



Platform Migration Avaya Virtual Services Platform 9000

3.2
NN46250-107, 02.02
March 2012

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 generally makes available to users of its products. 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 Web sites 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 its Hardware and Software ("Product(s)"). 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 Web site: <http://support.avaya.com>. Please note that if you acquired the Product(s) from an authorized Avaya reseller outside of the United States and Canada, the warranty is provided to you by said Avaya reseller and not by Avaya.

Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/](http://support.avaya.com/licenseinfo/) ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. 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 AUTHORIZED RESELLER; 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").

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, or Hardware provided by Avaya. All content on this site, the documentation 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

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: <http://support.avaya.com/Copyright>.

Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the Documentation 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 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, and "Linux" is a registered trademark of Linus Torvalds.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support Web site: <http://support.avaya.com>.

Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your Product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: <http://support.avaya.com>.

Contents

Chapter 1: Purpose of this document.....	5
Navigation.....	5
Chapter 2: New in this release.....	7
Features.....	7
Other changes.....	8
Chapter 3: Hardware considerations.....	9
Quick reference.....	9
Power requirements.....	10
Lift the chassis.....	11
Chassis measurements.....	11
Module requirements.....	12
Removing external storage devices from the CP module.....	12
Protecting modules.....	14
Optical transceivers.....	15
Chapter 4: Platform functionality and system handling.....	19
Quick reference.....	19
Upgrading the software.....	20
Committing an upgrade.....	22
Resetting the platform.....	22
Restarting the platform.....	23
Configuring system flags.....	24
High Availability mode.....	29
Out-of-band configuration.....	32
Key Health Indicators.....	33
Local alarms.....	33
Rapid Failure Detection and Recovery.....	33
Chapter 5: Software considerations.....	35
Key software differences.....	35
Spanning tree.....	36
VLANs.....	36
Policy-based VLANs.....	37
IP routing and VLANs.....	39
VLAN implementation.....	39
VLAN configuration rules.....	40
VLAN MAC security.....	40
Link aggregation and loop prevention.....	41
QoS and traffic filters.....	44
Queuing.....	46
Internal QoS level.....	46
Ingress mappings.....	46
CPU protection.....	48
Traffic management profiles.....	48
ACTs and ACLs.....	49
Access control entries.....	50

Remote mirroring.....	58
IP routing.....	59
IPv6 routing.....	59
IP multicast.....	60
Software scaling comparison.....	60
Chapter 6: Customer service.....	65
Getting technical documentation.....	65
Getting product training.....	65
Getting help from a distributor or reseller.....	65
Getting technical support from the Avaya Web site.....	65

Chapter 1: Purpose of this document

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. You must read this document to understand the differences between the two products before you begin to migrate to the Virtual Services Platform 9000.

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
<i>Hardware</i>	
VSP 9000 uses different power supplies.	Power requirements on page 10
The VSP 9012 chassis is larger and weighs more than the ERS 8000 series.	Lift the chassis on page 11 Chassis measurements on page 11
VSP 9000 has specific module installation and handling requirements.	Module requirements on page 12 Protecting modules on page 14
VSP 9000 supports different SFP transceivers than ERS 8000 series. VSP 9000 uses SFP+ transceivers rather than XFP.	Optical transceivers on page 15
<i>Platform functionality and system handling</i>	
VSP 9000 introduces ACLI differences, software patching, new features to monitor health, alarms, and failure detection and recovery.	Quick reference on page 19
<i>Software</i>	
Specific protocols and features are not supported on both products.	Key software differences on page 35
Support for spanning tree protocols is different.	Spanning tree on page 36

Purpose of this document

Difference	Document location
VLAN classification and port membership in VLANs is different.	VLANs on page 36
SMLT and SLT configuration is different.	Link aggregation and loop prevention on page 41
VSP 9000 uses ACLs to assign MAC or VLAN QoS levels, and does not use egress queue sets.	QoS and traffic filters on page 44
Attributes and operators for traffic filtering are different.	QoS and traffic filters on page 44
VSP 9000 supports Layer 3 remote mirroring.	Remote mirroring on page 58
VSP 9000 supports mixed-AS peer communication for BGP.	IP routing on page 59
IPv6 routing support is different.	IPv6 routing on page 59
VSP 9000 provides full IGMPv3 support.	IP multicast on page 60

Chapter 2: New in this release

The following sections describe what is new in *Avaya Virtual Services Platform 9000 Platform Migration*, NN46250–107 for Release 3.2 .

