`filter acl port 901 6/2`
`filter acl ace 901 1`
`filter acl ace ip 901 1 ip-protocol-type eq tcp`
`filter acl ace protocol 901 1 src-port mask 256 0xff`
`filter acl ace action 901 1 deny count`
`filter acl ace 901 1 enable`
This mask implies packets with TCP source port 256–511 match the filter, while 0–255 and > 511 miss the filter.
	TCP source port 256
	TCP source port 356
	TCP source port 511
		TCP source port 255
		TCP source port 512

Attributes

This section identifies the support differences between the attributes and operators to which an ACE rule can apply for each product.

Layer 2 Ethernet attributes

This section identifies the support differences between the Layer 2 Ethernet attributes and operators to which an ACE rule can apply for each product.

Table 29: Layer 2 Ethernet attributes

Attribute	ERS 8000 operator	VSP 9000 operator
Destination MAC	Equal-to, Less-than or equal-to, greater-than or equal-to, not-equal-to	Equal-to, Mask
Source MAC	Equal-to, Less-than or equal-to, greater-than or equal-to, not-equal-to	Equal-to, Mask

Table continues…

Attribute	ERS 8000 operator	VSP 9000 operator
VLAN ID	Equal-to	Equal-to, Mask
.1p bits	Equal-to, not-equal-to	Equal-to, Mask
Ether type	Equal-to, not-equal-to	Equal-to
Source port	Equal-to	Equal-to

The following commands show a Layer 2 Ethernet attribute example for each product.

ERS 8000: **`filter acl ace ethernet 1 1 vlan-id eq 4,5,6,7`**

VSP 9000: **`filter acl ace ethernet 1 1 vlan-id mask 4 0x3`**

The mask operator masked the last two bits of the VLAN ID as “don’t care” , which matches on values 4 , 5 , 6, and 7.

Layer 3 IP attributes

This section identifies the support differences between the Layer 3 IP attributes and operators to which an ACE rule can apply for each product.

Table 30: Layer 3 IP attributes

Attribute	ERS 8000 operator	VSP 9000 operator
Source IP	Equal-to, Less-than or equal-to, greater-than or equal-to, not-equal-to	Equal-to, Mask
Destination IP	Equal-to, Less-than or equal-to, greater-than or equal-to, not-equal-to	Equal-to, Mask
DSCP	Equal-to, not-equal-to	Equal-to, Mask
IP fragmentation	Equal-to	Equal-to Only checks for noFragment and anyFragment
Packet type	Any	Equal-to
IP options	Equal-to, not-equal-to	Equal-to

To use a Layer 3 attribute on VSP 9000, you must create an ACE rule with ether-type equal-to ip.

The following commands show a Layer 3 IP attribute example for each product.

ERS 8000:

```
filter acl 30 type inVlan act 1  name "Subnet30-IN"
filter acl vlan 30 30
filter acl ace 30 10 name "NoSpoofing"
filter acl ace action 30 10 deny
filter acl ace action 30 10 deny stop-on-match enable
filter acl ace ip 30 10 src-ip ne 155.247.30.1-155.247.30.255
flter acl ace 30 10 enable
```

VSP 9000:

```
filter acl 30 type inVlan name "Subnet30-IN"
filter acl set 30 default-action deny
filter acl vlan 30 30
filter acl ace 30 10 name "NoSpoofing"
filter acl ace action 30 10 permit
filter acl ace ethernet 30 10 ether-type eq ip
filter acl ace ip 30 10 src-ip mask 154.247.30.1  0.0.0.255
filter acl ace 30 10 enable
```

Layer 4 protocol attributes

This section identifies the support differences between the Layer 4 protocol attributes and operators to which an ACE rule can apply for each product.

Table 31: Layer 4 protocol attributes

ERS 8000 Attribute	ERS 8000 operator	VSP 9000 attribute	VSP 9000 operator
TCP source port	Equal-to, Less-than or equal-to, greater-than or equal-to, not-equal-to	Source port	Equal-to, Mask
TCP destination port	Equal-to, Less-than or equal-to, greater-than or equal-to, not-equal-to	Destination port	Equal-to, Mask
UDP source port	Equal-to, Less-than or equal-to, greater-than or equal-to, not-equal-to	TCP flags	Equal-to, Mask
UDP destination port	Equal-to, Less-than or equal-to, greater-than or equal-to, not-equal-to	ICMP message type	Equal-to
TCP flags	Match-any, match-all	—	—
ICMP message type	Equal-to, not-equal-to	—	—

VSP 9000 uses the protocol type in the IP attribute to define the protocol type rather than specifying TCP or UDP ports.

The following commands show a Layer 4 protocol attribute example for each product.

ERS 8000:

```
filter acl ace 770 20 name "AllowHosts-TCP"
filter acl ace action 770 20 permit
filter acl ace ip 770 20 src-ip eq 129.32.20.4,129.32.20.102-129.32.20.103
filter acl ace protocol 770 20 tcp-src-port eq 20-23
filter acl ace 770 20 enable
```

VSP 9000:

```
filter acl ace 770 10 name "AllowHosts-TCP"
filter acl ace action 770 10 permit
filter acl ace ethernet 770 10 ether-type eq ip
filter acl ace ip 770 10 src-ip eq 129.32.20.4
filter acl ace ip 770 10 ip-protocol-type eq tcp
filter acl ace protocol 770 10 src-port mask 20 0x3
filter acl ace 770 10 enable
filter acl ace 770 20 name "AllowHosts-TCP_102-103”
filter acl ace action 770 20 permit
filter acl ace ethernet 770 20 ether-type eq ip
filter acl ace ip 770 20 src-ip mask 129.32.20.102 0.0.0.1
filter acl ace ip 770 20 ip-protocol-type eq tcp
filter acl ace protocol 770 20 src-port mask 20 0x3
filter acl ace 770 20  enable
```

Actions

The types of actions filter configuration can execute are split into two categories for VSP 9000:

- Security actions supported by the ACE IDs in the range of 1 to 1,000
- QoS actions supported by the ACE IDs in the range of 1,001 to 2,000

Filter rules supporting Security actions and QoS actions are stored separately. When an ACL filter applies to a traffic flow, the Virtual Services Platform 9000 performs a parallel search on both Security and QoS ACE lists, resulting in distinct and non-conflicting actions. The following table provides the supported Virtual Services Platform 9000 actions.

