Avaya Virtual Services Platform 7000 Series Release Notes

Release 10.2
NN47202-400
Issue 03.02
October 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”).

Avaya grants you a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to you. “Designated Processor” means a single stand-alone computing device. “Server” means a Designated Processor that hosts a software application to be accessed by multiple users.

Licence 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: Introduction

Purpose

This document provides overview information about the new features supported in this software release for the Avaya Virtual Services Platform 7000 Series.

Related resources

Documentation

For a list of the documentation for this product, see *Avaya Virtual Services Platform 7000 Documentation Roadmap* (NN47202–103).

Training

Ongoing product training is available. For more information or to register, see http://avaya-learning.com/.

Enter the course code in the **Search** field and click **Go** to search for the course.

Course code	Course title
7D00080W	Avaya Stackable ERS and VSP Product Overview
7D00085V	Stackable ERS & VSP Installation, Configuration and Maintenance
7D00085I	Stackable ERS & VSP Installation, Configuration and Maintenance

Avaya Mentor videos

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

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

To find videos on the Avaya support site, select the product name, and check the *videos* checkbox to see a list of available videos.

Note:

Videos are not available for all products.

To find the Avaya Mentor videos on YouTube, go to http://www.youtube.com/AvayaMentor and perform one of the following actions:

- Enter a key word or key words in the Search Channel to search for a specific product or topic.
- Scroll down Playlists, and click the name of a topic to see the available list of videos posted on the site.

Support

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

Subscribing to e-notifications

Subscribe to e-notifications to receive an email notification when documents are added to or changed on the Avaya Support web site.

About this task

You can subscribe to different types of general notifications, for example, Product Correction Notices (PCN), that apply to any product or a specific product. You can also subscribe to specific types of documentation for a specific product, for example, Application & Technical Notes for Ethernet Routing Switch 8800.

Procedure

1. In an Internet browser, go to https://support.avaya.com

2. Type your username and password, and then click **LOG IN**.
3. Click **MY PROFILE**.

4. On the site toolbar, click your name, and then select **E Notifications**.

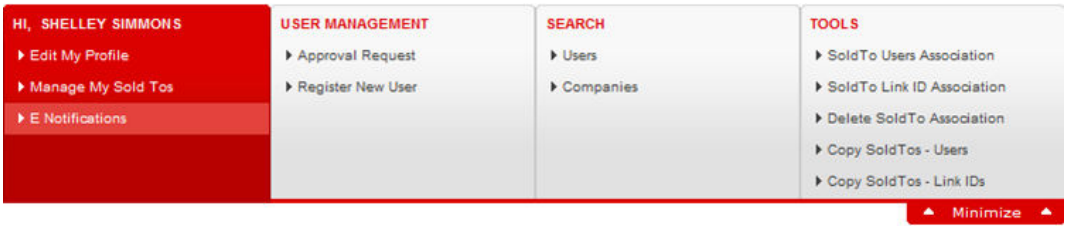

5. In the GENERAL NOTIFICATIONS area, select the required documentation types, and then click **UPDATE**.

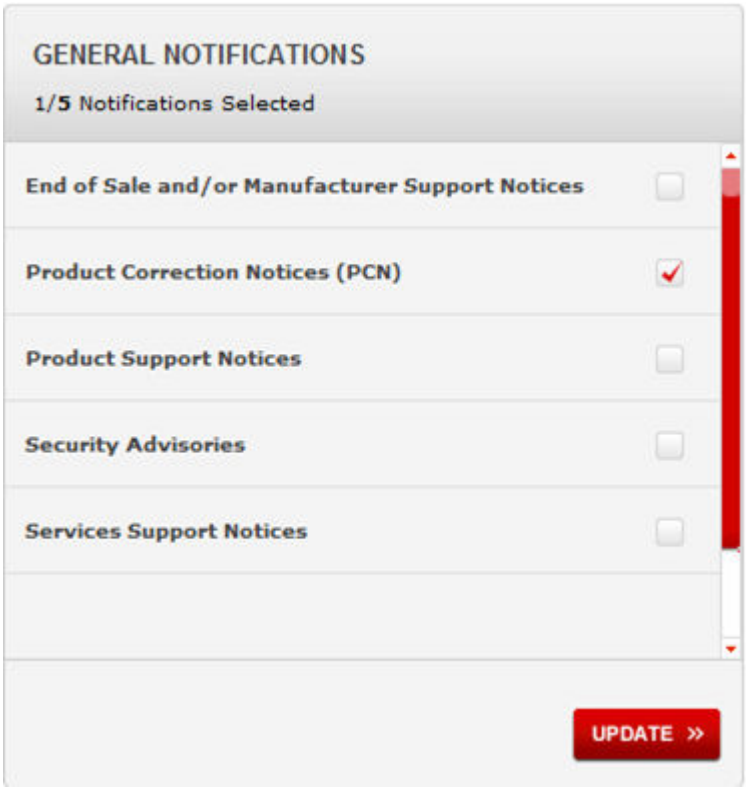

6. Click **OK**.
7. In the PRODUCT NOTIFICATIONS area, click **Add More Products**.

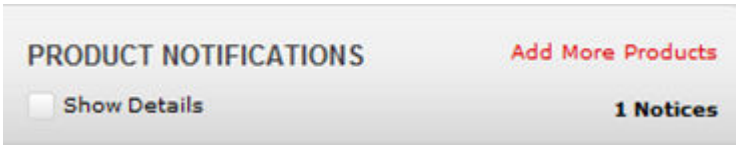

8. Scroll through the list, and then select the product name.
9. Select a release version.
10. Select the check box next to the required documentation types.