Features

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

BGP 4–byte AS

Release 3.2 adds support for 4–byte autonomous systems. For more information about the support differences from Ethernet Routing Switch 8000 series, see [IP routing](#) on page 59.

External flash memory card

Beginning with Release 3.2, you can hot swap the external flash memory card on a 9080CP module. For more information about the 9080CP module, see [Module requirements](#) on page 12.

IGMPv3

Release 3.2 adds support for full IGMPv3. For more information, see [IP multicast](#) on page 60.

IPv6

Release 3.2 adds IPv6 routing support. For more information about support differences from Ethernet Routing Switch 8000 series, see [IPv6 routing](#) on page 59 and [Software scaling comparison](#) on page 60.

Lossless Ethernet

Release 3.2 adds support for Lossless Ethernet, which the Ethernet Routing Switch 8000 series does not currently support. For configuration information, see *Avaya Virtual Services Platform 9000 Configuration — Ethernet Modules*, NN46250–508. For conceptual information, see *Avaya Virtual Services Platform 9000 Planning and Engineering — Network Design*, NN46250–200.

VSP Talk

Release 3.2 introduces the VSP Talk application. Use VSP Talk to monitor the status of a VSP 9000 device remotely through an instant messaging (IM) client. Ethernet Routing Switch 8000 series does not support a similar application. For more information, see *Avaya Virtual Services Platform 9000 Administration*, NN46250–600.

Other changes

See the following section for information about change that is not feature-related:

Introduction chapter

Introduction chapter is renamed to Purpose of the document that states the purpose of referring to that document.

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	VSP 9000 uses AC power supplies.
–	VSP 9000 uses a new chassis and new modules.
ERS 8000 uses combined SF/CPU modules	VSP 9000 uses separate SF and CP modules with specific requirements for each.
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 chassis is both front-to-back and side-to-side. Avaya recommends 36 inches (in.) (91 centimeters [cm]) of free space in both the front and the back of the machine, and also 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

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
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.

Power requirements

Virtual Services Platform 9000 supports up to six 1200–2000 Watt AC power supplies. Virtual Services Platform 9000 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.

You can operate the 9006AC power supplies separately, or in parallel, or parallel redundant configurations. You can use the 9006AC 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

N supplies connect to one power phase; N supplies connect to the other power phase. A loss of a phase affects only half of the power supplies.

! Important:


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

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 www.avaya.com/support.

Lift the chassis

About this task

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



Use the recessed handles at the top and bottom of the chassis sides to lift the chassis. From the rear: Lift the chassis from the bottom only.

! 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.

Chassis measurements

The 9012 chassis is a different size than the Ethernet Routing Switch 8000 series chassis, particularly the depth of the chassis. The 9012 chassis is deeper than the Ethernet Routing Switch 8000 series chassis. 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	VSP 9012	ERS 8010co	ERS 8010	ERS 8006	ERS 8803-R
Height	24.375 in. (61.91 cm)	35 in. (88.9 cm)	22.9 in. (58.2 cm)	15.8 in.(40.1 cm)	12.25 in. (31.1 cm)
Width	17.5 in. (44.45 cm)	19 in (48.26 cm)	17.5 in. (44.5 cm)	17.5 in. (44.5 cm)	19 in. (48.3 cm)
Depth	32.5 in. (82.55 cm)	23.7 in. (60.19 cm)	19.9 in. (50.5 cm)	19.9 in. (50.5 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 12.

9090SF module

Beginning with Release 3.1, you must install a minimum of three 9090SF modules in the chassis. Install an SF module in slots SF1 and SF4; install a third SF module in one of the remaining slots. Releases prior to 3.1 require a minimum of four SF modules.

Take care when you insert the SF modules as it is possible to misalign them, and potentially damage the connectors in the chassis.

+ Tip:

It is easier to install the SF modules 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 USB and 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 VSP 9000 stop command will 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 USB or 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.