Table 32: Security ACE Actions

Security ACE Actions	User supplied parameters	Comments
Mode	Permit or Deny	This action applies to both ingress and egress ACLs.
PCAP	None	Packet Capture: A copy of the packet is sent to the secondary CPU. This action applies to both ingress and egress ACL.
Log	None	Packet header is logged and it can be seen using an ACLI command along with time stamp and actions taken. This action applies to both ingress and egress ACLs.
• Redirect Next-Hop • Unreachable	IP address, Mode	Redirects the packet to the user supplied IP address if the user supplied IP address is unreachable, the user may specify a mode action. If mode is Deny, the packet is dropped; else the packet forwarded as normal. These actions apply to ingress ACLs only.
MLT index	Index value	The user supplied index overrides the computed hash ID based on fields within the packet. This action applies to ingress ACLs.

Table continues…

Security ACE Actions	User supplied parameters	Comments
Count	None	Collect ACE statistics. This action applies to ingress and egress ACLs.
Mirror	Port or list of ports, VLAN-ID, MLT-ID, or IP address	This action applies to ingress and egress ACLs.
IPFIX	None	Configures IPFIX metering. This action applies to ingress ACLs.

Table 33: QoS ACE Actions

QOS ACE Actions	User supplied parameters	Comments
Remark	• DCSP • .1P • Internal-qos	This action applies to ingress ACLs.
Police	Profile ID	The policer profile ID refers to a user defined profile. This action applies to ingress ACLs.
Count	None	This action applies to ingress and egress ACLs.
Log	None	Packet header is logged and it can be seen using an ACLI command along with time stamp and actions taken. This action applies to both ingress and egress ACLs.

TACACS+

The following table provides a quick reference to the Terminal Access Controller Access Control System plus (TACACS+) differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series. For more detail about the Virtual Services Platform 9000 implementation, see the section that follows.

Table 34: TACACS+ quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports Point-to-Point Protocol (PPP) functionality with TACACS+.	Does not support PPP with TACACS+.

Point-to-Point protocol (PPP)

Virtual Services Platform 9000 does not support Point-to-Point protocol (PPP) with TACACS+.

Ethernet Routing Switch 8800 supports PPP with TACACS+. PPP is a basic protocol at the data link layer that provides its own authentication protocols, with no authorization stage. PPP is often used to form a direct connection between two networking nodes.

Remote mirroring

The following table provides a quick reference to the remote mirroring differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series.

Table 35: Mirroring quick reference

Ethernet Routing Switch 8000 Series	Virtual Services Platform 9000
Supports Layer 2 remote mirroring	Supports Layer 2 and Layer 3 remote mirroring

Virtual Services Platform 9000 supports Layer 3 remote mirroring for ports and flows. Layer 3 remote mirroring monitors traffic from multiple network devices across an IP network, and sends that traffic in an encapsulated form to the destination analyzers. For configuration information, see *Troubleshooting Avaya Virtual Services Platform 9000,* NN46250-700.

IP routing

The following table provides a quick reference to the IP routing protocol differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series.

Table 36: IP routing protocol quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports BGP 4–byte AS only if the peers use the same configuration.	Supports mixed peer communication of 4– and 2–byte AS, as documented in RFC4893.

BGP 4–byte AS

Ethernet Routing Switch 8000 series supports communication between peers of the same type only. If a new 4–byte AS has to communicate with an old 2-byte AS, you assign a 2-byte AS number to the new AS. ERS currently supports communication between the following peer types only:

- 2-byte peer to 2-byte peer
- 4-byte peer to 4-byte peer

Virtual Services Platform supports communication between mixed peers, which means you can have a combination of peers that use a 4–byte AS and peers that use a 2–byte AS.

For more information, see *Configuring BGP Services on Avaya Virtual Services Platform 9000,* NN46250-507.

IPv6 routing

The following table provides a quick reference to the IPv6 routing differences between Ethernet Routing Switch 8000 series and Virtual Services Platform 9000 for this release.

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports IPv6 filters	Does not support IPv6 filters
Supports IPv6 MLD router	Supports MLD host mode but you cannot configure it

When you configure an ACL on Ethernet Routing Switch 8000 series, you can specify a packet type of either IPv4 or IPv6. You can also configure an ACE to monitor for specific IPv6 header attributes. Virtual Services Platform 9000 does not support IPv6 filtering in this release.

Ethernet Routing Switch 8000 series supports configurable Multicast Listener Discovery (MLD) router. You can enable and configure MLD on a VLAN or a port. On Virtual Services Platform 9000, the current release supports MLD host mode but you cannot enable or configure it either globally, or on an interface.

IP multicast

The following table provides a quick reference to the IP multicast differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series.

Table 37: IP multicast quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Partial support for IGMPv3	Full support for IGMPv3 RFC3376

IGMPv3

Virtual Services Platform 9000 is fully-compliant with IGMPv3 RFC3376. The enhancements over the partial support on Ethernet Routing Switch 8000 series include the following:

- Adds support for source filtering — The system can report interest in receiving packets from *only* a specific source address (INCLUDE), from all *but* specific source addresses (EXCLUDE), or sent to specific multicast addresses. IGMPv3 interacts with PIM-SM, PIM-SSM, and snooping to provide source filtering.
- Allows multiple sources for the same group in the ssm-map
- Enables you to configure the IGMP version of an interface to version 3 regardless of the PIM or snooping mode
- Adds explicit host tracking to allow fast-leave functionality

For more information, see *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504.

Virtualization support

You can configure multicast routing support with the Virtual Routing and Forwarding (VRF) Lite feature and you can use VRF Lite to emulate many virtual routers with one router.

