Release Notes — Release 3.3
Avaya Virtual Services Platform 9000

Release 3.3
NN46250-401
Issue 04.05
June 2013

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

Warranty

Avaya provides a limited warranty on 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 website: http://support.avaya.com. Please note that if you acquired the Product(s) from an authorized Avaya Channel Partner outside of the United States and Canada, the warranty is provided to you by said Avaya Channel Partner and not by Avaya. “Software” means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products or pre-installed on hardware products, and any upgrades, updates, bug fixes, or modified versions.

Licenses

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

License types

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

Heritage Nortel Software

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

Copyright

Except where expressly stated otherwise, no use should be made of materials on this site, the Documentation, Software, 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

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

Preventing Toll Fraud

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

Avaya Toll Fraud intervention

If you suspect that you are being victimized by Toll Fraud and you need technical assistance or support, call Technical Service Center Toll Fraud Intervention Hotline at +1-800-643-2353 for the United States and Canada. For additional support telephone numbers, see the Avaya Support website: http://support.avaya.com. Suspected security

vulnerabilities with Avaya products should be reported to Avaya by sending mail to: securityalerts@avaya.com.

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. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Downloading Documentation

For the most current versions of Documentation, see the Avaya Support website: http://support.avaya.com.

Contact Avaya Support

See the Avaya Support website: http://support.avaya.com for product notices and articles, or to report a problem with your Avaya product. For a list of support telephone numbers and contact addresses, go to the Avaya Support website: http://support.avaya.com, scroll to the bottom of the page, and select Contact Avaya Support.

Contents

Chapter 1: Purpose of this document

This document describes new features and important information about the latest release. Release notes include a list of known issues (including workarounds where appropriate) and a list of resolved issues. This document also describes known limitations and expected behaviors that may first appear to be issues.

Chapter 2: New in this release

The following sections describe what's new in *Avaya Virtual Services Platform 9000 Release Notes* (NN46250–401) for Release 3.3.

Features

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

Shortest Path Bridging MAC

This release adds support for Shortest Path Bridging MAC (SPBM). For more information, see New features on page 11, SPBM restrictions and limitations on page 28, Supported standards on page 41, and Software scaling capabilities on page 36.

SSL certificates

Beginning with this release, you can upload your own certificates to Virtual Services Platform 9000. For more information, see User configurable SSL certificates on page 27.

OM4

The maximum distance for 10GBASE-SR/SW is updated to reflect support for OM4.

Other changes

See the following sections for information about changes that are not feature-related.

Known issues

Known issues on page 53 is updated.

mib.txt and mib_vsp.txt

MIB errors on page 68 identifies errors in the MIB files.

Other documents

Other documents on page 32 is added to reference Technical Configuration and Technical Solutions guides for Virtual Services Platform 9000.

Resolved issues

Resolved issues on page 69 is updated for issues fixed in this release.

Cooling modules

The model names for cooling modules are updated in software and documentation. For more information, see:

- Supported hardware on page 13
- Hardware and software compatibility on page 30

Chapter 3: Important notices and new features

This section describes the supported hardware and software features of the Avaya Virtual Services Platform and provides important information for this release.

New features

The following sections highlight the feature support added in this release.

Connectivity Fault Management

Use Connectivity Fault Management (CFM) to debug connectivity issues and isolate faults in a Shortest Path Bridging MAC (SPBM) network. CFM operates at Layer 2 and provides an equivalent of the **`ping`** and **`traceroute`** commands. To support troubleshooting of the SPBM cloud, this release supports a subset of CFM functionality. CFM is based on the IEEE 802.1ag standard.

For more information about CFM, see *Avaya Virtual Services Platform 9000 Configuration — Shortest Path Bridging MAC (SPBM)*, NN46250–510.

Shortest Path Bridging MAC

Release 3.3 supports the IEEE 802.1aq standard of Shortest Path Bridging MACinMAC (SPBM). SPBM makes network virtualization much easier to deploy within the enterprise environment, reducing the complexity of the network while at the same time providing greater scalability.

SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core control plane to a single protocol that can provide virtualization services for both Layer 2 and Layer 3, on a common Ethernet infrastructure using a pure Ethernet technology base. SPBM separates the Ethernet network into edge and core domains with complete isolation between their MAC addresses. This technology provides all the features and benefits required by carrier-grade, enterprise, and service provider deployments without the complexity of alternative technologies, for example, Multiprotocol Label Switching (MPLS). SPBM integrates into a single control plane all the functions that MPLS requires multiple layers and protocols to support.

SPBM provides any-to-any connectivity in a network in an optimized, loop-free manner. SPBM employs shortest-path trees to each destination, without the long convergence delays experienced with Spanning Tree Protocol (STP). To do this, SPBM uses Intermediate System-to-Intermediate System (IS-IS) link state routing protocol to learn and distribute network information. IS-IS dynamically learns the topology of a network and uses its inherent knowledge to construct shortest path unicast and multicast trees from every node to every

other node in the network. Also, unlike STP, IS-IS does not block ports to provide a loop free topology, so bandwidth is not wasted.

The SPBM components introduced in this release are:

- Shortest Path Bridging (IEEE 802.1aq) for simple and safe VLAN extensions across a network. SPBM does not use spanning tree, and all its links are active.
- SPBM support of dual-homing non-SPBM switches to an IST pair.
- SPBM IP for simple and safe VRF extensions across a network infrastructure without Open Shortest Path First (OSPF) or Border Gateway Protocol (BGP).
- Inter-VSN routing for routing of Layer 2 VPNs in the SPB domain.

You must purchase and install the Premier License to use SPBM. For more information about SPBM, see *Avaya Virtual Services Platform 9000 Configuration — Shortest Path Bridging MAC (SPBM)*, NN46250–510.

Important:

Different from the Ethernet Routing Switch 8800 SPBM implementation, on Virtual Services Platform 9000 when you map a customer VLAN (C-VLAN) to a Service Instance Identifier (I-SID), you must include the IST MLT peer switches in the VLAN, which means all IST member ports must be members of this C-VLAN. In ERS 8800, you cannot configure IST MLT to be part of the C-VLAN. For more information about this and other differences between the two products, see *Avaya Virtual Services Platform 9000 Platform Migration*, NN46250–107.

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

Table 1: SPBM quick reference

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports IP VPN Lite over SPBM.	Does not support IP VPN Lite over SPBM.
Reserves 519 multicast group IDs (MGIDs) for SPBM operation.	Reserves 100 mulitcast group IDs (MGIDs) for SPBM operation.
Does not require you to configure C-VLANs on the IST MLT.	Requires the inclusion of IST MLT peer switches in the C-VLAN.
Traffic can pass between single-homed VLANs attached to IST peers if the IST is down.	Traffic cannot pass between single-homed VLANs attached to IST peers if the IST is down.
Decapsulates MAC-in-MAC traffic at the primary BEB or secondary BEB irrespective of whether the traffic is from the primary B-VLAN or secondary B-VLAN.	Decapsulates MAC-in-MAC traffic at the primary BEB from the primary B-VLAN and traffic at the secondary BEB from the secondary B-VLAN. Requires the IST to be up to pass traffic between both IST switches for single-homed VLANs.

Ethernet Routing Switch 8000 series	Virtual Services Platform 9000
Supports NNI functionality on all R and RS module ports.	Supports SPBM NNI on the 9024XL module.
Disables tagging if you delete IS-IS from an interface.	Keeps tagging enabled, if you delete IS-IS from an interface.
The **show isis spbm multicast-fib** command for both the primary B-VLAN on the secondary IST switch, and the secondary B-VLAN on the primary IST switch, lists the following: • All active single-homed UNIs • An SMLT if all ports with the SMLT are down on the partner (or the partner is down), and at least one port within the SMLT is up locally • An SLT if the SLT port is down on the partner (or the partner is down) but up locally	The **show isis spbm multicast-fib** command for both the primary B-VLAN on the secondary IST switch, and the secondary B-VLAN on the primary IST switch, lists no UNI ports if the partner is up, but all UNI ports if the partner is down.
Does not have the **spbm-tunnel-as-mac** option for **show vlan mac-address-entry** command and the **show ip arp** command.	Adds the option **spbm-tunnel-as-mac** to the **show vlan mac-address-entry** command and the **show ip arp** command.
Ethernet Routing Switch 8800 supports CFM enhancements that make CFM much simpler to configure and behave more like native ping and traceroute.	Virtual Services Platform 9000 will support these CFM enhancements in a future release.

Compatibility with Ethernet Routing Switch 8800 Release 7.1.3:

Release 3.3 fully supports the IEEE 802.1aq standard. Ethernet Routing Switch 8800 Releases 7.1.0.0 and 7.1.1.0 support a pre-standard version of 802.1aq so cannot interoperate with the Virtual Services Platform 9000.

Virtual Services Platform 9000 can interoperate with Ethernet Routing Switch 8800 Release 7.1.3 if you configure the **spbm version** on the Ethernet Routing Switch 8800 to `802.1aq`. For more information, see *Avaya Ethernet Routing Switch 8800/8600 Configuration — Shortest Path Bridging MAC (SPBM)*, NN46205–525.

Supported hardware

This section describes the hardware components of Avaya Virtual Services Platform 9000, and important notices or restrictions.

New hardware

No new hardware is added for this release.

9012 Chassis

The 9012 chassis has 12 slots and 8 bays in the front and 8 slots and 2 bays in the back. Two front slots are for the Control Processor (CP) module and ten front slots are for the interface modules. Six front bays are for the power supplies and two front bays are for interface cooling modules. Six back slots are for Switch Fabric (SF) modules and two back slots are for future development. The two back bays are for SF cooling modules.