Procedure

1. 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.
2. 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

```
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 VSP 9000 automatically recognizes them. The devices are accessible within seconds after insertion.

After you insert the external Compact Flash, you should enable logging to the external Compact Flash by using 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:

- Do not touch the top of the module or you can damage pins, components, and connectors. Doing so can result in equipment damage.
- Damage to a module can occur if it is bumped into another object, including other modules installed in a chassis. Use both hands to support modules.
- Always place the modules on appropriate antistatic material.
- Support the module from underneath with two hands. Do not touch the top of the module. Do not touch the pins or electrical connections.

- 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.
- Be careful not to bump module connectors against the action levers of an adjacent module. Damage to connectors can result.

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).

9024XL interface module

The 9024XL interface module is a 24 port 10 gigabit per second (Gb/s) small form-factor pluggable plus (SFP+) interface module. The module supports a maximum throughput of 105 Mpps over 24 ports of 10 Gb/s Ethernet traffic using standard SFP+ fiber connectors. The module supports SR, LR, LRM, and ER SFP+ format.

The following table details the multimode fiber (MMF) and single-mode fiber (SMF) SFP and SFP+ fiber connectors supported by the 9024XL module.

Table 4: Supported SFP and SFP+ fiber connectors for the 9024XL module

Model number	Product number	Description
10GBASE-SR/SW	AA1403015-E6	850 nanometers (nm). The range is up to the following: <ul style="list-style-type: none"> • 22 m using 62.5 micrometer (μm), 160 megaHertz times km (MHz-km) MMF • 33 m using 62.5 μm, 200 MHz-km MMF • 66 m using 62.5 μm, 500 MHz-km MMF • 82 m using 50 μm, 500 MHz-km MMF • 300 m using 50 μm, 2000 MHz-km MMF

Model number	Product number	Description
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-LR/LW	AA1403011-E6	1310 nm SMF. The range is up to 10 km.
10GBASE-ER/EW	AA1403013-E6	1550 nm SMF. The range is up to 40 km.
10GBASE-CX	AA1403018-E6 to AA1403021-E6	4-pair twinaxial copper cable to connect 10 Gb ports. The maximum range is 15 m.
1000BASE-SX	AA1419048-E6	Well-suited for campus local area networks (LAN) and intrabuilding links. Up to 275 or 550 m reach (fiber-dependent) over a fiber pair.
1000BASE-LX	AA1419049-E6	The range is up to 10 km reach over a single mode fiber (SMF) pair. The range is up to 550 m reach over a multimode fiber (MMF) pair.
1000BASE-XD	AA1419050-E6	1310 nm,. The range is up to 40 km over SMF pair.
	AA1419051-E6	1550 nm (non-CWDM). The range is up to 40 km over SMF pair.
1000BASE-ZX	AA1419052-E6	1550 nm (non-CWDM). The range is up to 70 km over SMF pair.
1000BASE CWDM	AA1419053-E6 to AA1419060-E6	1470 nm to 1610 nm (CWDM). The range is up to 40km over SMF pair.

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

9048GB interface module

The 9048GB interface module is a 48 port 1 Gb/s 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 connectors supported by the 9048GB module.

Table 5: Supported SFP connectors for the 9048GB module

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	AA1419048-E6	850 nm, up to 275 or 550 m
1000BASE-LX	AA1419049-E6	1310 nm, up to 10 km
1000BASE-XD	AA1419050-E6	1310 nm, up to 40 km
	AA1419051-E6	1550 nm, up to 40km (non-CWDM)
1000BASE-ZX	AA1419052-E6	1550 nm, up to 70 km (non-CWDM)
1000BASE-BX-U	AA1419069-E6	1310 nm, up to 10km
	AA1419076-E6	1310 nm, up to 40km
1000BASE-BX-D	AA1419070-E6	1490 nm, up to 10km
	AA1419077-E6	1490 nm, up to 40km
1000BASE-EX	AA1419071-E6	1550 nm, up to 120 km (non-CWDM)
1000BASE 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
	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

Hardware considerations

The 9048GB is 100/1000M capable.

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

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 6: 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

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 <i>Avaya Virtual Services Platform 9000 Troubleshooting</i> , NN46520–700.
-	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.
- You must log on to at least the Privileged EXEC mode in ACLI.

Procedure

1. Extract the release distribution files to the /intflash/release/ directory:

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

2. Install the image:

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

3. 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 auto-commit is not enabled, the system restarts with the last known working version after the commit timer has expired. This feature ensures you can regain control of the system if an upgrade fails.

4. Confirm the software is upgraded:

```
show software
```

Example

```
VSP-9012:1#software add VSP9K.3.2.0.0.tgz
```

```
VSP-9012:1#show software
```

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

```
VSP-9012:1#software activate 3.2.0.0.GA
```

Variable definitions

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

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

Variable	Value
remove <i>WORD</i> <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.

Before you begin

- You must log on to at least the Privileged EXEC mode in ACLI.

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

```
Commit an upgrade:  
software commit
```

Resetting the platform

Before you begin

- You must log on to Privileged EXEC mode in ACLI.

About this task

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

Procedure

Reset the switch:

```
reset [-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 7: Variable definitions

Variable	Value
-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

Before you begin

- You must log on to the Privileged EXEC mode in ACLI.

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. 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.

2. (Optional) Switch the primary CPU function to the backup CP module:

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

Variable definitions

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

Table 8: Variable definitions

Variable	Value
config <i>WORD</i> <1–99>	<p>Specifies the software configuration device and file name in one of the following formats:</p> <ul style="list-style-type: none"> • a.b.c.d:<file> • /intflash/ <file> • /extflash/ <file> • /usb/<file> <p>The file name, including the directory structure, can include up to 99 characters.</p>
-y	<p>Suppresses the confirmation message before the switch restarts. If you omit this parameter, you must confirm the action before the system restarts.</p>

Configuring system flags

Before you begin

- If you enable the `hsecure` flag, you cannot enable the flags for the Web server or SSH password-authentication.
- You must log on to Global Configuration mode in ACLI.

! Important:

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

About this task

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

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. Enable system flags:

```
boot config flags <block-snmp|debug-config [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|wdt>
```

2. 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|wdt>
```

3. Configure the system flag to the default value:

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

4. Save the changed configuration.

5. Restart the switch.

Example

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

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

Activate High Secure mode:

```
VSP-9012:1(config)#boot config flags hsecure
```

```
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 9: 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 [file]	Activates or disables run-time debugging of the configuration file. If you activate debugging, line-by-line configuration file processing appears on the console during CPU initialization. The default value is disabled. File logs the debug-config output to /extflash/debugconfig.txt. If you change the debug-config variable value, you must restart the switch.
debugmode	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. ! Important: Do not change this parameter unless directed by Avaya.
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 <ul style="list-style-type: none"> • 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

Variable	Value
	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: <ul style="list-style-type: none"> • 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 Secure mode, the switch prompts a password change if you enter invalid-length passwords.
logging	The logging command is used to activate or disable system logging. The default value is enabled. The system names log files according to the following: <ul style="list-style-type: none"> • File names appear in 8.3 (log.xxxxxxx.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

Variable	Value
	if the system reaches the maximum log file size.
reboot	<p>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.</p> <p>! Important: Do not change this parameter unless directed by Avaya.</p>
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	<p>Activates or disables the Telnet server service. The default is disabled.</p> <p>If you disable the Telnet server service in a dual CPU system, the Telnet server prevents a Telnet connection initiated from the other CPU.</p>
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	<p>Activates or disables the creation of trace logs. The default value is disabled.</p> <p>! Important: Do not change this parameter unless directed by Avaya.</p>

Variable	Value
verify-config	Activates syntax checking of the configuration file. The default value is enabled. If the system finds a syntax error, it loads the factory default configuration. If you disable this flag, the system logs syntax errors and the CPU continues to source the configuration file. Avaya recommends that you disable the verify-config flag. If you change this parameter, you must restart the switch.
wdt	Activates or disables the hardware watchdog timer monitoring a hardware circuit. The default value is activated. The watchdog timer restarts the switch based on software errors. If you change the wdt flag, you must restart the switch. ! Important: Do not change this parameter unless directed by Avaya.

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 are synchronized. In non HA-CPU mode, also called Warm Standby, the two CPUs are not synchronized.