VSP 9000

Multicast virtualization support includes:

- IGMP snooping
- IGMP in Layer 2 Virtual Services Networks (VSNs)
- IGMP in Layer 3 VSNs

IGMP snooping, Layer 2 VSN multicast over SPBM, and Layer 3 VSN multicast over SPB all support full HA.

Note:

You cannot configure PIM on a VRF for VSP 9000.

ERS 8800

In ERS 8800, VRF Lite virtualizes the following multicast protocols:

- IGMP
- PIM-SM
- PIM-SSM

For more information, see *Avaya Virtual Services Platform 9000 Configuration — IP Multicast Routing Protocols*.

Shortest Path Bridging MAC

The following table provides a quick reference to the Shortest Path Bridging MAC (SPBM) differences between Virtual Services Platform 9000 and Ethernet Routing Switch 8000 series. For more detail about the Virtual Services Platform 9000 implementation, see the sections that follow.

Table 38: SPBM quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports IP VPN Lite over SPBM.	Does not support IP VPN Lite over SPBM.
Reserves 519 multicast group IDs (MGIDs) for SPBM operation.	Reserves 100 mulitcast group IDs (MGIDs) for SPBM operation.
Does not require you to configure C-VLANs on the IST MLT.	Requires the inclusion of IST MLT peer switches in the C-VLAN.
If you enable SPB and the VLAN has an I-SID, then Traffic can pass between single-homed VLANs attached to IST peers if the IST is down.	Traffic cannot pass between single-homed VLANs attached to opposite IST peers if the IST is down. (Local forwarding is not effected.)

Table continues…

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Decapsulates MAC-in-MAC traffic at the primary BEB or secondary BEB irrespective of whether the traffic is from the primary B-VLAN or secondary B-VLAN.	Decapsulates MAC-in-MAC traffic at the primary BEB from the primary B-VLAN and traffic at the secondary BEB from the secondary B-VLAN. Requires the IST to be up to pass traffic between both IST switches for single-homed VLANs.
Supports NNI functionality on all R and RS module ports.	Supports SPBM NNI on the 9024XL module.
Disables tagging if you delete IS-IS from an interface.	Keeps tagging enabled, if you delete IS-IS from an interface.
The **`show isis spbm multicast-fib`** command for both the primary B-VLAN on the secondary IST switch, and the secondary B-VLAN on the primary IST switch, lists the following: • All active single-homed UNIs • An SMLT if all ports with the SMLT are down on the partner (or the partner is down), and at least one port within the SMLT is up locally • An SLT if the SLT port is down on the partner (or the partner is down) but up locally	The **`show isis spbm multicast-fib`** command for both the primary B-VLAN on the secondary IST switch, and the secondary B-VLAN on the primary IST switch, lists no UNI ports if the partner is up, but all UNI ports if the partner is down.
Does not have the **`spbm-tunnel-as-mac`** option for **`show vlan mac-address-entry`** command and the **`show ip arp`** command.	Adds the option **`spbm-tunnel-as-mac`** to the **`show vlan mac-address-entry`** command and the **`show ip arp`** command.

IP VPN Lite

Virtual Services Platform 9000 does not support IP VPN-Lite over SPBM.

Ethernet Routing Switch 8800 supports IP VPN-Lite over SPBM, which switches IP-in-IP packets in the SPBM core.

Multicast group ID

SPBM requires MGIDs for proper operation. The multicast group ID (MGID) is a hardware mechanism the switch uses to send data to several ports simultaneously. Instead of sending the data to a specific port number, the device directs the data to an MGID. The switch maintains a table that maps MGIDs to their member ports. Both virtual LAN (VLAN) and IP multicast (IPMC) use MGIDs. The system also reserves a small number of MGIDs.

After you enable SPB on the switch, Virtual Services Platform 9000 reserves 100 MGIDs for SPBM operation. Therefore, the number of MGIDs on the system available for IP multicast traffic is reduced by 100. Ethernet Routing Switch 8800 reserves 519 MGIDs.

To determine how many MGIDs are available on Virtual Services Platform 9000, enter `show sys mgid-usage`.

```
VSP-9012:1#show sys mgid-usage
        Number of MGIDs used for VLANs : (65)
        Number of MGIDs used for multicast : (0)
        Number of MGIDs used for SPBM : (1)
        Number of MGIDs remaining for VLANs : (4031)
        Number of MGIDs remaining for multicast : (7900)
        Number of MGIDs remaining for SPBM : (99)
```

Layer 2 VSN

Different from Ethernet Routing Switch 8800 SPBM implementation, on Virtual Services Platform 9000 when you map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID), you must include the IST MLT peer switches in the VLAN, which means all IST member ports must be members of this C-VLAN.

In ERS 8800, you cannot configure IST MLT to be part of a C-VLAN.

For more information and configuration examples, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

C-VLANs single-homed to one IST peer

On the VSP 9000, if you have a C-VLAN single-homed to an IST peer and the IST link goes down, traffic from the single-homed C-VLAN cannot reach the far-end peer or any other single-homed C-VLANs connected to it. This operation is the same whether SPBM is enabled or disabled.

On the ERS 8800/8600, the IST behavior is the same as the VSP 9000 when SPBM is disabled. However, when SPBM is enabled on the ERS 8800/8600, C-VLANs single-homed to an IST peer can still reach the far-end IST peer and connected C-VLANs even if the IST is down. With the ERS 8800/8600 implementation of SPBM, the IST ports are not members of the C-VLAN. As a result, the traffic still routes to the far-end peer through the SPBM cloud after IST failure.

This forwarding behavior is not possible with SPBM on the VSP 9000 because the IST ports must be members of the C-VLAN.