The following table details the weight and dimensions of the 9012 chassis.

Table 2: Dimensions and weight

Width	17.5 in. (44.45 cm)
Height	24.375 in. (61.91 cm)
Depth	32.5 in. (82.55 cm) plus cable management system
Weight (chassis and midplane)	160 lb (73 kg)
Weight (chassis, midplane, and cooling modules)	183 lb (83 kg)
Weight (chassis, midplane, cooling modules, and filler panels)	288 lb (131 kg)

The 9012 chassis also has the following features:

- Airflow is both side to side and front to back.
- Every module and power supply is hot swappable.
- The chassis has side handles on the bottom and top, and the front and back panels have handles for lifting the 9012 chassis.

9006AC power supply

The 9006AC power supply accepts 100 to 120 volts alternating current (V AC) and 200 to 240 V AC input voltage using 20 amps (A) plugs and drawing approximately twelve A. With 100 to 120 V AC nominal input voltage conditions, the power supply produces a maximum of 1200 Watts of 48 volts direct current (V DC) power. With 200 to 240 V AC nominal input voltage conditions, the power supply can output 2000 Watts of 48 V DC power.

9080CP Control Processor module

The 9080CP CP module runs all high level protocols, and distributes the results (routing updates) to the rest of the system, manages and configures the interface and SF modules, and maintains and monitors the health of the chassis.

The 9080CP module contains two 8542E Control Processor Units (CPU) running at 1.33 gigahertz (GHz). Each processor comes with two double-data-rate two (DDR2) dual in-line memory module (DIMM) of memory, for a maximum of 4 GB Random Access Memory (RAM) for each processor. The 9080CP module measures 23 inches in length.

The 9080CP module supports the following interfaces:

- console port, DB9
- ethernet management, RJ45
- Universal Serial Bus (USB) type A (Master)
- external Compact Flash

The external Compact Flash card is mandatory. The following table lists the external storage devices you can order to use with the CP module.

Table 3: Supported external storage devices

Model number	Description
EC1411011-E6	2GB USB memory flash drive
EC1411010-E6	2GB Compact Flash memory card

Note:

Avaya recommends that you use the Avaya Compact Flash device (EC1411010-E6) with the VSP 9000 because it has been validated for proper operation on the VSP 9000. Avaya recommends that you do not use other Compact Flash devices because they have not been verified for VSP 9000 compatibility, and can result in loss of access to the Compact Flash device.

You can hot swap the external storage devices but you must follow a specific procedure to avoid data loss or hardware damage. To properly remove an external storage device, see Removing external storage devices from the CP module on page 20.

The 9080CP module has light-emitting-diodes (LED) duplicating the LEDs of the modules in the back of the chassis. The LEDs that show on the CP module are

- six for the power supplies
- six for the Switch Fabrics
- two for the auxiliary modules
- four for the fans
- one for the central processor 1 on the CP module
- one for the central processor 2 on the CP module
- one for the AUX processor on the CP module, for future use
- one for master
- one for system temperature

You can see the following intermittent error on the CP module after a reboot or an HA failover:

```
ERROR: [bcmScoreboard.0]SB_FLIB(V0): [bcmScoreboard.0]QM_ERROR2,
FB_TAIL_CACHE_OVERFLOW Set
```

You can ignore this message. The message has no functional impact on the system.

Redundancy:

The 9080CP module architecture provides redundancy if you use two CP modules in a system.

Caution:

Risk of file system corruption

To remove a master CP module from the chassis, you must follow the Avaya recommended procedure, Removing a master CP module with CPU-HA mode activated on page 20. Failure to follow this procedure can result in file system corruption.

9090SF Switch Fabric module

The 9090SF Switch Fabric (SF) modules provide the back end switching solution in the midplane chassis. Each 9090SF module connects to ten different interface modules and two CP modules simultaneously. Each chassis has slots for five operational SF modules plus one hot backup. The 9090SF modules measure 8.5 inches in length.

While the 9090SF modules are physically the same, the functions they perform vary based on the slots in which you install them. The SF slots 1 and 4 are for system operations and are called bandwidth managers.

The following table details the functions of the SF modules based on their slot location.

Table 4: Switch Fabric module functions based on slot location

Slots	Function
SF1 and SF4	bandwidth manager
SF2, SF3, SF5, and SF6	X-bar switching function

Note:

You must install a minimum of three SF modules in the chassis. Install an SF module in both slots SF1 and SF4. Install a third SF module in one of the remaining slots.

If you install a second bandwidth manager after a bandwidth manager failover, there is a small window where data can be lost. Perform this action during a maintenance window.

The 9024XL has a throughput capability of 105 Mpps. The 9048 modules have a throughput capability of 70 Mpps.

9012RC Switch Fabric cooling module

The 9012RC SF cooling modules contain two fans each to cool the SF modules. You install them in the back of the chassis.

The output for the **`show sys-info card`** command can incorrectly identify the module description as 9012SC. This description was programmed in the SEEPROM during manufacturing of early 9012RC cooling modules. These modules will continue to report this information. Newly manufactured modules correctly display 9012RC as the module description.

9012FC IO cooling module

The 9012FC IO cooling modules contain eight fans each to cool the interface modules. You install them from the front of the chassis.

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 5: 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: • 26 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, 400 MHz-km MMF • 82 m using 50 µm, 500 MHz-km MMF • 300 m using 50 µm, 2000 MHz-km MMF • 400 m using 50 µm, 4700 MHz-km MMF (OM4)
10GBASE-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.

Model number	Product number	Description
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 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.

9048GT interface module

The 9048GT interface module is a 48 port 10/100/1000M Ethernet Copper interface module with RJ45 connectors.

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

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 6: 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.

Model	ROHS product number	Description
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

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.

Removing a master CP module with CPU-HA mode activated

Perform this procedure, if the system operates in CPU-HA mode, to properly remove the master CP module. You must perform this procedure to avoid jeopardizing the integrity of the file system.

Procedure

1. Log on to Global Configuration mode:
   ```
   enable
   configure terminal
   ```
2. Use the `sys action cpu-switch-over` command to fail over to another CP.
3. Use the slot power commands to power down the module.
4. Remove the CP module.
   This action removes the original master.

   Important:

   Do not reinsert a CP module until at least 15 seconds have elapsed. This is long enough for another CP module to become master.