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

Table 10: Feature support for HA-CPU

Feature	Release 3.2
Modules	Yes
Platform	Yes
Layer 2	Yes
Layer 3	Yes; partial-HA for Border Gateway Protocol and IPv6

Feature	Release 3.2
Multicast	Partial HA
Security	Yes
Applications	Partial HA

Table 11: Synchronization capabilities in HA-CPU mode

Synchronization of	Release 3.2
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
SMLT parameters	Yes
VLAN parameters	Yes
Layer 3	
ARP entries	Yes
Border Gateway Protocol (BGP)	Partial (configuration only)
Dynamic Host Configuration Protocol (DHCP) Relay	Partial (configuration only)
Internet Group Management Protocol	Partial (configuration only)
Internet Group Management (IGMP) Snooping	Yes
IP 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

Synchronization of	Release 3.2
Routed Split Multi-Link Trunking (RSMLT)	Yes
RSMLT edge support	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
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 *Avaya Virtual Services Platform 9000 Administration*, NN46250-600.

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
- Internet Group Management Protocol (IGMP)
- IPv6
- Protocol Independent Multicast-Sparse Mode (PIM-SM)
- Protocol Independent Multicast-Source Specific Mode (PIM-SSM)
- Border Gateway Protocol (BGP)
- 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 is not intended to provide a comprehensive debugging solution. Instead, KHI identifies key information that could lead Avaya support personnel towards discovery of a specific failure. After the technician assesses the KHI information, further debugging is required 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 *Avaya Virtual Services Platform 9000 Fault Management*, NN46250–703.

Local alarms

The Avaya Virtual Services Platform 9000 contains a local alarms mechanism. Local alarms are raised and cleared by applications running on the switch. Active alarms are viewed using the `show alarm database` command in the CLI. Local alarms are an automatic mechanism run by the system that do not require any 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 *Avaya Virtual Services Platform 9000 Troubleshooting*, 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 is updated every 3.3 ms.

All packets that are forwarded 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 7.1. The following table identifies key support differences for software features and protocols.

Feature	Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Distance Vector Multicast Routing Protocol (DVMRP)	Yes	No
IGMPv3 (full)	No — partial only	Yes
IP Virtual Private Network (IPVPN)	Yes	No
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
Shortest Path Bridging (SPB)	Yes	No
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 12: 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.

For more information about RSTP and MSTP, see *Avaya Virtual Services Platform 9000 Configuration — VLANs and Spanning Tree*, 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 13: 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 packet is dropped.

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 is changed 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 that has 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 14: 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	

Protocol	Protocol ID	Encapsulation	Number of records
	0x6003	Ether2	
	0x6005	Ether2	
	0x6006	Ether2	
	0x6007	Ether2	
	0x6008	Ether2	
	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 it 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 frame is dropped.

Because it 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 can support up to 4084 configurable VLANs. VLAN IDs range from 1 to 4084. VLAN IDs 4085 to 4094 are reserved for internal use.
- You can configure only one protocol-based VLAN for a given protocol. Up to 16 protocol-based VLANs are supported.
- The VLAN membership of a frame is determined by the following order of precedence, if applicable:
 - a. IEEE 802.1Q tagged VLAN ID
 - b. source MAC-based VLAN
 - c. IP subnet-based VLAN
 - d. protocol-based VLAN
 - e. 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 *Avaya Virtual Services Platform 9000 Configuration — VLANs and Spanning Tree*, 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 15: Link aggregation and loop prevention quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
<i>Loop prevention</i>	
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.
<i>SMLT, MLT, and IST</i>	
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 *Avaya Virtual Services Platform 9000 Configuration — VLANs and Spanning Tree*, 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.