NNI packet decapsulation

On Virtual Services Platform 9000, when the primary Backbone Edge Bridge (BEB) receives a MAC-in-MAC packet from the secondary B-VLAN, the BEB does not decapsulate the packet and send it out the User-to-Network Interface (UNI) port. Instead, the MAC-in-MAC packet forwards on the Network to-Network Interface (NNI) port over the IST to the secondary BEB. The MAC-in-MAC packet only forwards to the secondary BEB if the outgoing interface exists in the multicast Forwarding Information Base (FIB) table for the secondary B-VLAN. When the MAC-in-MAC packet it decapsulates the MAC-in-MAC packet and sends it back to the primary BEB, which then forwards it to the UNI ports.

This process is similar on the secondary BEB: if the secondary BEB receives a MAC-in-MAC packet from the primary B-VLAN, it forwards the packet to the primary BEB for decapsulation and forwarding.

On Ethernet Routing Switch 8800, the primary and secondary BEB can decapsulate any MAC-in-MAC traffic, whether the traffic is from the secondary B-VLAN or the primary B-VLAN.

NNI support

Virtual Services 9000 supports Network-to-Network Interface (NNI) ports on the 9024XL module and the 9048XS-2 module. NNI ports connect to the core. The 9024XL module provides 24 10GBASE-X, small form factor pluggable plus (SFP+) ports, and the 9048XS-2 module provides 48 10GBASE-X SFP+ ports. SFP+ is an enhanced version of the SFP that supports rates of up to 10 Gbps.

Virtual Services Platform 9000 is primarily intended for 10 GbE aggregations. However, using an SFP+ you can operate a 10 GbE port on the 9024XL module or 9048XS-2 module in either 1 GbE or 10 GbE mode. The port reverts to the highest common speed.

Ethernet Routing Switch 8800 supports NNI on all R and RS modules.

For more information on supported SFP+, see *Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000,* NN46250-305.

For more information on SPBM, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 9000,* NN46250-510.

VLAN tagging

The Virtual Services Platform 9000 and Ethernet Routing Switch 8800/8600 both support the IEEE 802.1Q specification for tagging frames. On both Avaya Virtual Services Platform 9000 and Ethernet Routing Switch 8800, when you add IS-IS to an interface, with the command `isis`, the device automatically enables tagging for that interface.

The difference occurs when the user deletes IS-IS from an interface, with the command `no isis`:

- For ERS 8800, the device disables tagging.
- For VSP 9000, tagging remains enabled.

Multicast FIB rules

IS-IS programs B-MAC addresses into the B-VLAN Forwarding Information Bases (FIBs) instead of the traditional VLANs flooding and learning approach. Each SPBM network instance is associated with at least one backbone VLAN (B-VLAN) in the core SPBM network.

The IS-IS link-state database carries B-MAC addresses. After the IS-IS link-state database discovers and stores the network topology, each node calculates shortest path trees for each source node, so that a unicast path now exists from every node to every other node. With this information, each node populates unicast information received from SPBM into the FIB for forwarding purposes. The multicast FIB is not produced until virtual services are configured and learned. The multicast addresses are built out of two pieces: the instance-ID (nickname) and the I-SID ID converted to hexadecimal format to form the multicast MAC address.

SPBM runs in the core on the ports that connect to the core. These ports are Network to-Network Interface (NNI) ports. Ports facing a customer VLAN are User-to-Network Interface (UNI) ports.

The following table describes the differences in information within the multicast FIB between the VSP 9000 and ERS 8800:

	VSP 9000	ERS 8800
Primary B-VLAN — Primary IST switch	Lists all active User-to-Network Interface (UNI) ports on the customer (C-VLAN.)	Lists all active User-to-Network Interface (UNI) ports on the customer (C-VLAN.)
Primary B-VLAN — Secondary IST switch	No UNI ports are listed — If the partner is UP (either the IST is up or reachable through IS-IS). All UNI ports are listed – If the partner is DOWN (IST is down and is not reachable through IS-IS).	Lists all active single-homed UNIs. Lists a Split MultiLink Trunking (SMLT) if all ports within that SMLT are down on the partner and at least one port within the SMLT is up locally.

Table continues…

	VSP 9000	ERS 8800
		Lists a Single Link Trunking (SLT) – if the SLT port is down on the partner and up locally.
Secondary B-VLAN — Primary IST switch	No UNI ports are listed - If the partner is UP (either IST is up or reachable through IS-IS) All UNI ports are listed – If the partner is DOWN (IST is down and is not reachable through IS-IS)	Lists all active single-homed UNIs. Lists an SMLT – If all ports within that SMLT are down on the partner and at least one port within the SMLT is up locally. Lists an SLT – If the SLT port is down on the partner and up locally.
Secondary B-VLAN — Secondary IST switch	Lists all active UNI ports on the C-VLAN.	Lists all active UNI ports on the CVLAN.

Show commands

VSP 9000:

The commands **show vlan mac-address-entry** and **show ip arp** produce different outputs for Virtual Services Platform 9000 and Ethernet Routing Switch 8800.

For Virtual Services Platform 9000, the command **show vlan mac-address-entry**:

- Shows the actual outgoing interface for a customer MAC (C-MAC).
- Shows the tunnel endpoint MAC if that endpoint is not a virtual MAC.

```
VSP-9012:1>show vlan mac-address-entry
================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                           SMLT
ID   STATUS     ADDRESS            INTERFACE  REMOTE TUNNEL
--------------------------------------------------------------------------------
216  self       58:16:26:b1:7a:02  Port-cpp   false  -
304  learned    00:14:0d:2a:02:0c  Port-11/18 false  C1
304  learned    00:1b:4f:64:fa:0c  Port-11/13 false  C0
304  learned    00:80:2d:30:fa:03  Port-11/18 false  ER0
304  self       58:16:26:b1:7a:03  Port-cpp   false  -
404  learned    00:14:0d:2a:02:10  Port-11/18 false  C1
404  learned    00:1b:4f:64:fa:10  Port-11/13 false  C0

7 out of 7 entries in all fdb(s) displayed.
```