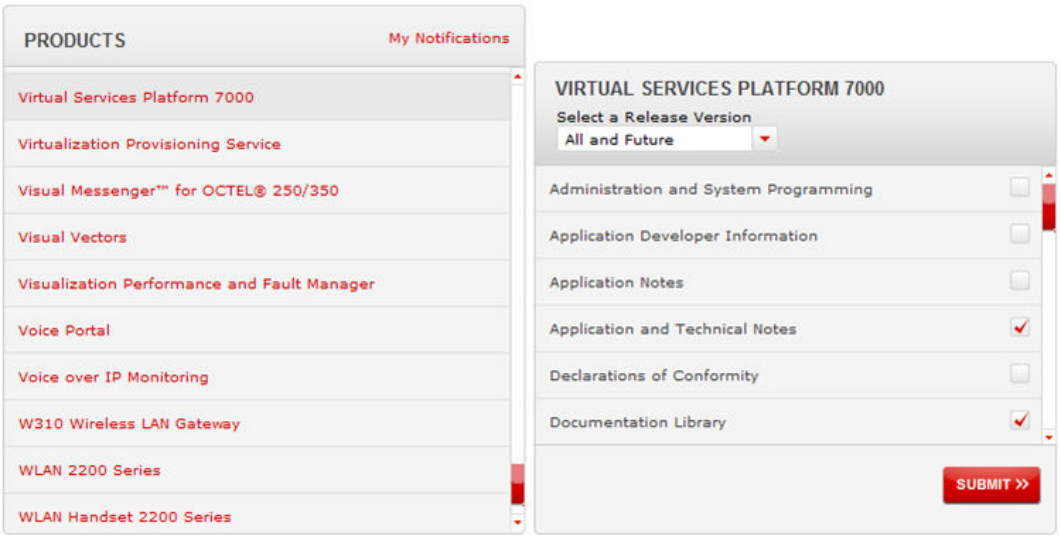

11. Click **Submit**.

Chapter 2: New in this release

The following sections detail what is new in Avaya Virtual Services Platform 7000 Series Release 10.2

Features

See the following sections for information about new features.

SPBM and IS-IS

Shortest Path Bridging MAC (SPBM) is a next generation virtualization technology that revolutionizes the design, deployment, and operations of enterprise campus core networks along with the enterprise data center. SPBM provides massive scalability while at the same time reducing the complexity of the network.

SPBM simplifies deployments by eliminating the need to configure multiple points throughout the network. When you add new connectivity services to an SPBM network you do not need intrusive core provisioning. The simple endpoint provisioning is done where the application meets the network, with all points in between automatically provisioned through the robust link-state protocol, Intermediate-System-to-Intermediate-System (IS-IS).

SPBM uses IS-IS to discover and advertise the network topology, which enables computation of the shortest path to all nodes in the SPBM network. SPBM uses IS-IS shortest path trees to populate forwarding tables for the individual B-MAC addresses of each participating node.

VSP 7000 Series supports the IEEE 802.1aq standard of SPBM, which allows for larger Layer 2 topologies and permits faster convergence. SPBM requires a Premier feature license.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - SPBM* (NN47202–510).

CFM

Connectivity Fault Management (CFM) operates at Layer 2 and provides an equivalent of ping and traceroute for SPBM that you can use to debug connectivity issues and to isolate faults. To support troubleshooting of the SPBM cloud, Virtual Services Platform 7000 supports a subset of CFM functionality.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - SPBM* (NN47202–510).

SMLT

Split Multi-Link Trunking (SMLT) is an extension of MLT that enables edge switches using MLT to dual-home to two SMLT aggregation switches. SMLT is transparent to the edge switches supporting MLT. In addition to link failure protection and flexible bandwidth scaling, SMLT improves the level of Layer 2 and Layer 3 resiliency by providing nodal protection.

Because SMLT inherently avoids loops, SMLT networks do not require the use of IEEE 802.1D Spanning Tree protocols to enable loop free triangle topologies.

SMLT avoids loops by enabling two aggregation switches to appear as a single device to edge switches, which are dual-homed to the aggregation switches. The aggregation switches are interconnected using an Inter-Switch Trunk (IST) that allows them to exchange addressing and state information (permitting rapid fault detection and forwarding path modification).

SMLT is primarily designed for Layer 2, however SMLT also provides benefits for Layer 3 networks.

Important:

When you enable SMLT, STP is automatically disabled on all SMLT, IST, and SLT ports in all spanning tree groups or instances to which the ports belong. You must also disable STP on all edge switch MLT ports that are connected to the SMLT or SLT.

Note:

Routing over SMLT is not supported in Release 10.2.

Rear port mode does not support SMLT or IST in Release 10.2

SMLT is supported on standalone units or in a stack. Avaya recommends a stack of at least three units to improve failure recovery. You can configure a maximum of 31 SMLT trunks on one device.