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#sys action cpu-switch-over
```

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.

About this task

Note:

Avaya recommends that you use the Avaya Compact Flash device (EC1411010-E6) with the VSP 9000 because it has been validated for proper operation on the VSP 9000. Avaya recommends that you do not use other Compact Flash devices because they have not been verified for VSP 9000 compatibility, and can result in loss of access to the Compact Flash device.

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

File names for this release

This section describes the Avaya Virtual Services Platform 9000 software files.

Software files

The following table provides the details of the Virtual Services Platform 9000 software files. File sizes are approximate.

Table 7: Software files

File name	Description	Size
VSP9K.3.3.0.0.tgz	Release 3.3 archived distribution	104 MB
VSP9K.3.3.0.0_modules .tgz	Encryption modules	39 KB

Open Source software files

The following table gives the details of the Open Source software files distributed with the Virtual Services Platform 9000 software.

Table 8: Open Source software files

File name	Description	Size
VSP9K.3.3.0.0_oss-notice.html	Master copyright file. This file is located in the Licenses directory.	412 KB
VSP9K.3.3.0.0_OpenSource.zip	Open source base software for Virtual Services Platform 9000 Release 3.3.	96 MB

You can download Avaya Virtual Services Platform 9000 software and files, including MIB files, from the Avaya Support Portal at www.avaya.com/support. Click **Downloads**.

Important information and restrictions

This section contains important information and restrictions you must consider before you use the Avaya Virtual Services Platform 9000.

Protecting modules

Caution:

Risk of equipment damage. Do not touch the top of the module or you can damage pins, components and connectors.

Caution:

Modules are heavy. Damage to a module can occur if it bumps 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 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 ejector levers of an adjacent module. Damage to connectors can result.

Module installation precautions

You must take the following precautions while you install modules in the Virtual Services Platform 9000:

- Ensure the module sheet metal slides in the rails on the side of chassis.
- Modules come with screws embedded in the sheet metal. You must use the screws to keep the cards tightly in place.
- You must support the weight of the modules until they are inserted completely.

Resetting multiple modules

When you reset multiple modules in the system, it is important to make sure the module has fully recovered before you reset the next module. If the subsequent module is reset before the previous module has recovered, various error messages can appear as the system recovers through the system synchronization.

Supported browsers

Virtual Services Platform 9000 supports the following browsers to access the Enterprise Device Manager (EDM):

- Microsoft Internet Explorer 8.0
- Mozilla Firefox 7.x

Environmental specifications

The following table lists the minimum and maximum environmental specifications for Virtual Services Platform 9000 operation.

Table 9: Minimum and maximum operational environmental specifications

Operating environmental specification	Min	Max
Temperature	0°C	60°C
Relative humidity	10%	90%
Altitude	0 ft	10 000 ft
Thermal Shock	-40C	85C
Vibration, peak to peak displacement	—	0.005 in. (5 to 32 Hz)
Audible noise	—	60 DB

The following table lists the minimum and maximum environmental specifications for storage of the Virtual Services Platform 9000.

Table 10: Minimum and maximum storage environmental specifications

Storage environmental parameter	Min	Max
Temperature	-25°C	70°C
Relative Humidity	—	90%
Altitude	—	10 000 ft

The following table lists the operational requirements of an unpackaged Virtual Services Platform 9000.

Table 11: Unpackaged operational requirements

Unpackaged operational requirements	Standard specification	Note
Operational vibration (Sinusoidal)	European Telecommunications Standard (ETS) 300 019-1-3 and International Electrotechnical Commission (IEC) 68-2-6 test Fc	Shock of low significance, such as slamming door
Shock 30g 11ms	International Electrotechnical Commission (IEC) 68-2-27	Shock of low significance

The following table lists the requirements of an unpackaged Virtual Services Platform 9000 for storage and transport.

Table 12: Unpackaged nonoperational requirements

Unpackaged nonoperational Requirements (Storage and transport)	Standard specification	Note
Transportation Vibration (Sinusoidal)	Bellcore GR-63-Core issue 1 oct 1995	All kinds of trucks and trailers, trains, ships; well-developed roads
Transportation bounce (4 inch drop onto normal rest face, 2 edges and 2 corners)	IEC 68-2-31 or Bellcore GR-63-CORE, issue 1 oct 1995	Transportation handling
Package drop (Package weight less than 20 kg. Drop onto 3 faces, 3 edges and 3 corners from a heigth of 75 cm.)	Bellcore GR-63-CORE issue 1 oct 1995	Transportation handling

Reliability

The following table lists the mean time between failures of the various modules of the Virtual Services Platform 9000.

Table 13: Reliability

Component	Mean time between failures
Chassis	500,000 hours
Fan tray module	500,000 hours
Power supply module	1,500,000 hours
Switch Fabric module	1,447,178 hours
Control Processor module	777,001 hours
Interface modules	• 9024XL — 209,908 hours • 9048GB — 267,953 hours • 9048GT — 278,629 hours
Hard Drive and Auxiliary modules	200,000 hours

IPv6 addressing

Virtual Services Platform 9000 supports IPv6 addressing with the following access methods and features:

- DHCP Relay
- DNS client
- FTP client and server
- HTTP and HTTPS
- ping
- Rlogin
- RADIUS client
- SNMP
- SSH
- Syslog client
- Telnet
- TFTP client and server

User configurable SSL certificates

Virtual Services Platform 9000 does not generate SSL certificates with user-configurable parameters. You can, however, use your own certificate.

You can generate a certificate off the VSP 9000 system, and upload the key and certificate files to the `/intflash/ssh` directory. Rename the uploaded files to host.cert and host.key, and then reboot the system. The system loads the user-generated certificates during startup. If the system cannot find host.cert and host.key during startup, it generates a default certificate.

For more information about SSH and SSL certificates, see *Avaya Virtual Services Platform 9000 Administration*, NN46250–600.

EDM image management

EDM does not currently support image management functionality. You must perform all image management work through the ACLI. This includes, but is not limited to, software upgrades, software image management, and software patching. See *Avaya Virtual Services Platform*

9000 Upgrades and Patches , NN46250–400, for information and procedures about image management.

After you use ACLI to upgrade or downgrade the system software, before you connect to the device using EDM, Avaya recommends that you clear the browser cache. If you fail to clear the browser cache before you connect to the device, you can continue to see the previous software version in EDM.

Feature licensing

After you start a new system, the 60–day Premium Trial license countdown begins. You will see notification messages as the countdown approaches the end of the trial period. After 60 days, the Premium Trial license expires. You will see messages on the console and in the alarms database that the license has expired. The next time you restart the system after the license expiration, the system no longer supports Advanced or Premier services.

If you use a Base license, you do not need to install a license file. If you purchase an Advanced or Premier license, you must obtain and install a license file. For more information about how to generate and install a license file, see *Avaya Virtual Services Platform 9000 Administration*, NN46250–600.

SPBM restrictions and limitations

This section describes the restrictions and limitations associated with SPBM on Avaya Virtual Services Platform 9000.

RSTP and MSTP

The following list identifies restrictions and limitations associated with RSTP and MSTP:

- RSTP mode does not support SPBM.
- A C-VLAN-level loop across SPBM NNI ports cannot be detected and needs to be resolved at the provisional level.
- SPBM NNI ports are not part of the Layer 2 VSN C-VLAN, and BPDUs are not transmitted over the SPBM tunnel. SPBM can only guarantee loop-free topologies consisting of the NNI ports. Avaya recommends that you always use Simple Loop Prevention Protocol (SLPP) in an SMLT environment.

Note:

Avaya recommends that you deploy SLPP on C-VLANs to detect loops created by customers in their access networks. However, SLPP is not required on B-VLANs, and it is not supported. The B-VLAN active topology is controlled by IS-IS that has loop mitigation and prevention capabilities built into the protocol.

- SPBM uses STG 63 and MSTI 62 for internal use. STG 63 or MSTI 62 cannot be used by other VLANs or MSTIs.
- You must configure SPBM B-VLANs on all devices in the same MSTP region. MSTP requires this configuration to generate the correct digest.

SPBM IS-IS

The following list identifies restrictions and limitations associated with SPBM IS-IS:

- The current release does not support IP over IS-IS as defined by RFC 1195. IS-IS protocol is only to facilitate SPBM.
- The current release uses level 1 IS-IS. The current release does not support level 2 IS-IS. The ACLI command **`show isis int-l2-contl-pkts`** is not supported in the current release because the IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1.
- The IS-IS standard defines wide (32bit) metrics and narrow (8 bits) metrics. The current release supports the wide metric.
- SPBM supports full High Availability (HA). The SPBM and IS-IS configuration and dynamic information (such as adjacencies and LSPs) are all synchronized to the standby CPU to ensure seamless switchover. Because the HA framework cannot guarantee seamless switchover, there is a 6 to 7 seconds gap between the master CPU going down and the backup CPU coming up. To avoid IS-IS adjacencies bouncing during the switchover, the recommended hello interval is 9 seconds and the hello multiple is 3.