To see the virtual MAC, VSP 9000 adds the option **spbm-tunnel-as-mac** to the **show vlan mac-address-entry** command.

```
VSP-9012:1>show vlan mac-address-entry spbm-tunnel-as-mac
================================================================================
                                    Vlan Fdb
================================================================================
VLAN            MAC                           SMLT
ID   STATUS     ADDRESS            INTERFACE  REMOTE TUNNEL
--------------------------------------------------------------------------------
216  self       58:16:26:b1:7a:02  Port-cpp   false  -
304  learned    00:14:0d:2a:02:0c  Port-11/18 false  00:00:bc:b0:00:04
304  learned    00:1b:4f:64:fa:0c  Port-11/13 false  00:00:bc:b0:00:03
```

```
304  learned    00:80:2d:30:fa:03  Port-11/18 false  00:00:be:b0:00:07
304  self       58:16:26:b1:7a:03  Port-cpp   false  -
404  learned    00:14:0d:2a:02:10  Port-11/18 false  00:00:bc:b0:00:04
404  learned    00:1b:4f:64:fa:10  Port-11/13 false  00:00:bc:b0:00:03

7 out of 7 entries in all fdb(s) displayed.
```

The following table describes the fields in the output for the **`show vlan mac-address-entry`** command.

Parameter	Description
VLAN ID	Indicates the VLAN for this MAC address.
STATUS	Indicates the status of this entry: • other • invalid • learned • self • mgmt
MAC ADDRESS	Indicates the MAC address.
INTERFACE	Displays the network-to-network (NNI) interface.
SMLT REMOTE	Indicates the MAC address entry for the remote IST peer.
TUNNEL	Indicates the host name of the remote Backbone Edge Bridge (BEB).

For VSP 9000, the command **`show ip arp`**:

- Shows the tunnel endpoint if the endpoint is not a virtual MAC.

```
VSP-9012:1#show ip arp

================================================================================
                             IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN    PORT  TYPE    TTL(10 Sec) TUNNEL

--------------------------------------------------------------------------------
192.0.2.111    00:00:00:00:00:01  0       -     LOCAL   2160
198.51.100.1      00:00:00:00:00:02  2100   4/12   STATIC  2160

================================================================================
                           IP Arp Extn - GlobalRouter
================================================================================
MULTICAST-MAC-FLOODING    AGING(Minutes)       ARP-THRESHOLD
--------------------------------------------------------------------------------
disable                   360                  500

2 out of 14 ARP entries displayed
```

To see the virtual MAC, VSP 9000 adds the option **`spbm-tunnel-as-mac`** to the **`show ip arp`** command.

```
VSP-9012:1>show ip arp spbm-tunnel-as-mac
```

```
================================================================================
                             IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN    PORT  TYPE    TTL(10 Sec)  TUNNEL

--------------------------------------------------------------------------------
192.0.2.8       00:00:00:00:00:01  0       -     LOCAL   2160
198.51.100.8    58:16:26:b1:7a:01  616     -     LOCAL   2160
198.51.100.8    ff:ff:ff:ff:ff:ff  616     -     LOCAL   2160

================================================================================
                           IP Arp Extn - GlobalRouter
================================================================================
MULTICAST-MAC-FLOODING    AGING(Minutes)       ARP-THRESHOLD
--------------------------------------------------------------------------------
disable                   360                  500

3 out of 83 ARP entries displayed
```

The following table describes the fields in the output for the **`show ip arp`** command.

Parameter	Description
IP_ADDRESS	Indicates the IP address where ARP is configured.
MAC_ADDRESS	Indicates the MAC address where ARP is configured.
VLAN	Indicates the VLAN address where ARP is configured.
PORT	Indicates the port where ARP is configured.
TYPE	Indicates the type of learning (dynamic or local) where ARP is configured.
TTL (10 Sec)	Indicates the time-to-live as tenths of a second where ARP is configured.
TUNNEL	Displays the remote host name in the TUNNEL column for the SPBM ARP entry.
MULTICAST-MAC-FLOODING	Displays whether IP ARP multicast-MAC flooding is enabled or disabled. When enabled, the ARP entries for multicast MAC addresses are associated with the VLAN or port interface on which they were learned.
AGING (Minutes)	Displays when the ARP aging timer expires.
ARP-THRESHOLD	Displays the maximum number of outstanding ARP requests that a device can generate.

ERS 8800:

For Ethernet Routing Switch 8800, the command **`show vlan mac-address-entry`** shows the I-SID.

```
ERS-8800:5(config)#show vlan mac-address-entry
*******************************************************************************
Command Execution Time: WED APR 11 14:41:10 2012 UTC
*******************************************************************************

================================================================================
                                    Vlan Fdb
```

```
================================================================================
VLAN            MAC                                         QOS
 SMLT
ID   STATUS     ADDRESS            INTERFACE        MONITOR LEVEL
 REMOTE
--------------------------------------------------------------------------------
199  learned    00:0e:c0:c0:e2:00  MLT-199          false   1
 true
199  self       00:0e:c0:c0:f2:00  -                false   1
 false
1510 learned    00:04:75:dc:c3:4e  MLT-202          false   1
 false
1510 learned    00:04:75:f6:36:6d  MLT-202          false   1
 false
1510 learned    00:0c:f8:a9:22:0a  MLT-151          false   1
 true
1510 learned    00:15:9b:04:82:0f  MLT-151          false   1
 true
1510 learned    00:1b:4f:6a:70:40  MLT-202          false   1
 false
1510 learned    00:1c:eb:3c:e0:01  MLT-202          false   1
 false
1510 learned    00:1c:eb:ce:d4:01  MLT-201          false   1
 false

9 out of 9 entries in all fdb(s) displayed.
```

For Ethernet Routing Switch 8800, the command **show ip arp** output is the following:

```
ERS-8800:5(config)#show ip arp
*******************************************************************************
Command Execution Time: WED APR 11 14:41:31 2012 UTC
*******************************************************************************

================================================================================
                             IP Arp - GlobalRouter
================================================================================
IP_ADDRESS      MAC_ADDRESS        VLAN    PORT       TYPE    TTL(10 Sec)
--------------------------------------------------------------------------------
1.1.1.4         00:00:00:00:00:01  -       -         LOCAL   2160
10.30.199.1     00:0e:c0:c0:f2:00  199     -         LOCAL   2160
10.30.199.255   ff:ff:ff:ff:ff:ff  199     -         LOCAL   2160
10.30.199.2     00:0e:c0:c0:e2:00  199   MLT 199    DYNAMIC 2159

================================================================================
                           IP Arp Extn - GlobalRouter
================================================================================
   MULTICAST-MAC-FLOODING                AGING        ARP-THRESHOLD
--------------------------------------------------------------------------------
        disable                           360             500
4 out of 4 ARP entries displayed
```

Software scaling comparison

This following table compares the software scaling capabilities of the Virtual Services Platform 9000 and the Ethernet Routing Switch 8000 Series Release 7.2.