SMLT is supported on standalone or stacked units in triangle or square configuration. In a stack, the SMLT is active only on the base unit or the temporary base unit, and the base unit is solely responsible for the peer to peer switch communication.

You cannot configure SMLT data when SMLT is running. To modify an SLT or SMLT, you must disable SMLT on that port or trunk. IGMP over SMLT is not supported in Release 10.2.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Layer 2* (NN47202–502).

SMLT disables STP

Once you enable SMLT, the following actions occur automatically:

1. The current Spanning Tree Protocol (STP) administrative state of Inter Switch Trunk (IST), SMLT, and Split Link Trunk (SLT) ports is saved on the NVRAM
2. STP is disabled on IST, SMLT, and SLT ports.
3. Global IP Routing is enabled.

This feature reduces the differences of the SMLT configuration between the VSP 7000 Series and the Ethernet Routing Switch 8600 and 8800 Series.

LACP over SMLT (SMLT with LAGs)

Link Aggregation Control Protocol (LACP) over SMLT results in better recovery for SMLT and SLT-configured trunks in fail-over scenarios such as when a stack link breaks.

LACP dynamically creates and removes trunk groups. In the absence of STP on the SMLT network, configuration errors can easily introduce a loop. To limit loops, IST links do not support LACP: only SMLT and SLT links support LACP.

Note:

You must configure LACP timers from short to long when using LACP over SMLT.

SLPP

Simple Loop Prevention Protocol (SLPP) is a new feature designed to detect loops in a SMLT network. Not intended to replace STP as a comprehensive loop detection mechanism, SLPP acts as a secondary mechanism for detection and prevention of looping in a SMLT environment and can only be configured on SMLT networks. Since SMLT requires STP to be disabled on IST, SMLT and SLT ports for normal operation, loops may be introduced to a network. SLPP was designed to prevent such loops and resulting traffic disruptions.

When enabled, SLPP causes the switch to send a periodic SLPP PDU on the transmitting VLAN at a user defined or default (500 ms) transmission interval. If a loop is active in the network, the SLPP PDU is returned to the switch and the affected port is shutdown after the specified number of PDU has been received (Default is 5). If a port is shutdown as the result of a detected loop, it must be manually returned to an active state by the network administrator unless auto enable is configured. SLPP only sends a PDU to VLANs specified in the transmitting list configured by the user.

Note:

A maximum of 20 VLANs is supported by SLPP.

Note:

If you enable SLPP in addition to STP, STP operation takes precedence and SLPP is used as a supplementary measure for loop detection.

Rear port mode

Rear port mode configures the VSP 7000 Series Fabric Interconnect (FI) ports on the rear of the chassis for use as 40 Gbps Avaya interface ports. You can use rear port mode to interconnect standalone VSP 7000 switches using the various FI cables. Due to the different

bandwidth support of the FI port top and bottom connectors, you must connect the top connector to the top connector of another unit, and connect the bottom connector to the bottom connector of another unit. Depending on the mode of operation either 7 or 8 interfaces are presented over the four FI ports on the rear of the switch.

Each VSP 7000 provides an FI-up and FI-down port pair on the rear of the chassis. In each FI port pair there is a top and bottom connector. The top connector provides up to three 40 Gbps ports, and the bottom connector provides one 40 Gbps port.

You can configure rear port mode as on (enabled) or off (disabled). The default is off (disabled).

Note:

When you enable or disable rear port mode, the switch automatically initiates a reboot.

Warning:

Enabling rear port mode and answering yes to the confirmation prompt results in a switch configuration reset equivalent to a partial default command.

When you enable rear port mode, the switch applies the following default settings to all FI ports on the rear of the chassis:

- The LACP hashing mode is set to *advance*.
- VLAN tagging for rear ports is set to *tagAll*.
- The LACP administration key is set to *4095*.
- The LACP operating mode for rear ports is set to *active*.
- The LACP rear ports time-out value is set to *short*.
- LACP for rear ports is set to *enable*.

Note:

Rear port mode allows the VSP 7000 Series to automatically aggregate multiple connections between adjacent units in the FI mesh without additional configuration.

LACP mode is disabled on a rear port if you remove the port from the default VLAN (VLAN 1) and the port is not configured for any other VLAN. You can activate LACP mode once you add the port to a VLAN.

Note:

Rear port mode does not support SMLT or IST in Release 10.2

The following table provides detailed information about the FI ports and rear port modes:

FI port	Rear Ports Mode	Bandwidth	Ports
UP 1 (top right)	standard	120 Gbps	34,35,36

FI port	Rear Ports Mode	Bandwidth	Ports
UP 2 (bottom right)	standard	40 Gbps	33
DOWN 1 (top left)	standard	120 Gbps	38,39,40
DOWN 2 (bottom left)	standard	40 Gbps	37
UP 1 (top right)	SPB	120 Gbps	34, 35, 36
UP 2 (bottom right)	SPB	40 Gbps	33
DOWN 1 (top left)	SPB	80 Gbps	38, 39
DOWN 2 (bottom left)	SPB	40 Gbps	37

When you enable rear port mode Shortest Path Bridging (SPB), the switch configures port 40 as loopback. Port 40 in looback mode is not accessible by applications, or displayed on the port list.

Note:

Standard rear port mode does not support a SPB configuration. You must enable SPB rear port mode to support a SPB configuration. Changing between rear port modes results in a reboot and a partial configuration reset.