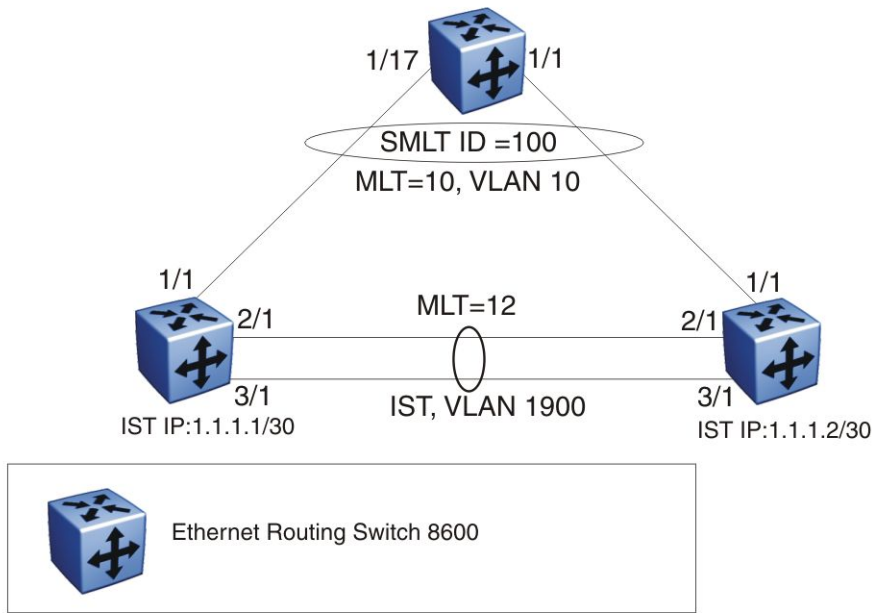


Figure 1: Ethernet Routing Switch 8600 Network topology

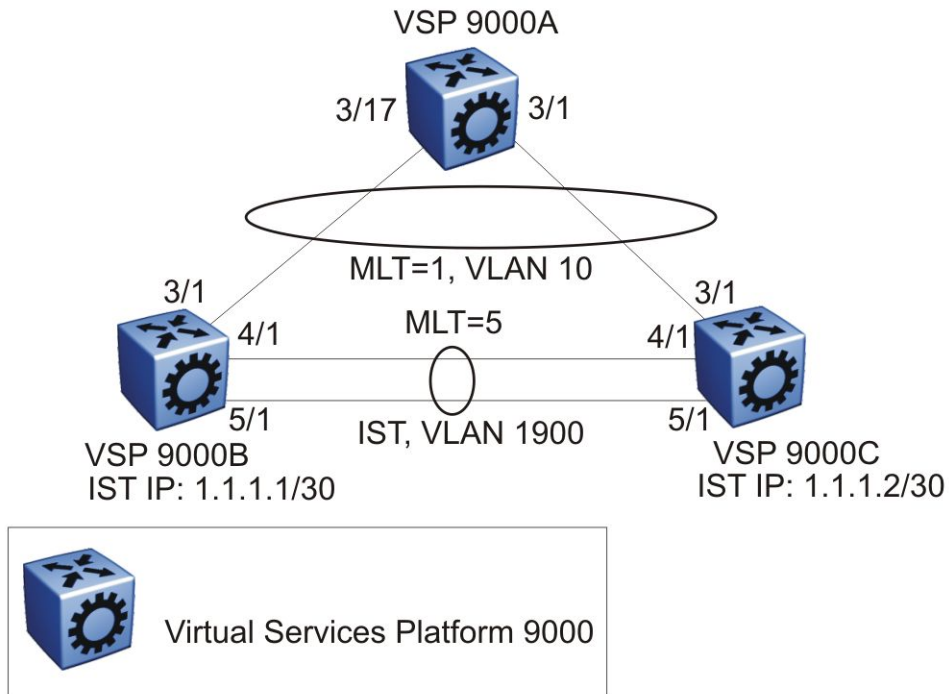


Figure 2: Virtual Services Platform 9000 Network topology

For information about MLT and SMLT, see *Avaya Virtual Services Platform 9000 Configuration – Link Aggregation, MLT, and SMLT*, NN46250–503. For information about Routed SMLT, see *Avaya Virtual Services Platform 9000 Configuration — IP Routing*, 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 16: 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 will 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.

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
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.
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 *Avaya Virtual Services Platform 9000 Configuration — QoS and ACL-Based Filtering*, NN46250–502.

Queuing

Weighted Random Early Detection (WRED) is supported 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 Behaviour (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 is derived 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 17: 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

DSCP	Layer 2 trusted	Layer 3 trusted and DiffServ enabled	IP packet	Routed packet	Ingress tagged	Internal 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
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 *Avaya Virtual Services Platform 9000 Administration*, 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 is dropped and what traffic is switched or routed by the system 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 60% of line rate, the system limits multicast traffic to 40%.