Pay special attention to the expected scaling of routes in the network when you select configuration values for the **`isis l1-hello-interval`** and **`isis l1-hello-multiplier`** commands on IS-IS interfaces. The default values for these commands work well for most networks, including those using moderately-scaled routes. In highly-scaled networks, you may need to configure higher values for these commands. For example, if the total number of non IS-IS routes on a given BEB exceeds 25,000 in combination with approximately 60,000 IS-IS routes that the BEB receives from other BEBs in the network, Avaya recommends that you configure a value of 12 for **`isis l1-hello-multiplier`**, rather than use the default value of 3.

SPBM NNI SMLT

For NNI-facing SMLT, the current release supports only one link between each IS-IS adjacency pair.

Only the 9024XL interface module supports NNI ports.

Multicast

The current release does not support enabling PIM or IGMP snooping on a C-VLAN.

Virtual Services Platform 9000 does not support NLB-multicast and NLB-multicast with IGMP for SPBM.

VLACP

VLACP is generally used when a repeater or switch exists between connected Virtual Services Platform 9000 switches to detect when a connection is down even when the link LED is lit. If you configure VLACP on an SPBM link that is also an IST link, during a connection fail over (where the link LED stays lit) the IS-IS hellos time out first (after 27 seconds, using default values) and take down the IS-IS adjacency. IS-IS then calculates the new shortest path and

fails over the SPBM traffic. 90 seconds after the connection failure (using default values), VLACP goes down but the IST link was already taken down by IS-IS.

In this scenario, there is no data traffic impact because IS-IS can find another path in the SPBM network before VLACP goes down.

SNMP traps

On each SPBM peer, if you configure the SPBM B-VLANs to use different VLAN IDs, for example, VLAN 10 and 20 on one switch, and VLAN 30 and 40 on the second, the system does not generate a trap message to alert of the mismatch because the two switches cannot receive control packets from one another. Configure the SPBM B-VLANs to use matching VLAN IDs.

Other

The following list identifies other restrictions and limitations:

- The current release does not support I-SID filters.
- You cannot enable C-VLAN and B-VLAN on the same port, except the IST ports.
- The current release supports NNI only on 10 GbE interface modules.

Legacy IS-IS

An SPBM node can form an adjacency with a legacy IS-IS router. This adjacency means you can introduce SPBM into existing networks and provide easy migration.

Fixes from previous releases

The VSP 9000 Software Release 3.3 incorporates all fixes from prior releases, up to and including, Release 3.2.0.0.

Hardware and software compatibility

The following tables describe the hardware and the minimum Avaya Virtual Services Platform 9000 software version required to support the hardware.

Table 14: Hardware and minimum software version

Chassis, switching fabrics, and control processors			Minimum software version	Part number
	9012VSP chassis	12-slot chassis	3.0	EC1402001-E6
	9090SF	Switch Fabric module	3.0	EC1404006-E6

Chassis, switching fabrics, and control processors			Minimum software version	Part number
	9080CP	Control Processor module	3.0	EC1404007-E6
Power Supplies				
	9006AC	1200–2000W AC Power Supply	3.0	EC1405A01-E6
Cooling modules				
	9012FC	Side Fan Tray	3.0	EC1411001-E6
	9012RC	Fabric Fan Tray	3.0	EC1411002-E6
Ethernet modules				
	9024XL	24-port 10GBASE-X SFP +/SFP	3.0	EC1404001-E6
	9048GB	48-port 1000BASE-X SFP	3.0	EC1404002-E6
	9048GT	48-port 10/100/1000BASE-T	3.0	EC1404003-E6
Compatible SFPs and SFP+s For more information about SFP and SFP+, see *Avaya Virtual Services Platform 9000 Installation — SFP Hardware Components, NN46250-305*				
	100BASE-FX SFP	1310 nm, 100 Mb/s Ethernet, multimode fiber, duplex LC connector	3.0	AA1419074-E6
	1000BASE-T SFP	Gigabit Ethernet, RJ-45 connector	3.0	AA1419043-E6
	1000BASE-SX DDI SFP	850 nm, Gigabit Ethernet, duplex LC connector	3.0	AA1419048-E6
	1000BASE-LX DDI SFP	1310 nm, Gigabit Ethernet, duplex LC connector	3.0	AA1419049-E6
	1000BASE-XD DDI SFP	1310 nm, Gigabit Ethernet, duplex LC connector	3.0	AA1419050-E6
		1550 nm, Gigabit Ethernet, duplex LC connector		AA1419051-E6
	1000BASE-ZX DDI SFP	1550 nm, Gigabit Ethernet, duplex LC connector	3.0	AA1419052-E6
	1000BASE-BX DDI SFP	1310 nm (tx) and 1490 nm (rx), 1490 nm (tx) 1310 nm (rx), Gigabit Ethernet, single-fiber LC connector,	3.0	AA1419069-E6 (10 km at 1310 nm) AA1419076-E6 (40 km at 1310 nm) AA1419070-E6 (10 km at 1490 nm)

Chassis, switching fabrics, and control processors			Minimum software version	Part number
				AA1419077-E6 (40 km at 1490 nm)
	1000BASE-EX DDI SFP	1550 nm, Gigabit Ethernet, duplex LC connector	3.0	AA1419071-E6
	1000BASE DDI CWDM 40 km SFP	Gigabit Ethernet, duplex LC connector	3.0	AA1419053-E6 to AA1419060-E6.
	1000BASE DDI CWDM 70 km SFP	Gigabit Ethernet, duplex LC connector	3.0	AA1419061-E6 to AA1419068-E6.
	10GBASE-SR/SW SFP+	400m, 850nm MMF	3.0	AA1403015-E6
	10GBASE-LRM SFP+	220 m, 1260 to 1355 nm; 1310 nm nominal MMF,	3.0	AA1403017-E6
	10GBASE-LR/LW SFP+	10km, 1310nm SMF	3.0	AA1403011-E6
	10GBASE-ER/EW SFP+	40km, 1550nm SMF	3.0	AA1403013-E6
	10GBASE-CX	4-pair twinaxial copper cable that plugs into the SFP+ socket and connects two 10 Gb ports.	3.0	AA1403018-E6 to AA1403021-E6

Other documents

In addition to the product documentation, Avaya provides Technical Configuration Guides and Technical Solution Guides. You can refer to these guides for more information about how to configure or use the Virtual Services Platform 9000 in specific scenarios. The following table lists the guides available for the Virtual Services Platform 9000.

Document title	Document number
Switch Clustering using Split-MultiLink Trunking (SMLT) with VSP 9000, ERS 8600/8800, 8300, and 5000 Technical Configuration Guide	NN48500-518
Switch Clustering Supported Topologies and Interoperability with Virtual Services Platform 9000 & Ethernet Routing Switches	NN48500-555

Document title	Document number
Super Large Campus Technical Configuration Guide	NN48500-609
Avaya Virtual Services Platform 9000 with Coraid EtherDrive SRX-Series Storage Appliances Technical Configuration Guide	NN48500-611
Avaya Virtual Services Platform 9000 and Avaya Virtual Services Platform 7000 with Coraid EtherDrive SRX-Series Storage Appliances Technical Configuration Guide	NN48500-629
Avaya Flare™ for Avaya Data Technical Configuration Guide	NN48500-613
IPv6 for VSP 9000 Technical Configuration Guide	NN48500–634

You can find these documents at www.avaya.com/support under the product Data Networking Solution, or by performing a search.

Chapter 4: Software and hardware scaling capabilities

This chapter details the software and hardware scaling capabilities of Avaya Virtual Services Platform 9000. The information in *Avaya Virtual Services Platform 9000 Release Notes*, NN46250–401, takes precedence over information in other documents.

Hardware scaling capabilities

This section lists hardware scaling capabilities of Avaya Virtual Services Platform 9000.

Table 15: Hardware scaling capabilities

	Maximum number supported
9024XL I/O module	
10GbE fiber connections	240 (10 x 24)
Processor	1 GHz
9048GB I/O module	
GbE fiber connections	480 (10 x 48)
Processor	1 GHz
9048GT I/O module	
10/100/1000 copper connections	480 (10 x 48)
Processor	1 GHz
9080CP CP module	
Processor	1.33 GHz
Console port	1 D-subminiature 25-pin shell 9 pin connector (DB9) per CP module
Ethernet management	1 Registered Jack (RJ) 45 per CP module
USB port	1 Universal Serial Bus (USB) Type A (Master) per CP module
External Compact Flash	1 per CP module
9012 Chassis	

	Maximum number supported
Control Processor (CP) modules	2
Interface modules	10
Switch Fabric (SF) modules	6 You must install a minimum of 3 SF modules in the chassis.
Auxiliary slots	2
Power supplies	6
Total power capacity	• 10 kW in 220 V AC mode • 6 kW in 110 V AC mode
Jumbo packets	9600 bytes for IPv4 9500 bytes for IPv6

Software scaling capabilities

This section lists software scaling capabilities of Avaya Virtual Services Platform 9000.

Table 16: Software scaling capabilities

	Maximum number supported
Layer 2	
IEEE/Port-based VLANs	4,084
Inter-Switch Trunk (IST)	1 group
Internet Protocol (IP) Subnet-based VLANs	256
LACP	512 aggregators
LACP ports per aggregator	8 active and 8 standby
Lossless Ethernet	2 ports for each 8–port cluster 6 ports for each 9024XL module
MACs in forwarding database (FDB)	128K
Multi-Link Trunking (MLT)	512 groups
Multiple Spanning Tree Protocol (MSTP)	64 instances
Protocol-based VLANs	16
Rapid Spanning Tree Protocol (RSTP)	1 instance
SLPP	500 VLANs

	Maximum number supported
Source MAC-based VLANs	100
Split Multi-Link Trunking (SMLT)	511 groups
SMLT ports per group	16
VLACP Interfaces	128
Layer 3	
Address Resolution Protocol (ARP) for each port, VRF, or VLAN	64,000 entries total
BGP Internet peers (full)	3
BGP peers	256
BGP routes	1.5 million
BGP+ routes	128,000
Circuitless IP interfaces	256
ECMP routes	64,000
ECMP routes (fastpath)	8
FIB IPv4 routes	500,000
FIB IPv6 routes	128,000
The fastpath forwarding table uses a common table for IPv4 and IPv6 forwarding records. IPv6 records are approximately four times the size of IPv4 records. The maximum number of 500,000 IPv4 routes is possible when no IPv6 routes are configured, and the maximum number of IPv6 routes is 128,000 when no IPv4 routes are configured.	
IPv4 interfaces	4,343
IP interfaces (Brouter)	480
IP prefix entries	25 000
IPv4 prefix list	500
IP routing policies	500 for each VRF 5,000 for each system
IPFIX flows	96,000 for each interface module 960,000 for each chassis
IPv4 or IPv6 FTP sessions	4 each, 8 total
IPv4 or IPv6 Rlogin sessions	8 each, 16 total
IPv4 or IPv6 SSH sessions	8 total (any combination of IPv4 and IPv6 up to 8)
IPv4 or IPv6 Telnet sessions	8 each, 16 total
IPv4 VRF instances	512

	Maximum number supported
IPv6 dynamic neighbors/interface	64K
IPv6 interfaces	4,087 (4,084 VLAN and 3 management [1/1, 2/1, virtual IP])
IPv6 routes (fastpath)	128,000
IPv6 static neighbors	1,000
IPv6 static routes	10,000
IPv6 tunnels	2,000
Multicast IGMP interfaces	4,084
Multicast source and group (S, G)	6,000
NLB Clusters — Multicast, with multicast MAC flooding disabled	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
NLB Clusters — Unicast	128 for each VLAN 2,000 for each system
OSPF adjacencies	512
OSPF areas	12 for each OSPF instance 80 for each system
OSPF instances	64 (one per VRF)
OSPF interfaces	512 active, 2000 passive
OSPF LSA packet size	Jumbo packets
OSPF routes	64,000
OSPFv3 adjacencies	512
OSPFv3 adjacencies per interface	256
OSPFv3 areas	64
OSPFv3 passive interfaces	1,000
OSPFv3 routers per area	250
OSPFv3 routes	64,000
PIM interfaces	512 active; 4084 passive
RIP IPv4 routes	3 * fastpath routes
RIP instances	64 (one for each VRF)
RIP interfaces	200