For more information about FI mesh cabling, see *Avaya Virtual Services Platform 7000 Series Installation* (NN47202–300).

For more information about configuring rear port mode, see *Avaya Virtual Services Platform 7000 Series Getting Started* (NN47202–303).

Fiber Fabric Interconnect cables

Fiber Fabric Interconnect (FFI) cables are available for use on the VSP 7000 Series for extended Fabric Interconnect stacking or rear-port mode. 50 m and 100 m FFI cables are available. Each FFI cable requires two Fabric Interconnect Transceivers.

AL7018006–E5	VSP 7000 Fabric Interconnect Transceiver (two required for each Fiber Fabric Interconnect cable)
AL7018007–E5	VSP 7000 Fiber Fabric Interconnect cable — 50 m
AL7018008–E5	VSP 7000 Fiber Fabric Interconnect cable — 100 m

A minimum of two FFI cables and four FI transceivers are required to establish a virtual fabric to an adjacent unit. For more information about FI stacking and cabling, see *Avaya Virtual Services Platform 7000 Series Installation* (NN47202–300).

Egress queue shaping

With egress queue shaping, you can specify the maximum and minimum egress shaping rates on an individual port and queue basis, for any or all egress queues associated with a switch port. The QoS agent egress queue set value determines the number of egress queues available for a port.

You can use QoS egress queue shaping to configure egress shaping for each queue without traffic interruption.

VSP 7000 Series switches support a maximum of 8 unicast queues and a maximum of 4 multicast queues. Known unicast packets are routed through unicast queues and unknown unicast, multicast, and broadcast packets are routed through multicast queues.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Quality of Service* (NN47202–504).

PFC-lite

PFC-lite is a simplified version of Priority Flow Control (PFC) that supports the Avaya VSP 7000 Series. PFC is a lossless QoS mode that provides a mechanism to stop ingress traffic of a given packet priority. Full PFC is described in IEEE 802.1Qbb. You can enable PFC-lite by configuring the QoS agent buffer to operate in lossless-pfc mode.

Important:

PFC-lite is a technology demonstration feature only. Lossless-pfc mode is not fully supported in Release 10.2.0.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Quality of Service* (NN47202–504).

CANA

With Custom Autonegotiation Advertisements (CANA), you can control the speed and duplex setting information that each switch port advertises as part of the autonegotiation process.

Without CANA, a port with autonegotiation enabled advertises all speed and duplex modes that are supported by the switch, and attempts to establish a link at the highest common speed and duplex setting. With CANA, you can configure the port can be configured to advertise only specific speed and duplex settings, thereby allowing links to be established only at these settings, regardless of the highest common supported operating mode.

CANA also provides control of the IEEE802.3x flow control settings advertised by switch ports, as part of the autonegotiation process

For more information, see *Avaya Virtual Services Platform 7000 Series Getting Started* (NN47202–303).

Out-of-band management

Out-of-band management allows IPv4 switch or stack management using the dedicated out-of-band management port. Out-of-band management supports Telnet, Secure Shell (SSH) protocol, Simple Network Management Protocol (SNMP), HTTP, or HTTPS, without requiring an in-band management VLAN.

To configure out-of-band management, you assign an IPv4 address to the RJ-45 Ethernet management port for a switch or stack. You can configure a specific out-of-band management default gateway, which takes precedence over the in-band default gateway. If you do not configure an out-of-band management default gateway, the in-band default gateway is used for out-of-band switch or stack management.

Note:

The out-of-band switch or stack management IPv4 address must be different than the in-band IPv4 address and belong to a different subnet.

You can use the out-of-band management port to perform tasks such as downloading software images and, when the SNMP server is enabled, access the Enterprise Device Manager (EDM) interface for a switch or stack. To access EDM, you type the out-of-band management IPv4 address in the address bar of an Internet browser.

The out-of-band management port supports full auto negotiation, which enables management stations to connect at any of the supported speeds or duplexes.

For more information, considerations, and limitations, see *Avaya Virtual Services Platform 7000 Series Getting Started* (NN47202–303).

Feature licenses

This release introduces new features that require a license. You purchase switches and licenses separately.

To enable certain features, Avaya VSP 7000 Series supports:

- Advanced License
- Premier License
- Trial License

Advanced Licenses support the following features:

- Split MultiLink Trunking (SMLT)
- Open Shortest Path First (OSPF)
- Virtual Router Redundancy Protocol (VRRP)
- Equal Cost Multi Path (ECMP)

Premier Licenses support all Advanced features plus:

- Shortest Path Bridging MAC (SPBM)

To enable License features on an Avaya Virtual Services Platform 7000 Series, you must:

- purchase a License Kit
- generate a license file on the Avaya electronic licensing portal
- install a license file on the switch

Caution:

If you reset a standalone device to the default configuration, you erase the license file.

For more information, see *Avaya Virtual Services Platform 7000 Series Fundamentals* (NN47202–101).

IPFIX

IP Flow Information Export (IPFIX) is a protocol you can use to export flow information from traffic observed on a switch. IPFIX implementation is based on Netflow Version 9.

Note:

IPFIX can also monitor IGMP traffic.