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

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 18: 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 19: Default action comparison

	Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Drop mode	Permit/Deny	Permit/Deny

	Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
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 will return a match (hit). If the rule does not match, the search continues and at the end of the search a miss is returned.

- Mask

ACL-based filters provide the mask operator to match on Layer 2, Layer 3, and Layer 4 packet fields. The mask operator is used 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 it 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 0x00:00:00:FF:FF:FF.

The following table explains the mask operator for MAC addresses.

Table 20: Mask operator for MAC address

Rule	Result
<pre>filter acl ace ethernet 10 10 dst- mac mask 0x01:00:5e:00:00:01 24</pre> <p>which is the same as</p> <pre>filter acl ace ethernet 10 10 dst- mac mask 01:00:5e:00:00:01 0x000000FFFFFF</pre>	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.
<pre>filter acl ace ethernet 10 10 dst- mac mask 0x01:00:5e:00:00:01 0xFFFFFFFF0000</pre>	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.
<pre>filter acl ace ethernet 10 10 dst- mac mask 0x01:00:5e:00:00:01 0xFF00FF0000FF</pre>	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 21: Mask operator for IP address

Rule	Result
<pre>filter acl ace ip 10 10 src-ip mask 2.10.10.12 24</pre> <p>which is the same as</p> <pre>filter acl ace ip 10 10 src-ip mask 2.10.10.12 0.255.255.255</pre>	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 .
<pre>filter acl ace ip 10 10 src-ip mask 3.4.5.6 255.255.255.0</pre>	The rule matches only the least significant 8 bits, for example, 6, and does not care about

Rule	Result
	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 interpret 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 22: Mask operator for Layer 4 source port

Rule	Result
<code>filter acl ace protocol 10 10 src-port mask 80 0xF</code>	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 23: Mask operator configuration examples

Filter configuration	Address examples that match the filter	Address examples that do not match the filter
<pre>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</pre>	<pre>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</pre>	<pre>Source MAC: 00:00:11:11:16:01 00:ff:11:11:16:f1</pre>
<pre>filter acl ace 1000 1000 filter acl ace ethernet 1000 1000 dst-</pre>	<pre>Destination MAC: 00:00:00:64:16:01</pre>	<pre>Destination MAC: 00:00:00:24:16:20</pre>

Filter configuration	Address examples that match the filter	Address examples that do not match the filter
<pre>mac mask 00:00:00:64:16:00 0x00000060001f filter acl ace action 1000 1000 deny log count filter acl ace 1000 1000 enable</pre>	<pre>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</pre>	<pre>00:00:00:64:16:20 00:00:00:63:16:01 00:00:00:65:16:01</pre>
<p>IP mask (dotted decimal notation):</p> <pre>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</pre>	<p>Source IP:</p> <pre>192.168.4.1 192.168.4.10 192.168.4.30 192.168.4.31</pre>	<p>Source IP:</p> <pre>192.168.3.1 192.168.4.32</pre>
<pre>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</pre>	<p>Destination IP:</p> <pre>192.168.7.1 192.168.7.3</pre>	<p>Destination IP:</p> <pre>192.168.7.4 192.168.7.5</pre>
<p>IP mask (decimal notation):</p> <pre>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</pre>	<p>Source IP:</p> <pre>192.168.4.1 192.168.4.10 192.168.4.30 192.168.4.31</pre>	<p>Source IP:</p> <pre>192.168.3.1 192.168.4.32</pre>
<pre>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</pre>	<p>Destination IP:</p> <pre>192.168.7.1 192.168.7.3</pre>	<p>Destination IP:</p> <pre>192.168.7.4 192.168.7.5</pre>
<p>Protocol mask:</p> <pre>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</pre>	<p>TCP source port</p> <pre>256 TCP source port 356 TCP source port 511</pre>	<p>TCP source port</p> <pre>255 TCP source port 512</pre>

Filter configuration	Address examples that match the filter	Address examples that do not match the filter
<pre>filter acl ace action 901 1 deny count filter acl ace 901 1 enable</pre> <p>This mask implies packets with TCP source port 256–511 match the filter, while 0–255 and > 511 miss the filter.</p>		

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 24: 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
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 25: 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
filter 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 26: 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-	Source port	Equal-to, Mask

ERS 8000 Attribute	ERS 8000 operator	VSP 9000 attribute	VSP 9000 operator
	than or equal-to, not-equal-to		
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 27: 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.
<ul style="list-style-type: none"> • 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.
Count	None	Collect ACE statistics. This action applies to ingress and egress ACLs.