Table 39: Software scaling comparison

	ERS 8000 7.2	VSP 9000 4.0
	Maximum number supported	
Layer 2		
IEEE-or Port-based VLANs	4,000 for port-, protocol-, and IEEE 802.1Q based-VLANs combined	4,084
Protocol-based VLANs	4,000 for port-, protocol-, and IEEE 802.1Q based-VLANs combined	16
Internet Protocol (IP) subnet-based VLANs	800	256
Source MAC-based VLANs	4,000	100
Multiple Spanning Tree Protocol (MSTP)	32 instances	64 instances
Rapid Spanning Tree Protocol (RSTP)	1 instance	1 instance
MACs in forwarding database (FDB)	64,000 (32,000 with SMLT)	128,000
Multi-Link Trunking (MLT)	128 groups	512 groups
Split Multi-Link Trunking (SMLT)	127 groups	511 groups
Inter-Switch Trunk (IST)	1 group	1 group
S/MLT ports for each group	8	16
LACP	128 aggregators	512 aggregators
LACP ports for each aggregator	8 active and 8 standby	8 active and 8 standby
VLACP Interfaces	96	128
SLPP	200 VLANs	500 VLANs
Layer 3		
Internet Protocol version 4 (IPv4) Interfaces	1,972 (VLAN-and brouter-based)	4,343
IP interfaces (Brouter)	1,972 (VLAN-and brouter-based)	480
Circuitless IP interfaces	256	256
ARP for each port, VRF, or VLAN	64,000 for each system	64,000 for each system
Static Address Resolution Protocol (ARP) entries	2,048 for each VRF 10,000 for each system	2,048 for each VRF 10,000 for each system
Static routes (IPv4)	2,000 for each VRF 10,000 total across VRFs	2,000 for each VRF 10,000 total across VRFs
FIB IPv4 routes	250,000	400,000 in first generation mode.

Table continues…

	ERS 8000 7.2	VSP 9000 4.0
	Maximum number supported	
		1,000,000 in second generation mode for second generation I/O modules. **Note:** • If you want to scale to 1,000,000, you must install second generation modules in the VSP 9010 or the VSP 9012, and you must run the chassis in second generation mode. • Both first generation I/O modules and second generation I/O modules require a Premier license.
RIB IPv4 routes	3 * fastpath routes	3 * fastpath routes
ECMP routes	5,000	64,000
ECMP routes (fastpath)	8	8
IPv4 VRF instances	256	512
RIP instances	64 (one for each VRF)	64 (one for each VRF)
RIP interfaces	200	200
RIP routes	2,500 for each VRF 10,000 for each system	2,500 for each VRF 10,000 for each system
OSPF instances	64 (one for each VRF)	64 (one for each VRF)
OSPF interfaces	500 for each system	512 active, 2000 passive
OSPF adjacencies	80 for each VRF 200 for each system	512
OSPF areas	5 for each VRF 24 for each system	12 for each OSPF instance 80 for each system
OSPF LSA packet size	6,000 bytes	Jumbo packets
OSPF routes	20,000 for each VRF 50,000 for each system	64,000
BGP peers	250	256
BGP Internet peers (full)	3	3
BGP routes	250,000	1.5 million
IP Routing policies (IPv4)	500 for each VRF	500 for each VRF

Table continues…

	ERS 8000 7.2	**VSP 9000 4.0**
	Maximum number supported	
	5,000 for each system	5,000 for each system
IP Prefix List	500	500
IP Prefix entries	25,000	25,000
RSMLT interfaces	500 RSMLT enabled VLANs on 128 SMLT interfaces	4,000 over 512 SMLT interfaces **Note:** VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
Multicast IGMP interfaces	1,980	4,084
Multicast source and group (S, G)	2,000 with SMLT 4,000 without SMLT	6,000
PIM interfaces	200 active; 1,972 passive	512 active; 4084 passive
VRRP interfaces	255	255 for each VRF 512 for each system **Note:** VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
VRRP interfaces fast timers (200ms)	12	24
UDP/DHCP Forwarding entries	512	512 for each VRF 1,024 for each system
NLB clusters — Unicast	128 for each VLAN 2,000 for each system	128 for each VLAN 2,000 for each system
NLB clusters — Multicast, with multicast MAC flooding disabled	1 for each VLAN 2,000 for each system	1 for each VLAN 2,000 for each system
NLB clusters — Multicast, with multicast MAC flooding enabled	128 for each VLAN 2,000 for each system	128 for each VLAN 2,000 for each system
IPv4/IPv6 Telnet sessions	8	8 each, 16 total **Note:** VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
IPv4/IPv6 FTP sessions	4	4 each, 8 total

Table continues…

	ERS 8000 7.2	VSP 9000 4.0
	Maximum number supported	
		Note: VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
IPv4/IPv6 Rlogin sessions	8	8 each, 16 total **Note:** VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
IPv6		
IPv6 interfaces	250	4,087 (4,084 VLAN and 3 management [1/1, 2/1, virtual IP]) **Note:** VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
IPv6 tunnels	350	2,000 **Note:** VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
IPv6 static routes	2,000	10,000 **Note:** VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
OSPFv3 areas	5	64
OSPFv3 adjacencies	80	512
OSPFv3 routes	5,000	64,000
Filters and QoS		
Flow-based policers	4,000	16,000 **Note:** VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
Port shapers	64 queues for each port	480