For the VSP 7000 Series, IPFIX supports the following external IPFIX collectors:

- NetQoS Harvester/Collector
- Avaya IP Netflow Version 9
- Avaya IP Flow Manager
- Fluke Collector

Note:

You cannot use IPFIX for secondary interfaces.

IPFIX shares resources with QoS. If you enable IPFIX, a QoS policy precedence is used.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - System Monitoring* (NN47202–505).

Trace

The Trace feature provides troubleshooting commands that you can use to gather information about errors and events occurring in real time. Trace supports OSPF, IGMP, RIP, SMLT, and NTP. You can configure Trace to display various levels of detail on the console, from very little information (Level 1, VERY TERSE) to very complete information (Level 4, VERY VERBOSE).

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - System Monitoring* (NN47202–505).

Change RADIUS password

Remote users can change their account passwords when RADIUS is configured and enabled in your network.

Once you configure RADIUS servers in your network to provide centralized authentication, authorization and accounting for network access, you can enable the MS-CHAPv2 encapsulation method, which permits the changing of the RADIUS password for user accounts.

Note:

Change RADIUS password is available only in secure software builds.

Change RADIUS password is disabled by default.

If you enable RADIUS encapsulation ms-chap-v2, when an account password expires the RADIUS server reports the expiry during the next logon attempt and the system prompts you to create a new password. Also, you can change your RADIUS password before expiry using ACLI.

To use change RADIUS password you must have:

- at least one configured and reachable RADIUS server in your network
- configured RADIUS encapsulation ms-chap-v2

Change RADIUS password is compatible with RADIUS password fallback.

Effects of software upgrade on RADIUS settings:

The system saves all RADIUS settings in NVRAM. If you upgrade the switch software, the new image loads the settings.

Effects of software downgrade on RADIUS settings:

- if change RADIUS password is available on the release it will be set to default.
- if change RADIUS password is not available in the release it will not be present in the image.

Configuration for the change RADIUS password feature save in both the binary and ASCII configuration files.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Security* (NN47202–501).

RADIUS use management IP

When the switch is operating in Layer 3 mode, by default, a RADIUS request uses one of the routing IP addresses on the switch. In Layer 3 mode, the RADIUS use management IP configuration ensures that the switch or stack generates RADIUS requests using the source IP address of the management VLAN. In some networks, the source IP in the RADIUS request is used to track management access to the switch.

When the switch is operating in Layer 2 mode, by default, all RADIUS requests generated by the switch use the stack or switch management IP address as the source address in RADIUS requests or status reports. The RADIUS use management IP configuration has no impact when the switch operates in Layer 2 mode.

Note:

If the management VLAN is not operational, the switch cannot send any RADIUS requests when:

- the switch is operating in Layer 2 mode.
- the switch is operating in Layer 3 (routing) and RADIUS use management IP is enabled.

This is normal behavior. In Layer 2 mode, if the management VLAN is unavailable, there is no active management IP instance. In Layer 3 mode, if RADIUS use management IP is enabled, then the switch does not use any of the other routing instances to send RADIUS requests when the management VLAN is inactive or disabled.

DHCP snooping

When you use Dynamic Host Configuration Protocol (DHCP) servers in a network to allocate IP addresses to network clients, you can configure DHCP snooping on network switches to allow only clients with specific IP or MAC addresses to access the network.

With DHCP snooping, you can prevent attackers from responding to requests from DHCP servers with false IP or MAC information. DHCP snooping defends against this type of attack,

known as DHCP spoofing, by performing as a firewall between untrusted hosts and the DHCP servers.

You must configure DHCP snooping individually for each VLAN. You can configure DHCP snooping using ACLI, EDM, or SNMP

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Security* (NN47202–501).

DHCP snooping binding table

When enabled, DHCP snooping dynamically creates and maintains a binding table. The DHCP snooping binding table includes the following information about DHCP leases on untrusted interfaces:

- source MAC address
- IP address
- lease duration
- time until expiration of current entry
- VLAN ID
- port
- source information that can be learned or is static

The maximum size of the DHCP snooping binding table is 1024 entries.

The DHCP snooping binding table is used by IP Source Guard to filter traffic. If the sending station is not in the binding table, no IP traffic is allowed to pass. When a connecting client receives a valid IP address from the DHCP server, IP Source Guard installs a filter on the port to allow only traffic from the assigned IP address.

Static DHCP binding table entries

You can manually add static entries in the DHCP binding table to protect IP devices using applications such as DAI and IPSG, that rely on DHCP snooping table entries. When the protection of these statically configured IP devices is no longer required, you can manually delete entries from the DHCP binding table.

Static DHCP binding table entries with infinite expiry time are stored in NVRAM and are saved across restarts.

DHCP snooping external save

You can use DHCP snooping external save to store the DHCP snooping database at predefined, 5 minute intervals, to an external TFTP server or USB drive.

You must enable SNTP/ NTP and synchronization. The lease expiry time the switch writes to the externally saved DHCP snooping database is the absolute lease expiry time, which can be accurately restored from the external save when you reboot the switch.

Important:

Any DHCP snooping database entries that you manually configure, or that the switch learns between the time of the last initiated external save and the beginning of the reboot process are lost and not available when the switch is again operational.

DHCP snooping option 82

When you enable DHCP snooping Option 82, the switch can transmit information about the DHCP client to the DHCP server. The server can use the information from the switch to locate the DHCP client in the network and allocate a specific IP address to the DHCP client.