RIP routes	2,500 for each VRF

	Maximum number supported
	10,000 for each system
RSMLT interfaces (IPv4/IPv6)	4,000 over 512 SMLT interfaces
Static ARP entries	2,048 for each VRF 10,000 for each system
Static routes (IPv4)	2,000 for each VRF 10,000 total across VRFs
UDP/DHCP forwarding entries	512 for each VRF 1,024 for each system
VRRP interfaces (IPv4)	255 for a VRF 512 for a system
VRRP interfaces (IPv6)	255 for a system
VRRP interfaces fast timers (200ms)	24
Diagnostics	
Mirrored ports	479
Remote Mirroring Termination (RMT) ports	32
Filters and QoS	
Flow-based policers (IPv4 and IPv6)	16,000
Port shapers (IPv4 and IPv6)	480
Access control lists (ACL) for each chassis	2,048 The current release does not support IPv6 filters
Access control entries (ACE) for each chassis (IPv4)	16,000
ACEs per ACL (a combination of Security and QoS ACEs)	1,000
Unique redirect next hop values for ACE Actions (IPv4)	2,000
SPBM	
ARP entries (routed)	64,000
MAC entries	128,000 (combination of ARP entries and Layer 2 MACs)
Backbone MAC	1,000
IP routes in the Global Router	100,000 (combination of OSPF and IS-IS)
IS-IS adjacencies	64
Layer 2 VSNs	4,000

	Maximum number supported
VLANs in VRF	1,600
Layer 3 VSNs	512

Chapter 5: Supported standards, request for comments, and Management Information Bases

This chapter details the standards, request for comments (RFC), and Management Information Bases (MIB) that Avaya Virtual Services Platform 9000 supports.

Supported standards

The following table details the standards that Avaya Virtual Services Platform 9000 supports.

Table 17: Supported standards

Standard	Description
802.1ah	Provider Backbone Bridges (MacInMac encapsulation)
802.1aq	Shortest Path Bridging (SPB)
802.1AX	Link Aggregation Control Protocol (LACP)
802.1p	VLAN prioritization
802.1Q	Virtual Local Area Network (VLAN) tagging
802.1Qbb	Virtual Bridged Local Area Networks - Amendment: Priority-based Flow Control
802.1s	Multiple Spanning Tree Protocol
802.1t	802.1D maintenance
802.1w-2001	Rapid Spanning Tree protocol (RSTP)
802.1X	Extended Authentication Protocol (EAP), and EAP over LAN (EAPoL)
802.1X-2004	Port Based Network Access Control
802.3 CSMA/CD Ethernet ISO/IEC 8802	International Organization for Standardization (ISO) /International Eletrotechnical Commission (IEC) 8802-3

Standard	Description
802.3ab	Gigabit Ethernet 1000BaseT 4 pair Category 5 (Cat5) Unshieled Twisted Pair (UTP)
802.3ae	10 Gigabit Ethernet
802.3an	10 Gigabit Copper
802.3i	10BaseT
802.3u	100BaseT
802.3x	flow control
802.3z	Gigabit Ethernet

Supported RFCs

The following table and sections list the RFCs that Avaya Virtual Services Platform 9000 supports.

Table 18: Supported request for comments

Request for comment	Description
RFC768	UDP Protocol
RFC783	Trivial File Transfer Protocol (TFTP)
RFC791	Internet Protocol (IP)
RFC792	Internet Control Message Protocol (ICMP)
RFC793	Transmission Control Protocol (TCP)
RFC826	Address Resolution Protocol (ARP)
RFC854	Telnet protocol
RFC894	A standard for the Transmission of IP Datagrams over Ethernet Networks
RFC896	Congestion control in IP/TCP internetworks
RFC903	Reverse ARP Protocol
RFC906	Bootstrap loading using TFTP
RFC950	Internet Standard Subnetting Procedure
RFC951	BootP
RFC959, RFC1350, and RFC2428	IPv6 FTP and TFTP client and server

Request for comment	Description
RFC1027	Using ARP to implement transparent subnet gateways/Nortel Subnet based VLAN
RFC1058	RIPv1 Protocol
RFC1112	IGMPv1
RFC1122	Requirements for Internet Hosts
RFC1253	OSPF
RFC1256	ICMP Router Discovery
RFC1258	IPv6 Rlogin server
RFC1305	Network Time Protocol v3 Specification, Implementation and Analysis
RFC1340	Assigned Numbers
RFC1519	Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy
RFC1541	Dynamic Host Configuration Protocol1
RFC1542	Clarifications and Extensions for the Bootstrap Protocol
RFC1583	OSPFv2
RFC1587	The OSPF NSSA Option
RFC1591	DNS Client
RFC1723	RIP v2 – Carrying Additional Information
RFC1745	BGP / OSPF Interaction
RFC1771 and RFC1772	BGP-4
RFC1812	Router requirements
RFC1866	HyperText Markup Language version 2 (HTMLv2) protocol
RFC1965	BGP-4 Confederations
RFC1966	BGP-4 Route Reflectors
RFC1981	Path MTU discovery
RFC1997	BGP-4 Community Attributes
RFC1998	An Application of the BGP Community Attribute in Multi-home Routing
RFC2068	Hypertext Transfer Protocol
RFC2131	Dynamic Host Control Protocol (DHCP)

Request for comment	Description
RFC2138	RADIUS Authentication
RFC2139	RADIUS Accounting
RFC2178	OSPF MD5 cryptographic authentication / OSPFv2
RFC2236	IGMPv2 for snooping
RFC2270	BGP-4 Dedicated AS for sites/single provide
RFC2328	OSPFv2
RFC2338	VRRP: Virtual Redundancy Router Protocol
RFC2362	PIM-SM
RFC2385	BGP-4 MD5 authentication
RFC2439	BGP-4 Route Flap Dampening
RFC2453	RIPv2 Protocol
RFC2460	IPv6 base stack
RFC2464	Transmission of IPv6 packets over Ethernet networks
RFC2545 and RFC4710	IPv6 capable BGPv4+
RFC2616	IPv6 HTTP server
RFC2710 and RFC3810	MLD (host-mode only)
RFC2740	OSPFv3
RFC2796	BGP Route Reflection – An Alternative to Full Mesh IBGP
RFC2819	RMON
RFC2918	Route Refresh Capability for BGP-4
RFC2992	Analysis of an Equal-Cost Multi-Path Algorithm
RFC3046	DHCP Option 82
RFC3065	Autonomous System Confederations for BGP
RFC3162	IPv6 RADIUS client
RFC3315	IPv6 DHCP Relay
RFC3376	Internet Group Management Protocol, v3
RFC3411 and RFC2418	SNMP over IPv6 networks

Request for comment	Description
RFC3513	Internet Protocol Version 6 (IPv6) Addressing Architecture
RFC3569	An overview of Source-Specific Multicast (SSM)
RFC3587	IPv6 Global Unicast Address Format
RFC3768 and draft-ietf-vrrp-ipv6-spec-08.txt	IPv6 capable VRRP
RFC4213	IPv6 configured tunnel
RFC4250–RFC4256	SSH server and client support
RFC4443	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
RFC4861	IPv6 Neighbor discovery
RFC4862	IPv6 stateless address autoconfiguration
RFC4893	BGP support for Four-octet AS Number Space

Quality of service

Table 19: Supported request for comments

Request for comment	Description
RFC2474 and RFC2475	DiffServ Support
RFC2597	Assured Forwarding PHB Group
RFC2598	An Expedited Forwarding PHB

Network management

Table 20: Supported request for comments

Request for comment	Description
RFC1155	SMI
RFC1157	SNMP
RFC1215	Convention for defining traps for use with the SNMP

Request for comment	Description
RFC1269	Definitions of Managed Objects for the Border Gateway Protocol: v3
RFC1271	Remote Network Monitoring Management Information Base
RFC1305	Network Time Protocol v3 Specification, Implementation and Analysis3
RFC1350	The TFTP Protocol (Revision 2)
RFC1354	IP Forwarding Table MIB
RFC1389	RIP v2 MIB Extensions
RFC1757	Remote Network Monitoring Management Information Base
RFC1907	Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)
RFC1908	Coexistence between v1 & v2 of the Internet-standard Network Management Framework
RFC1930	Guidelines for creation, selection, and registration of an Autonomous System (AS)
RFC2541	Secure Shell Protocol Architecture
RFC2571	An Architecture for Describing SNMP Management Frameworks
RFC2572	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
RFC2573	SNMP Applications
RFC2574	User-based Security Model (USM) for v3 of the Simple Network Management Protocol (SNMPv3)
RFC2575	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
RFC2576	Coexistence between v1, v2, & v3 of the Internet standard Network Management Framework
RFC2819	Remote Network Monitoring Management Information Base

MIBs

Table 21: Supported request for comments

Request for comment	Description
RFC1156	MIB for network management of TCP/IP
RFC1212	Concise MIB definitions
RFC1213	TCP/IP Management Information Base
RFC1354	IP Forwarding Table MIB
RFC1389	RIP v2 MIB Extensions
RFC1398	Ethernet MIB
RFC1442	Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)
RFC1450	Management Information Base for v2 of the Simple Network Management Protocol (SNMPv2)
RFC1573	Interface MIB
RFC1650	Definitions of Managed Objects for the Ethernet-like Interface Types
RFC1657	BGP-4 MIB using SMIv2
RFC1724	RIPv2 MIB extensions
RFC1850	OSPF MIB
RFC2021	RMON MIB using SMIv2
RFC2096	IP Forwarding Table MIB
RFC2452	IPv6 MIB: TCP MIB
RFC2454	IPv6 MIB: UDP MIB
RFC2466	IPv6 MIB: ICMPv6 Group
RFC2578	Structure of Management Information v2 (SMIv2)
RFC2674	Bridges with Traffic MIB
RFC2787	Definitions of Managed Objects for the Virtual Router Redundancy Protocol
RFC2863	Interface Group MIB

Request for comment	Description
RFC2925	Remote Ping, Traceroute & Lookup Operations MIB
RFC2932	IPv4 Multicast Routing MIB
RFC2933	IGMP MIB
RFC2934	PIM MIB
RFC3416	v2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)
RFC4022	Management Information Base for the Transmission Control Protocol (TCP)
RFC4113	Management Information Base for the User Datagram Protocol (UDP)

Standard MIBs

The following table details the standard MIBs that Avaya Virtual Services Platform 9000 supports.

Table 22: Supported MIBs

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STDMIB2— Link Aggregation Control Protocol (LACP) (802.3ad)	802.3ad	ieee802-lag.mib
STDMIB3—Exensible Authentication Protocol Over Local Area Networks (EAPoL) (802.1x)	802.1x	ieee8021x.mib
STDMIB4—Internet Assigned Numbers Authority (IANA) Interface Type	—	iana_if_type.mib
STDMIB5—Structure of Management Information (SMI)	RFC1155	rfc1155.mib
STDMIB6—Simple Network Management Protocol (SNMP)	RFC1157	rfc1157.mib

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STDMIB7—MIB for network management of Transfer Control Protocol/Internet Protocol (TCP/IP) based Internet MIB2	RFC1213	rfc1213.mib
STDMIB8—A convention for defining traps for use with SNMP	RFC1215	rfc1215.mib
STDMIB9—Routing Information Protocol (RIP) version 2 MIB extensions	RFC1389	rfc1389.mib
STDMIB10—Definitions of Managed Objects for Bridges	RFC1493	rfc1493.mib
STDMIB11—Evolution of the Interface Groups for MIB2	RFC2863	rfc2863.mib
STDMIB12—Definitions of Managed Objects for the Ethernet-like Interface Types	RFC1643	rfc1643.mib
STDMIB13—Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIv2	RFC1657	rfc1657.mib
STDMIB14—RIP version 2 MIB extensions	RFC1724	rfc1724.mib
STDMIB15—Remote Network Monitoring (RMON)	RFC2819	rfc2819.mib
STDMIB16—Open Shortest Path First (OSPF) Version 2	RFC1850	rfc1850.mib
STDMIB17—Management Information Base of the Simple Network Management Protocol version 2 (SNMPv2)	RFC1907	rfc1907.mib
STDMIB21—Interfaces Group MIB using SMIv2	RFC2233	rfc2233.mib
STDMIB26a—An Architecture for Describing	RFC2571	rfc2571.mib

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
SNMP Management Frameworks		
STDMIB26b—Message Processing and Dispatching for the SNMP	RFC2572	rfc2572.mib
STDMIB26c—SNMP Applications	RFC2573	rfc2573.mib