Table continues…

	ERS 8000 7.2	**VSP 9000 4.0**
	Maximum number supported	
		Note: VSP 9000 does not support IPv6 in Release 4.0 and Release 4.0.1.
Access control lists (ACL) for each chassis	4,000	2,048
Access control entries (ACE) for each chassis	10,000	16,000
ACEs for each ACL	1,000	1,000 (a combination of Security and QoS ACEs)
Unique redirect next hop values for ACE Actions	2,000	2,000
Diagnostics		
Mirroring ports	150 non R mode 384 R mode	479
Remote Mirroring Termination (RMT) ports	16	32
Operations, Administration, and Maintenance		
IPFIX flows	384,000 flows for each chassis	96,000 for each interface module 960,000 for each chassis
Shortest Path Bridging MAC (SPBM)		
ARP entries	6,000 (with SPBM enabled)	64,000 (routed)
IP routes with SPBM enabled	8,000	100,000 (combination of OSPF and IS-IS)
Layer 2 VSNs	2,000	4,000

Glossary

4–byte AS

4-byte Autonomous System (AS) numbers is the solution to the soon depleting 2-byte AS numbers. It provides a theoretical 4,294,967,296 unique AS numbers in BGP. 4-byte AS numbers are backward compatible with 2-byte AS numbers.

access control entry (ACE)

One of the filter rules that comprise an access control list (ACL). An ACE statement defines pattern match criteria for a packet and the desired behavior for packets that carry the pattern. When the packets match an ACE rule, the specified action executes.

access control list (ACL)

An ordered list of filter rules referred to as access control entries. The ACEs provide specific actions, such as dropping packets within a specified IP range, or a specific Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port or port range. When an ingress or egress packet meets the match criteria specified in one or more ACEs within an ACL, the corresponding action executes.

Avaya command line interface (ACLI)

A textual user interface. When you use ACLI, you respond to a prompt by typing a command. After you enter the command, you receive a system response.

Backbone Core Bridge (BCB)

Backbone Core Bridges (BCBs) form the core of the SPBM network. The BCBs are SPBM nodes that do not terminate the VSN services. BCBs forward encapsulated VSN traffic based on the Backbone MAC Destination Address (B-MAC-DA). A BCB can access information to send that traffic to any Backbone Edge Bridges (BEBs) in the SPBM backbone.

Backbone Edge Bridge (BEB)

Backbone Edge Bridges (BEBs) are SPBM nodes where Virtual Services Networks (VSNs) terminate. BEBs handle the boundary between the core MAC-in-MAC Shortest Bath Bridging MAC (SPBM) domain and the edge customer 802.1Q domain. A BEB node performs 802.1ah MAC-in-MAC encapsulation and decapsulation for the Virtual Services Network (VSN).

Backbone MAC (B-MAC)

Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation encapsulates customer MAC addresses in Backbone MAC (B-MAC) addresses. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that SPBM uses for delivery from end to end. As the MAC header stays the same across the network, no need exists to swap a label or perform a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end. In Shortest Path Bridging MAC (SPBM), each node has a System ID, which is used in the topology announcement. This same System ID also serves as the switch Backbone MAC address (B-MAC), which is used as the source and destination MAC address in the SPBM network.

Backbone VLAN identifier (B-VID)

The Backbone VLAN identifier (B-VID) indicates the Shortest Path Bridging MAC (SPBM) B-VLAN associated with the SPBM instance.

Border Gateway Protocol (BGP)

An inter-domain routing protocol that provides loop-free inter-domain routing between Autonomous Systems (AS) or within an AS.

Control Processor (CP) module

The Control Processor module runs all high level protocols (BGP, OSPF) and distributes the results (routing updates) to the rest of the system. The CP manages and configures the IO and Switch Fabric modules, and maintains and monitors the health of the chassis.

Control Processor Unit High Availability (CPU-HA)

CPU-HA activates two CP modules simultaneously. The CP modules exchange topology data so, if a failure occurs, either CP module can take precedence in less than one second with the most recent topology data.

cooling module (9010CM)

The cooling module is a hot swappable fan tray used to cool the Control Processor, I/O, and Switch Fabric modules in the Virtual Services Platform 9010. Two cooling modules are installed horizontally in the front of the chassis.

Enterprise Device Manager (EDM)

A Web-based embedded management system to support single-element management. EDM provides complete configuration management functionality for the supported devices and is supplied to the customer as embedded software in the device.

equal cost multipath (ECMP)

Distributes routing traffic among multiple equal-cost routes.

forwarding database (FDB)

A database that maps a port for every MAC address. If a packet is sent to a specific MAC address, the switch refers to the forwarding database for the corresponding port number and sends the data packet through that port.

I/O cooling module (9012FC)

The I/O cooling module is a hot swappable fan tray used to cool the I/O and CP modules in the Virtual Services Platform 9012.

Internet Group Management Protocol (IGMP)

IGMP is a host membership protocol used to arbitrate membership in multicast services. IP multicast routers use IGMP to learn the existence of host group members on their directly attached subnets.

Internet Protocol Flow Information eXport (IPFIX)

An IETF standard that improves the Netflow V9 protocol. IPFIX monitors IP flows.

interswitch trunking (IST)

A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.

IP multicast over SPBM

With IP multicast over SPBM, Avaya introduces extensions to the SPBM IS-IS control plane to exchange IP multicast stream advertisement and membership information. These extensions, combined with the use of IGMP snooping and querier functions at the edge of the SPBM cloud, efficiently transport IP multicast data by using sub-trees of the VSN shortest path tree per IP multicast group.