DHCP snooping Option 82 function is controlled by the one switch at the edge of a network and not by any switches located between the network edge switch and the DHCP server.

DHCP snooping Option 82 cannot function independently from DHCP snooping (Layer 2 mode).

Dynamic ARP inspection

Dynamic Address Resolution Protocol (ARP) inspection is a security feature that inspects and validates ARP packets in a network.

Dynamic ARP inspection can prevent attacks by intercepting, logging, and discarding ARP packets with invalid IP-to-MAC address bindings. Without dynamic ARP inspection, a malicious user can attack hosts, switches, and routers connected to the Layer 2 network by poisoning the ARP caches of systems connected to the subnet, and by intercepting traffic intended for other hosts on the subnet.

The address binding table is dynamically built from information gathered in the DHCP request and reply when DHCP snooping is enabled. The MAC address from the DHCP request is paired with the IP address from the DHCP reply to create an entry in the DHCP binding table.

When you enable dynamic ARP inspection, ARP packets on untrusted ports are filtered based on the source MAC and IP addresses seen on the switch port. The switch forwards an ARP packet when the source MAC and IP address matches an entry in the address binding table. ARP packets for which the source MAC and IP address do not match an entry in the address binding table are dropped.

The function of dynamic ARP inspection is dependant on DHCP snooping being globally enabled.

Dynamic ARP inspection is configured on a VLAN to VLAN basis.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Security* (NN47202–501).

IP Source Guard

IP Source Guard (IPSG) is a Layer 2 port-to-port basis feature that provides security to a network by using information in the DHCP snooping binding table to filter clients with invalid IP addresses. When you enable IPSG on an untrusted port with DHCP snooping enabled, an IP filter entry is created or deleted for that port automatically, based on IP information stored in the corresponding DHCP binding table entry. When a connecting client receives a valid IP address from the DHCP server, a filter is installed on the port to allow traffic only from the assigned IP address. A maximum of 10 IP addresses are allowed on each IPSG-enabled port. When this number is reached, no more filters are set up and traffic is dropped.

Important:

Enable IPSG only on an untrusted DHCP snooping port.

Important:

Avaya recommends that you do not enable IPSG on MLT, DMLT and LAG ports.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Security* (NN47202–501).

Unicast storm control

To control unicast traffic congestion on the switch you can use unicast storm control. Unicast storm control blocks all known and unknown unicast traffic when traffic rates exceed a user-configured threshold, also known as the high water mark.

You can also configure the rate at which traffic can resume, also known as the low water mark. Once the unicast traffic rate drops below the low water mark, all unicast traffic can pass through the switch

Regardless of the blocking state of unicast traffic, all broadcast and multicast traffic continues to pass, or forward, unless it is blocked or limited by other means, such as broadcast rate limiting for example.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Security* (NN47202–501).

MAC address-based security

You can safeguard your Ethernet networks from unauthorized surveillance and intrusion with the Media Access Control (MAC) address-based security feature, a real-time security system based on Avaya local area network (LAN) access for Ethernet.

You can use MAC address-based security to set up network access control based on the source MAC addresses of authorized stations and to perform the following activities:

- Create a list of up to 10 MAC addresses to filter as:
  - destination addresses (DA)—the system drops all packets containing one of the specified MAC addresses, regardless of the ingress port, source address intrusion, or virtual local area network (VLAN) membership.
  - source addresses (SA)—the system drops all packets containing one of the specified MAC addresses.

  Important:

  Do not use the MAC address for the stack or units in the stack.
- Create a list of up to 448 MAC source addresses.
  - Specify source addresses authorized to connect to the switch or stack.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Security* (NN47202–501).

MAC security autolearning

You can use the MAC address-based security autolearning feature to automatically add allowed MAC addresses to the MAC Security Address table.

When you use MAC address-based security autolearning you can:

- Specify the number of addresses that can be learned on the ports, to a maximum of 25 addresses for each port. The switch forwards only traffic for those MAC addresses statically associated with a port or those learned with the autolearning process.
- Configure an aging time, in minutes, after which the system refreshes autolearned MAC address entries in the MAC Security Address Table. If you set the aging time value to 0, the entries never age out. To force relearning of entries in the MAC Security Address Table you must reset learning for the port.
- Reset the MAC address table for a port by disabling the security on the port and then enabling it.

MAC security port lockout

You can use MAC security port lockout to exclude specified ports from participating in MAC-based security. MAC security port lockout can simplify switch operation and provide protection against improper configurations. You can lock out:

- uplink ports
- MLT ports
- remote-administration ports

MAC security port lockout can prevent accidental loss of network connectivity caused by improper MAC security settings.

Sticky MAC address

Sticky MAC address provides a high level of control, and simpler configuration and operation for MAC address security.

You can use the Sticky MAC address feature on either a standalone switch or a stack unit.

With Sticky MAC address, you can secure the MAC address to a specified port, so that if the MAC address moves to another port, the system raises an intrusion event.

When you enable Sticky MAC address, the switch performs the initial autolearning of MAC addresses and can store the automatically learned addresses across switch reboots.