STDMIB26d—User-based Security Model (USM) for version 3 of the SNMP	RFC2574	rfc2574.mib
STDMIB26e—View-based Access Control Model (VACM) for the SNMP	RFC2575	rfc2575.mib
STDMIB26f —Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework	RFC2576	rfc2576.mib
STDMIB29—Definitions of Managed Objects for the Virtual Router Redundancy Protocol	RFC2787	rfc2787.mib
STDMIB31—Textual Conventions for Internet Network Addresses	RFC2851	rfc2851.mib
STDMIB32—The Interface Group MIB	RFC2863	rfc2863.mib
STDMIB33—Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations	RFC2925	rfc2925.mib
STDMIB34—IPv4 Multicast Routing MIB	RFC2932	rfc2932.mib
STDMIB35—Internet Group Management Protocol MIB	RFC2933	rfc2933.mib
STDMIB36—Protocol Independent Multicast MIB for IPv4	RFC2934, RFC2936	rfc2934.mib, rfc2936.mib

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STDMIB38—SNMPv3 These Request For Comments (RFC) make some previously named RFCs obsolete	RFC3411, RFC3412, RFC3413, RFC3414, RFC3415	rfc2571.mib, rfc2572.mib, rfc2573.mib, rfc2574.mib, rfc2575.mib
STDMIB39—Entity Sensor Management Information Base	RFC3433	
STDMIB40—The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model	RFC3826	rfc3826.mib
STDMIB41—Management Information Base for the Transmission Control protocol (TCP)	RFC4022	rfc4022.mib
STDMIB43—Management Information Base for the User Datagram Protocol (UDP)	RFC4113	rfc4113.mib
STDMIB44—Entity MIB	RFC4133	rfc4133.mib
STDMIB46—Definitions of Managed Objects for BGP-4	RFC4273	rfc4273.mib

Proprietary MIBs

The following table details the proprietary MIBs that Avaya Virtual Services Platform 9000 supports.

Table 23: Proprietary MIBs

Proprietary MIB name	File name
PROMIB1 - Rapid City MIB	rapid_city.mib
PROMIB 2 - SynOptics Root MIB	synro.mib
PROMIB3 - Other SynOptics definitions	s5114roo.mib
PROMIB4 - Other SynOptics definitions	s5tcs112.mib

Proprietary MIB name	File name
PROMIB5 - Other SynOptics definitions	s5emt103.mib
PROMIB6 - Avaya RSTP/MSTP proprietary MIBs	nnrst000.mib, nnmst000.mib
PROMIB7 - Avaya IGMP MIB	rfc_igmp.mib
PROMIB8 - MIAvayal IP Multicast MIB	ipmroute_rcc.mib
PROMIB9 - Avaya PIM MIB	pim-rcc.mib
PROMIB11 - Avaya MIB definitions	wf_com.mib

Chapter 6: Known issues and limitations

This section details the known issues and limitations of the Avaya Virtual Services Platform 9000. Where appropriate, use the workarounds provided.

Known issues

The following sections identify the known issues in this release of the Avaya Virtual Services Platform 9000.

Alarm, logging, and error reporting

Table 24: Known issues

Issue number	Description	Workaround
wi00519967	Log filtering based on date and time is currently unavailable.	—
wi00768362	The **show alarm database** command does not show the CPU from which the alarm originated. The log report contains this information.	Use the **show logging file alarm** command to check the complete details and the CPU from which the alarm originated.
wi00980601	Enabling tracing to log can affect system performance in scaled environments.	Avaya recommends that you capture tracing by using a Telnet session, and capture to a file on the local management system to avoid potential performance issues. Verify that **boot config flags trace-logging** is not present in the configuration file to assure this feature is not enabled before you enable tracing.

Issue number	Description	Workaround
wi00986085	After a CP switch-over the following message can appear in the logs: `smltTick: Initial MAC/ARP tbl completed, unlocked SMLT/SLT ports.` The SMLT/SLT ports were not locked during the switch-over. This message is incorrectly reused in this situation.	—
wi01004076	You can see the following error message when you boot the VSP 9000: `HW ERROR framework_process_entity_data: Application Sync failed for entity: 0x414c524d representing Module ALARM`	This message has no functional impact and can be ignored.
wi01005062	You can see the following message on a running system: `GlobalRouter KHI ERROR Could not read /proc/3824/task/3824/stat`. This message is a result of a race condition in the KHI polling process.	This message has no functional impact and can be ignored.

Applications

Table 25: Known issues

Issue number	Description	Workaround
wi00940744	The command **`show application vsptalk client`** does not show connection status.	View this information in the log file.

Issue number	Description	Workaround
wi01002102	ACLI does not enforce the VSP Talk restriction of no all-numeric passwords.	Do not use all-numeric passwords.
wi01002448	Deleting VSP Talk does not disable the messaging to the client.	Disable the IM connection first, and then delete the gtalk or Avaya configuration.

Chassis operations

Table 26: Known issues

Issue number	Description	Workaround
wi00564595	If there is not enough power available in the chassis to power all cards when the system is powered up, one or more cards will not be powered on. Configuration for these cards will be ignored. When enough power is available the cards will be automatically powered up but they will not automatically receive their configuration.	To download the configuration to any cards that experience delayed boot up, source the configuration for that card.
wi00891718	Unable to access `/usb` from the peer CP.	Use TFTP from the peer or configure the network management port and use a transport protocol like FTP or TFTP, directly to the secondary CP.
wi00930215	After you use the **`extflash-stop`** command to unmount the external Compact Flash, you must remove and reinsert the Compact Flash to remount the volume.	—
wi00969922	If you remove the backup CP module, you can see the following output on the console:	This message has no functional impact and can be ignored.

Issue number	Description	Workaround
	`fbuf allocated in "/vob/cb/nd_platform/chassis/lib/ch_sync.c" at line 341 is freed` This message occurs if an application tries to synchronize data to the backup CP module at the same time that you remove the module.	
wi00970236	The default value for the loadingconfig time is 15 minutes. The configurable range for the **`boot config loadconfigtime`** command is 0 to 300 seconds. If you configure a value that is less than the default, the device still uses the default value to validate the loading time. Because the maximum configurable value is 300 seconds, the value is always less than the default and does not take effect. The intent of the parameter is the time to load the configuration. The timer that runs in the VSP 9000 actually tracks the full start time, for example, the time spent waiting for other IO ready and to download port MAC.	—

COM

Table 27: Known issues

Issue number	Description	Workaround
wi00949006	You cannot view more than 50 ARP entries using COM.	In ACLI, use **`show ip arp`**. In EDM, use **Configuration** > **IP** > **IP** > **ARP**.

EDM

Table 28: Known issues

Issue number	Description	Workaround
wi00668629	After deleting a VRF, all open EDM sessions that are using that VRF need to be closed manually by the user.	—
wi00830411	Error messages may be displayed on the console when connected to EDM using Internet Explorer 6.0. Internet Explorer 6.0 is not supported.	Discontinuing use of this browser will discontinue the error messages.
wi00948868	EDM can take a significant amount of time to capture and display the MAC and ARP table with the maximum number of 128K MAC and 64K ARP entries.	—
wi00956046	You cannot use FireFox 7.x to connect to an IPv6 address with HTTPS. The connection appears as untrusted, and if you select the option to add a security exception, the browser displays an error. This issue is a known Mozilla bug (633001).	Use FireFox 3.x or Internet Explorer.

Issue number	Description	Workaround
wi00962229	EDM does not show the TTL field in the ARP table.	Use ACLI.
wi00965260	Do not use **VLAN** > **VLANs** > **IP** > **RSMLT** tab to configure RSMLT hold-down timer and hold-up timer parameters for IPv6.	Use **IPv6** > **RSMLT**to configure the IPv6 only interface RSMLT hold-down timer and hold-up timer interfaces.
wi00968931	You cannot configure EDM to use an IPv6 address for the Help directory on a TFTP server.	Use an IPv4 TFTP server. Store the Help files on the IPv4 TFTP server, and then configure EDM to use the IPv4 server address.
wi00969780	EDM does not show the tunnel field in the ARP table.	Use ACLI.
wi00994286	EDM does not support cross VRF route redistribution. You can use only the current VRF context as the destination and source VRF, and you can use only the route policies you create under the current VRF context. For example, if you use VRF1 as the VRF context, the destination and source VRF is VRF1. For route redistribution, you can use only policies created under the VRF1 context.	Use ACLI.
wi01004052	EDM does not support the `/intflash`, `/extflash`, or `/usb` path options for the EDM Help files.	You must install the EDM Help files on a TFTP server.

HA operations

Table 29: Known issues

Issue number	Description	Workaround
wi00937861	The **`exception dump max-disk-space`** configuration ACLI	—

Issue number	Description	Workaround
	command is not available on the standby CPU.	

Hardware

Table 30: Known issues

Issue number	Description	Workaround
wi01005012	If you hot insert an SFP or SFP+ into a 9024XL or 9048GB module, the port is erroneously enabled even if the port is administratively disabled. The LED changes to solid amber, which indicates the port is enabled, or solid green if it connects to a live port.	The problem corrects itself after you administratively enable the port before or after you insert the SFP or SFP+. If you need to keep the port administratively disabled, enable and disable the port.
wi01006381	A port can go into the blocking state if you change its configuration while the interface module to which it belongs is not present, for example, is removed or powered down.	Do not make configuration changes that can affect a port while the interface module is in this state.

Management and general administration

Table 31: Known issues

Issue number	Description	Workaround
wi00509904	File transfer may fail when attempting to move large files with TFTP.	Use FTP for transfer of files larger than 32MB.
wi00510551	Compression options are not supported in SSHv2 but no error message is displayed when they are used.	Do not use compression options with SSHv2.

Issue number	Description	Workaround
wi00520113	Transferring files using passive FTP may fail when using a Windows PC.	Use active mode when transferring files with FTP.
wi00932777	The current release does not support IPv6 addresses for the **`traceroute`** command. VSP 9000 will add IPv6 traceroute support in a future release.	—