Layer 2 Virtual Services Network

The Layer 2 Virtual Services Network (L2 VSN) feature provides IP connectivity over SPBM for VLANs. Backbone Edge Bridges (BEBs) handle Layer 2 virtualization. At the BEBs you map the end-user VLAN to a Service Instance Identifier (I-SID). BEBs that have the same I-SID configured can participate in the same Layer 2 Virtual Services Network (VSN).

link aggregation group (LAG)

A group that increases the link speed beyond the limits of one single cable or port, and increases the redundancy for higher availability.

multicast group ID (MGID)

The multicast group ID (MGID) is a hardware mechanism the switch uses to send data to several ports simultaneously. Instead of sending the data to a specific port number, the switch directs the data to an MGID. The switch maintains a table that maps MGIDs to their member ports. Both virtual LAN (VLAN) and IP multicast (IPMC) use MGIDs.

MultiLink Trunking (MLT)

A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.

Multiple Spanning Tree Protocol (MSTP)

Configures multiple instances of the Rapid Spanning Tree Protocol (RSTP) on the switch.

Open Shortest Path First (OSPF)

A link-state routing protocol used as an Interior Gateway Protocol (IGP).

out of band (OOB)

Network dedicated for management access to chassis.

Packet Capture Tool (PCAP)

A data packet capture tool that captures ingress and egress (on Ethernet modules only) packets on selected ports. You can analyze captured packets for troubleshooting purposes.

Point-to-Point Protocol (PPP)

Point-to-Point Protocol is a basic protocol at the data link layer that provides its own authentication protocols, with no authorization stage. PPP is often used to form a direct connection between two networking nodes.

quality of service (QoS)

QoS features reserve resources in a congested network, allowing you to configure a higher priority for certain devices. For example, you can configure a higher priority for IP deskphones, which need a fixed bit rate and split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.

Rapid Spanning Tree Protocol (RSTP)

Reduces the recovery time after a network breakdown. RSTP enhances switch-generated Topology Change Notification (TCN) packets to reduce network flooding.

remote mirroring

A mirroring port that encapsulates traffic into a Layer 2 header and transmits it to a remote mirror target (RMT) for decapsulation. The packet transmits over a Layer 2 network and preserves the original packet.

Routed Split MultiLink Trunking (RSMLT)

Provides full router redundancy and rapid failover in routed core SMLT networks and as RSMLT-edge in routed SMLT edge applications; eliminating routing protocol timer dependencies when network failures occur.

Routing Information Protocol (RIP)

A distance vector protocol in the IP suite, used by IP network-layer protocol, that enables routers in the same AS to exchange routing information by means of periodic updates. You often use RIP as a very simple interior gateway protocol (IGP) within small networks.

Shortest Path Bridging MAC (SPBM)

Shortest Path Bridging MAC (SPBM) uses the Intermediate-System-to-Intermediate-System (IS-IS) link-state routing protocol to provide a loop-free Ethernet topology that creates a shortest-path topology from every node to every other node in the network based on node MAC addresses. SPBM uses the 802.1ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header. SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based link-state protocol, which can provide virtualization services, both layer 2 and layer 3, using a pure Ethernet technology base.

Simple Loop Prevention Protocol (SLPP)

Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).

Simple Network Management Protocol (SNMP)

SNMP administratively monitors network performance through agents and management stations.

small form-factor pluggable (SFP)

A hot-swappable input and output enhancement component used with Avaya products to allow gigabit Ethernet ports to link with other gigabit Ethernet ports over various media types.

small form-factor pluggable plus (SFP +)

SFP+ transceivers are similar to SFPs in physical appearance but SFP+ transceivers provide Ethernet at 10 gigabits per second (Gbps).

SMLT aggregation switch

One of two IST peer switches that form a split link aggregation group. It connects to multiple wiring closet switches, edge switches, or customer premise equipment (CPE) devices.

spanning tree

A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.

Spanning Tree Group (STG)

A collection of ports in one spanning-tree instance.

Split MultiLink Trunking (SMLT)

An Avaya extension to IEEE 802.1AX (link aggregation), provides nodal and link failure protection and flexible bandwidth scaling to improve on the level of Layer 2 resiliency.

Switch Fabric (SF) cooling module (9012RC)

The SF cooling module is a hot swappable fan tray used to cool the Switch Fabric (SF) modules in the Virtual Services Platform 9012.

Terminal Access Controller Access Control System plus

Terminal Access Controller Access Control System plus (TACACS+) is a security protocol that provides centralized validation of users who attempt to gain access to a router or network access server. TACACS+ uses Transmission Control Protocol (TCP) for its transport to ensure reliable delivery and encrypts the entire body of the packet. TACACS+ provides separate authentication, authorization, and accounting services. TACACS+ is not compatible with previous versions of TACACS.

tunnel

An end-to-end unidirectional tunnel between MPLS-enabled routers.

Virtual Enterprise Network Architecture (VENA)

Virtual Enterprise Network Architecture (VENA) is a virtualization architecture for next generation Enterprise data networks. VENA provides the data infrastructure for the private cloud by leveraging an open and interoperable IEEE technology, called Shortest Path Bridging (SPB) standard (802.1aq). VENA allows for the aggregation of multiple independent virtual servers to exist on a physical server and decouples the physical infrastructure from the connectivity services making the network adaptive and dynamic with simple one-touch provisioning.

Virtual Local Area Network (VLAN)

A Virtual Local Area Network is a group of hosts that communicate as if they are attached to the same broadcast domain regardless of their physical location. VLANs are layer 2 constructs.

Virtual Private Network (VPN)

A Virtual Private Network (VPN) requires remote users to be authenticated and ensures private information is not accessible to unauthorized parties. A VPN can allow users to access network resources or to share data.

virtual router forwarding (VRF)

Provides traffic isolation between customers operating over the same node. Each virtual router emulates the behavior of a dedicated hardware router by providing separate routing functionality, and the network treats each VRF as a separate physical router.

Virtual Router Redundancy Protocol (VRRP)

A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.

Weighted Random Early Detection (WRED)

A mechanism that provides congestion avoidance capabilities. The basic operating philosophy of WRED is that it detects the onset of congestion and starts dropping packets in random fashion before queue overflow leads to tail drops.