MAC Flush

You can use the MAC Flush feature to clear MAC Address entries directly from the MAC Address Table or Forwarding Database. For dynamically learned addresses, if you do not use the MAC Flush feature commands, you can use the following indirect methods:

- power cycling the switch
- deleting, and then recreating the VLAN
- unplugging, and then replugging the port to flush out all addresses learned on the port

MAC Flush provides the following options to flush out MAC address entries:

- clear a single MAC address
- clear all MAC addresses from a port (or list of ports)
- clear all MAC addresses from a trunk (MLT or LAG)
- clear all MAC addresses from a particular VLAN
- clear only dynamic or only static MAC addresses from a port
- clear only dynamic or only static MAC addresses from a VLAN
- clear only dynamic or only static MAC addresses from a trunk
- clear all static MAC addresses
- clear all dynamic MAC addresses
- clear all MAC addresses

MAC Flush clears only dynamically learned or statically entered MAC Addresses. MAC Flush does not delete MAC Addresses created by MAC Security or Port Mirroring because deletion of these MAC Addresses can affect the MAC Security or Port Mirroring function.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Layer 2* (NN47202–502).

FDB Disable MAC Learning

You can use Forwarding Database (FDB) Disable MAC Learning to disable MAC address learning on specific ports.

Disabling MAC learning is useful in situations where you want to control the Layer 2 FDB entries. For example, when you deploy the switch in metro environments, nodes on the network might flood traffic, and the MAC tables can fill rapidly. Disabling MAC learning gives you control to prevent the MAC tables from filling unnecessarily in a situation like this.

You cannot control the learning behavior for ports per VLAN due to hardware limitations.

You use FDB Disable MAC Learning in combination with the FDB Static MAC Entry feature when you want to add a MAC address in the MAC address table, by adding it statically using FDB Static MAC Entry. FDB Disable MAC Learning interacts with the Layer 2 application, MAC Security.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Layer 2* (NN47202–502).

FDB Static MAC Entry

The Forwarding Database (FDB) contains information that maps the MAC address of each known device to the switch port where the device address was learned.

When you use the FDB Static MAC Entry feature you can configure static MAC address entries in the FDB, or MAC address table. Once you configure a static MAC address entry in the FDB, the static MAC address does not age out like a dynamically learned address. A static address from the FDB is a unicast address and the system does not erase it after switch resets or when link-down events occur.

You can configure up to 1,024 static MAC addresses in the FDB. FDB Static MAC Entry works in conjunction with the FDB Disable MAC Learning feature.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Layer 2* (NN47202–502).

IP routing

The VSP 7000 Series supports IP routing. To configure IP routing, you must create a virtual router interface by assigning an IP address to a virtual local area network (VLAN). You can create a virtual router interface for a specified VLAN by associating an IP address with the VLAN. The VSP 7000 Series supports wire-speed IP routing between VLANs.

Because a physical switch port can be associated with multiple routable VLANs, a virtual router interface is not associated with any specific switch port. Network traffic can reach the VLAN IP address through any of the VLAN port members. The assigned IP address also serves as the gateway through which packets are routed out of that VLAN. Routed traffic can be forwarded to another VLAN within a switch or stack.

When you configure a switch to route IP traffic between different VLANs, that switch operates in Layer 3 mode; otherwise, the switch operates in Layer 2 mode. When you assign an IP address to a Layer 2 VLAN, that VLAN becomes a routable, Layer 3 VLAN. You can assign a unique IP address to each VLAN.

You can configure all IP routing parameters on a switch before you globally enable or disable IP routing for that switch. IP routing is globally disabled by default.

The VSP 7000 Series supports the following IP route formats:

- local routes: the switch automatically creates routes to each of the local Layer 3 VLAN interfaces.
- static routes: you must manually enter the routes to the destination IP addresses.
- dynamic routes: are identified using a routing protocol such as RIP or OSPF.

For more information about IP routing and the following Layer 3 features, see *Avaya Virtual Services Platform 7000 Series Configuration- IP Routing* (NN47202–511).

Switch management with IP routing

When IP routing is enabled for a switch or stack, you can use any of the virtual router IP addresses for switch or stack management over IP. Any routable Layer 3 VLAN can carry the management traffic, including Telnet, Web, Simple Network Management Protocol (SNMP), BootP and Trivial File Transfer Protocol (TFTP).

You can also use the management VLAN for switch or stack management over IP. If IP routing is not enabled for a switch or stack, the management VLAN is reachable only through the switch or stack IP address, and only through ports that are members of the management VLAN. When you enable IP routing for a switch or stack, the management VLAN behaves similar to other routable VLANs. The IP address is reachable through any virtual router interface, as long as a route is available.

You can configure a management route from the management VLAN to a particular subnet. The management route is a static route that allows incoming management connections from the remote network to the management VLAN.

Brouter port

A brouter port is a single-port VLAN that can route IP packets as well as bridge all non-routable traffic. The difference between a brouter port and a standard IP protocol-based VLAN configured for routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for non-routable traffic and still be able to route IP traffic. This feature removes any interruptions caused by Spanning Tree Protocol recalculations in routed traffic. Because a brouter port is a one-port VLAN, each brouter port decreases the number of available VLANs by one, and uses one VLAN ID.

CLIP

Circuitless IP (CLIP) is a virtual IP (VIP), or loopback interface that provides a method of assigning one or more IP addresses to a routing switch, without the requirement of binding the IP address to a physical interface.