wi00979353	The ACLI command to configure the SNMPv3 trap target entry does not support the entry name configuration. The name is derived internally from the IP address and port number by using the MD5 hash. If you use EDM to create the trap target entry, the specified entry name is not retained after you use the **`save config`** command and restart the system. The name will be derived from the host IP address and port number.	—
wi01002533	Executing a MIB walk on the proprietary Routing Table MIB .1.3.6.1.4.1.2272.1.8.7 when alternative routes are present, and the alternative route nexthop is less than the best route nexthop, will fail.	Use the Standard MIB RFC1213 ipRouteTable MIB .1.3.6.1.2.1.4.21: for access to the routing table.

MLT, SMLT, and link aggregation

Table 32: Known issues

Issue number	Description	Workaround
wi00822560	Disable member ports before deleting an MLT.	—

Issue number	Description	Workaround
wi00822571	In rare occurrences traffic loops can be introduced if ports are removed from and MLT before being disabled.	The operator must disable participating ports before removing them from the MLT, or deleting the MLT completely.

Multicast

Table 33: Known issues

Issue number	Description	Workaround
–	The ACLI query-interval parameter for PIM interfaces and ports is changed to hello-interval. You can still use the query-interval parameter, but the configuration file saves this information as hello-interval and query-interval does not appear in ACLI Help text.	—
–	If you need to upgrade the IGMP version of routers in a network that operates using a lower version of IGMP (caused by the dynamic downgrade function), perform the following actions on all routers: 1. Disable PIM on the interface on which you need to upgrade the IGMP version. 2. Change the IGMP version to the required value. 3. Enable PIM on the interface you just upgraded.	—

Patching and upgrading

Table 34: Known issues

Issue number	Description	Workaround
wi00511642	The **software patch commit** and **software patch remove** commands will not display messages such as Syncing release directory on backup CP card in slot 2 while executing the command in a Telnet session.	—
wi00888516	If you apply multiple patches using the patch-id parameter, and at least one patch is a candidate and at least one patch is a non-candidate, the system returns an error message. The error message identifies the non-candidate patch but does not indicate the other patches were applied, even though they were.	Use the **show software patch** command to see the status of the patches.
wi00989121	When you upgrade the software image, a slight chance exists that one of the Switch Fabric or interface modules can fail to upgrade, which results in a rollback to the previous release.	After the upgrade, use the **show system software** command to verify that the upgrade was successful. If the upgrade was not successful, activate the Release 3.3 software again.
wi00990115	The console can display error messages when you revert a patch.	The patch revert will be successful and you can ignore the error messages.
wi01007353	Modules with factory images do not upgrade if the CP module in slot 2 is the master and the CP module in slot 1 is not operational.	Ensure the CP module in slot 1 is operational.

Routing

Table 35: Known issues

Issue number	Description	Workaround
wi00703966	There is a problem that can cause the standby CP to reboot during the synchronization of a large number of BGP routes from the master CP. If this does occur, the standby CP will reboot and attempt to synchronize with the master CP again. It will very likely synchronize correctly on the second attempt. During the time of standby CP reboot and resynchronization the traffic flow through the system is not effected because the master CP stays operational.	—
wi00974143	You cannot change the OSPF area for an IPv6 interface. You must delete the interface from one OSPF area, and then create the new OSPF area.	—
wi01014704	The **`show ip route`** ACLI command shows incorrect routing protocols for inter-VRF redistributed routes. This issue is a display issue only, and does not have functional impact on inter-VRF route redistribution.	—

SPBM and IS-IS

Table 36: Known issues

Issue number	Description	Workaround
wi01000756	Removing a 9024XL module that was configured for IS-IS/SPBM does not properly remove all IS-IS configuration. This situation is not an issue if you replace the module with another 9024XL module.	If you install a 1 GB module after you remove a 9024XL module that runs IS-IS, you must remove IS-IS from the slot configuration before you replace the module.
wi01004034	The ERS **`show isis spbm show-all`** command is not available on VSP 9000.	—

VLAN operations

Table 37: Known issues

Issue number	Description	Workaround
wi01005885	If you configure NLB unicast mode, you can ping the NLB cluster address, but cannot ping the physical server address.	To ping the physical IP address, use a device that is on the local VLAN.

Interoperability issues

The following table lists the known issues between this release of the Avaya Virtual Services Platform 9000 and other Avaya products. The following table also identifies if the issue is fixed in a specific release. Note that the issue will still exist in previous releases.

Table 38: Interoperability issues

Issue number	Description	Workaround	Fixed in release
wi00511257	If you change the priority of, and then disable and enable the MLT port on an Ethernet Routing Switch 8600, the port takes 35 seconds to become the designated forwarding port on the root bridge. This condition causes traffic interruption for 35 seconds.		ERS 8600 5.1.7.0
wi00565499	If you use VSP and Ethernet Routing Switch 8600 on a VLAN, and all systems operate in MSTP mode, a loop can be generated if you restart a VLAN port on the Ethernet Routing Switch.	Disable the links on the Ethernet Routing Switch 8600 to remove the loop.	ERS 8600 5.1.8.0
wi00689238	If a VSP 9000 aggregation switch sends a high volume (more than 3000) of OSPF or RIP routes to an Ethernet Routing Switch 8600 edge device to redistribute into an OSPF domain, the CPU utilization of the edge device can increase, which results in dropping all VLACP packets from the VSP device. The VLACP link operational state is down.		ERS 8600 5.1.6.0
wi00691506	A topology change of an SMLT link between VSP 9000 systems and Ethernet Routing Switch 8600 Release 5.1.3 results in dropped packets. This problem occurs when one of the two MLT ports of the ERS is not a designated port or a root port. Topology changes make this port a blocking port		ERS 8600 5.1.4.0

Issue number	Description	Workaround	Fixed in release
	and also other ports of the MLT change to the same state (blocking).		

Limitations

This section lists known limitations and expected behaviors that may first appear to be issues. The following table provides a description of the limitation or behavior and the work around, if one exists.

Table 39: Limitations and expected behaviors

Issue number	Description	Workaround
wi00511527	MSTP bridges may not learn the correct CIST regional root.	If you encounter this problem you can change the bridging priority of the switch to make sure the root selection occurs as desired.
wi00664833	The MAC DA filter only applies for traffic that is bridged through the device. If the packet is routed, then the legacy MAC DA filter does not apply for traffic that is routed through the box.	Use ACL-based filters to implement the MAC DA filter. The ACL-based filter works correctly regardless of whether the packet is bridged or routed.
wi00732215	When all members of an LACP aggregation go down, the ARP record corresponding to the aggregation gets deleted and needs to be re-ARPed to forward traffic across IST.	To work around this problem use regular MLT interfaces instead of LACP interfaces. If LACP is required, the traffic recovery time will be between 1-12 seconds based on volume of re-ARPing required.
wi00733551	The Bandwidth Allocation Group (BAG) rate configuration of all ports is based on the maximum port speed of the module during the system bootup time. When you configure an interface shaper and it is lower than the maximum port speed, the BAG rate	

Issue number	Description	Workaround
	becomes larger than the port forwarding rate. This condition is an incorrect Qos configuration. As a result, low priority traffic is not dropped as expected.	
wi00820028	You should clear the cache of the browser used to configure and monitor the device after an image upgrade. If this is not done incorrect screen displays can result.	Clearing the browser cache is found in **Tools** > **Internet Options** > **Browser History** > **Delete** > **Delete all ….** in Internet Explorer 7.0 and in **Tools** > **Clear Recent History** > **Select all options** > **Clear Now** in Firefox 3.6.x.
wi00854206	For VLACP enabled links, recommended values exist for the configuration of the VLACP timers. However, in an SMLT topology, with VLACP and multicast both enabled on the SMLT link, you may need to adjust or increase the VLACP timers on that link to accommodate for a scaled multicast environment where there is a higher processing load on the CP, especially during failover events. This higher load can affect the ability to process VLACP keep-alive messages in a timely manner, which can cause the link to flap. You may need to configure timers proportionately to the anticipated multicast route load.	Avaya recommends that you configure IST links with a VLACP timeout of long, timeout scale of 3, and slow-periodic-time of 30 000 ms. These links are not impacted by multicast scaling considerations.
wi00981875	VSP 9000 management-plane-initiated applications that do not have VRF specific context are biased toward the Management Router routing table. When you configure a default route on both the Management Router and the Global	If you require both in-band management (Global Router) and out-of-band management (Management Router), the default route should not be present on the Management Router. Configure static routes for specific management

Issue number	Description	Workaround
	Router, the default route on the Management Router takes precedence.	networks in the desired VRF instead.

MIB errors

The following errors exist in the mib.txt and mib_vsp.txt file in the VSP 9000 3.3 Release. These errors do not impact the functionality of the device.

- The mib.txt file and mib_vsp.txt file specify an unsupported mib object, `rcIpRoutePolicySetLevel`. Ignore or remove this object when you use the mib.txt and mib_vsp.txt files.
- A line in the mib_vst.txt that includes the igmpInterfaceQuerier mib object is misspelled:

  ```
  igmpInterfaceIfIndex, igmpIntrfaceQuerier          FROM IGMP-MIB
  ```

  should be replaced with:

  ```
  igmpInterfaceIfIndex, igmpInterfaceQuerier        FROM IGMP-MIB
  ```

- A line in the mib_vsp.txt has a truncated mib object name.

  ```
  ForwardingCbZagrosPool9Peak  OBJECT-TYPE
  ```

  should be replaced with:

  ```
  rcKhiForwardingCbZagrosPool9Peak  OBJECT-TYPE
  ```

- A line in the mib_vsp.txt is missing a comma after an enumeration value (bySrcMac) for rcVlanType mib object.

  ```
  bySrcMac(4)
  spbm-bvlan(11)
  ```

  should be replaced with:

  ```
  bySrcMac(4),
  spbm-bvlan(11)
  ```

Chapter 7: Resolved issues

This section identifies the issues resolved in this release.

Alarm, logging, and error reporting

Table 40: Resolved issues

Issue number	Issue description
wi00828941	After an HA failover, if KHI polling polls the other CP, but that CP is still starting, the KHI poll times out and a letter sync error message appears.
wi00944915	The message `License trial period has expired` appears in the log every 30 minutes even if only the Basic feature set is used. This fix changes the frequency of this message to every 24 hours.
wi00969817	When you configure an access policy for SSH access, and an SSH session is attempted and blocked due to the access policy, there is no message logged in the VSP 9000 system log to indicate that a session was blocked.
wi00974862	The **`sys ecn-compatibility`** command to enable explicit congestion notification does not apply to VSP 9000.
wi00975675	When you disable and then enable VSP Talk, no log message indicates that VSP Talk is enabled.
wi00980491	The **`show alarm database`** command shows a link down even though it is operational.

Applications

Table 41: Resolved issues

Issue number	Issue description
wi00975662	When you use the **`vsptalk gtalk enable`** command, the system checks for an IPv4 DNS configuration. If a DNS primary name server is not configured, VSP Talk does not connect to the Google Talk IM server.

Chassis operations

Table 42: Resolved issues

Issue number	Issue description
wi00967400	You can see the following message at system start: `GlobalRouter HW ERROR portGetPortNum: invalid physical port 0724`. This message has no functional impact and can be ignored.
wi00968592	A KHI warning message can appear after you stop and remove a USB device. This message is not functionally impacting and can be ignored.
wi00988403	If you configure the IPFIX collector parameters, the configuration does not save correctly. After you save the configuration, if the verify_config flag is TRUE, the configuration will not boot. If the verify_config flag is not TRUE, the parameter is skipped on boot.
wi00998449	A malformed DHCP Option 82 reply packet received without an end option in the packet can cause the CP to crash.

EDM

Table 43: Resolved issues

Issue number	Issue description
wi00825617	EDM may display port 1/48 in the VLAN forwarding table.
wi00832390	Embedded EDM does not work if you do not configure SNMP community names.
wi00889572	When you use EDM to view a specific VRF context, only VRF 0 appears.
wi00975479	The VLAN range for a brouter port should be 2 to 4084.
wi00975511	You cannot delete a static IGMP group using EDM.
wi00978158	You can configure IPv6 **ReversePathCheckEnable** and **ReversePathCheckMode** on an IPv6 management interface even though this configuration has no meaning. The message `undoFailed` appears on EDM, and SW errors appear on the CP console. You can ignore these errors.
wi00984551	You cannot configure an IPv6 tunnel using EDM.

Management and general administration

Table 44: Resolved issues

Issue number	Issue description
wi00972377	Do not use spaces in an SNMP group. If you do use spaces, two errors print several times.
wi00974849	You cannot disable the UDP checksum.

Issue number	Issue description
wi00977397	You may be unable to use the **ls** command if you create an FTP connection over IPv6 from a Windows client. The display freezes for 30 seconds and the list does not display.
wi00980950	If the link from the VRRP master back to the edge goes down, the ICMP ping from the edge to the VRRP virtual interface fails in the following scenario: • A triangle SMLT configuration with VRRP, where ICMP ping to the virtual IP interface from the edge hashes to two different paths: a. Traffic to the core to the VRRP backup master. b. The return path from the VRRP master back to the edge. There is no impact to other data packets. After the link is restored, ICMP ping to the VRRP virtual interface will work as expected.
wi00981422	Some ACLI commands allow you to configure and enable the feature in a single command. However, some of these command result in an error because certain entries must be created before they can be enabled. For example, the **ipv6 dhcp-relay fwd-path 9:1:1::2 1::2 enable** command results in the following error: `Error: rcIpv6DhcpForwardTblSet failed for IfIndex 2050. Make sure the path is created before you enable it.`
wi00993044	A configured DHCP-Relay forward path is not always saved in the configuration file.
wi00996576	If you use the ACLI **login** command to login again with a different user name on an SSH session, the system crashes and reboots. If you use the ACLI **login** command to login again with a different user name on a Telnet or Rlogin session, an invalid login

Issue number	Issue description
	configures the access level of the current session to be NONE. Some ACLI commands are not accessible. You must logout, and login again to correct the access level. Log messages are not correct for relogin success or failure.

MLT, SMLT, and link aggregation

Table 45: Resolved issues

Issue number	Issue description
wi00976726	The **`show lacp int mlt x`** command will show information for nonexistent MLTs (information is correct), rather than not showing any information at all.
wi00976736	The maximum range for the IfIndex in the **`show lacp interface`** command is 6399 when it is actually 6655.
wi00998070	If a trunk port is part of an MLT and is configured to discard untagged frames, and you delete the MLT, the port becomes untagged but still discards untagged frames.

Patching

Table 46: Resolved issues

Issue number	Issue description
wi00732444	After a patch is reverted the patch action should be committed before moving on with more actions. The commit may be explicitly typed or the commit can occur via auto-commit after the configured timeout period. In this case the remove command was

Issue number	Issue description
	allowed and operated correctly without an error or warning indicating that it would not take effect until the commit occurred explicitly or via timeout.
wi00966662	If you issue a hitless patching revert command against an active applied hitless patch, there may be a short processing delay that can affect VLACP.
wi00971012	If you reset the chassis prior to committing a patch, the system clears the patch. A reset patch does not require an **`abort`** command. You do not need to enter an **`abort`** command after a reset even for a hitless patch. The reset clears the patch. If you do use an **`abort`** command after the reset, you will see a message that instructs you to reset the shelf to rollback the patch. The chassis has already rolled-back and no further reset is required.

QoS and filters

Table 47: Resolved issues

Issue number	Issue description
wi00970134	If you use egress filters with multicast traffic, using the VLAN attribute in the ACE to match the outgoing VLAN does not work.
wi00970587	If a configuration file with filter configuration is sourced on a switch that already has filters configured, FILTER and HW errors can occur. This does not cause any functional impact and can be ignored.

Routing

Table 48: Resolved issues

Issue number	Issue description
wi00962267	IPv6 Duplicate Address Detection (DAD) does not work on 10 Gb interface modules. If you configure duplicate IPv6 addresses between the VSP and another node on a 10 Gb link, the VSP software does not detect the duplicate scenario.
wi00997200	After a CP switchover, a mismatch between the ARP and IP forwarding tables on the IO can occur. This mismatch causes error messages each time the ARP is unresolved or ages out.
wi00976377	When you configure VRRP on an IPv6 interface, two router advertisements will be sent to the host instead of a single router advertisement.

Chapter 8: 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.