Security ACE Actions	User supplied parameters	Comments
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 28: QoS ACE Actions

QOS ACE Actions	User supplied parameters	Comments
Remark	<ul style="list-style-type: none"> • 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.

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 29: Mirroring quick reference

Ethernet Routing Switch 8000 Series	Virtual Services Platform 9000
Supports Layer 2 remote mirroring	Support 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 *Avaya Virtual Services Platform 9000 Troubleshooting*, 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 30: 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 *Avaya Virtual Services Platform 9000 Configuration — BGP Services*, 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 31: 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 *Avaya Virtual Services Platform 9000 Configuration — IP Multicast Protocols*, NN46250–504.

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 7.1.

Table 32: Software scaling comparison

	ERS 8000 7.1	VSP 9000 3.2
	Maximum number supported	
<i>Layer 2</i>		
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)	128K
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
<i>Layer 3</i>		
Internet Protocol version 4 (IPv4) Interfaces	1,972 (VLAN-and router-based)	4,343
IP interfaces (Router)	1,972 (VLAN-and router-based)	480
Circuitless IP interfaces	256	256
ARP for each port, VRF, or VLAN	64,000 for each system	64,000 for each system

	ERS 8000 7.1	VSP 9000 3.2
	Maximum number supported	
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	500,000
RIB IPv4 routes	3 * fastpath routes	3 * fastpath routes
ECMP routes	5,000	64,000
ECMP routes (fastpath)	8	8
IPv4 VRF instances	255	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 5,000 for each system	500 for each VRF 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
Multicast IGMP interfaces	1,980	4,084
Multicast source and group (S, G)	2,000 with SMLT 4,000 without SMLT	6,000

	ERS 8000 7.1	VSP 9000 3.2
	Maximum number supported	
PIM interfaces	200 active; 1,972 passive	512 active; 4084 passive
VRRP interfaces	255	255 for each VRF 512 for each system
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
IPv4/IPv6 FTP sessions	4	4 each, 8 total
IPv4/IPv6 Rlogin sessions	8	8 each, 16 total
<i>IPv6</i>		
IPv6 interfaces	250	4,087 (4,084 VLAN and 3 management [1/1, 2/1, virtual IP])
IPv6 tunnels	350	2,000
IPv6 static routes	2,000	10,000
OSPFv3 areas	5	64
OSPFv3 adjacencies	80	512
OSPFv3 routes	5,000	64,000
<i>Filters and QoS</i>		
Flow-based policers	4,000	16,000
Port shapers	64 queues for each port	480
Access control lists (ACL) for each chassis	4,000	2,048
Access control entries (ACE) for each chassis	10,000	16,000

	ERS 8000 7.1	VSP 9000 3.2
	Maximum number supported	
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
<i>Diagnostics</i>		
Mirroring ports	150 non R mode 384 R mode	479
Remote Mirroring Termination (RMT) ports	16	32
<i>Operations, Administration, and Maintenance</i>		
IPFIX flows	384,000 flows for each chassis	96,000 for each interface module 960,000 for each chassis

Chapter 6: Customer service

Visit the Avaya Web site to access the complete range of services and support that Avaya provides. Go to www.avaya.com or go to one of the pages listed in the following sections.

Getting technical documentation

To download and print selected technical publications and release notes directly from the Internet, go to www.avaya.com/support.

Getting product training

Ongoing product training is available. For more information or to register, you can access the Web site at www.avaya.com/support. From this Web site, you can locate the Training contacts link on the left-hand navigation pane.

Getting help from a distributor or reseller

If you purchased a service contract for your Avaya product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

Getting technical support from the Avaya Web site

The easiest and most effective way to get technical support for Avaya products is from the Avaya Technical Support Web site at www.avaya.com/support.