Because the IP address assigned to a CLIP interface does not map to a specific physical interface, if one or more physical IP interfaces on a routing switch fails, the CLIP interface ensures connectivity if an actual path is available to reach the device.

The system treats a CLIP interface the same as any IP interface. The network associated with a CLIP is treated as a local network connected to the switch and is always reachable through a VLAN interface. This route always exists and the circuit is always up because there is no physical attachment.

Note:

CLIP interfaces are disabled by default on VSP 7000 Series devices.

DHCP relay

The VSP 7000 Series supports DHCP relay, which forwards the DHCP broadcasts to the IP address of the DHCP server. Network managers prefer to configure a small number of DHCP servers in a central location to lower administrative overhead. Routers must support DHCP relay so that hosts can access configuration information from servers several router hops away.

When you enable DHCP relay, the switch can relay client requests to DHCP servers located on different Layer 3 VLANs or in remote networks, and can also relay server replies back to clients.

To relay DHCP messages, you must create two Layer 3 VLANs. One VLAN is connected to the client and the other provides a path to the DHCP server. You enable DHCP relay on each VLAN individually.

Important:

DHCP relay uses a hardware resource that is shared by switch Quality of Service (QoS) applications. When you enable DHCP relay globally, the QoS filter manager cannot use precedence 8 for configurations. For the filter manager to be able to use this resource, you must disable DHCP relay for the switch unit or entire stack.

UDP broadcast forwarding

By default, User Datagram Protocol (UDP) broadcast frames received on one VLAN are not routed to another VLAN. Because some network applications, such as the NetBIOS name service, rely on UDP broadcasts to request a service or locate a server, the VSP 7000 Series supports UDP broadcast forwarding, which forwards the broadcasts to the server through a Layer 3 VLAN interface.

UDP broadcast forwarding selectively forwards limited UDP broadcasts received on an IP interface to a configured server IP address, as a unicast packet.

Important:

UDP broadcast forwarding shares resources with Quality of Service (QoS). When UDP forwarding is enabled, the switch dynamically assigns the highest available precedence value to the UDP broadcast forwarding feature. You can view the assigned precedence value after you enable UDP forwarding by displaying QoS diagnostics entries on the switch.

IGMPv3 snooping and IGMP Querier

IGMPv3 provides the ability to reduce the amount of network traffic by packing multiple group members in a single report message. Also, IGMPv3 allows a host to include or exclude a list of source IP addresses for each multicast group of which the host is a member. Routers merge the source IP address requirements of different hosts for each group.

The VSP 7000 Series switch supports IGMPv3 source filtering capability with IGMPv3 snooping. IGMPv3 snooping maintains backward compatible with IGMPv1 and IGMPv2.

IGMPv3 Snooping operates independently of PIM-SSM.

With the VSP 7000 Series, IGMPv3 snooping-enabled interfaces process IGMP reports as follows:

- process IGMPv3 reports independent of PIM-SSM
- process all six IGMPv3 group record types
- process all IGMPv3 source information
- backward compatible with IGMPv1 and IGMPv2 reports

The IGMP Querier functionality allows a switch or stack to be configured as an active query router without the need for dedicating a standalone switch in each network to the task. Avaya recommends that each VLAN using IGMP multicast have a router performing multicast queries. This router typically has PIM or DVMRP enabled.

A multicast query router communicates with hosts on a local network by sending IGMP queries. The multicast router periodically sends a general query message to each local network of the router. This process is standard multicast behavior.

OSPF

Open Shortest Path First (OSPF) is a classless Interior Gateway Protocol (IGP) that distributes routing information among switches that belong to a single autonomous system (AS). An OSPF AS can be defined as a group of switches in a network that run OSPF and operate under the same administration. Intended for use in large networks, OSPF is a link-state protocol that supports variable length subnet masking (VLSM) and tagging of externally-derived routing information.

Important:

The VSP 7000 Series implementation of OSPF supports only broadcast and passive interfaces. Point-to-point and Non-Broadcast Multi-Access (NBMA) interfaces are not supported.

RIP

Routing Information Protocol (RIP) is a standards-based, dynamic routing protocol based on the Bellman-Ford (distance–vector) algorithm. RIP functions as an Interior Gateway Protocol

(IGP) and allows switches to exchange information for computing the shortest routes through an IPv4-based network.

RIP-enabled switches maintain routing tables that list the optimal routes to every destination device in the network. Each RIP-enabled switch advertises its routing information throughout the network at regular intervals. Network neighbor devices use the information advertised by the switches to recalculate routing tables, and then retransmit the routing information.

RIP uses User Datagram Protocol (UDP) data packets to exchange routing information.

VRRP

Virtual Router Redundancy Protocol (VRRP) allows the use of a virtual IP address (transparent to users) shared between two or more routers connecting a common subnet to the enterprise network. With end hosts using the virtual IP address as the default gateway, VRRP provides dynamic default gateway redundancy in the event of failure. You can use the VRRP, defined in RFC 3768, to eliminate the single point of failure that can occur when the single static default gateway router for an end station is lost.

Once you initialize a VRRP router, if there are no other VRRP routers enabled within the VLAN, the initialized router assumes the role of the VRRP master router. When you enable additional VRRP routers within the VLAN, an election process takes place among the routers to elect a master router, based on their priority.

ECMP

With Equal Cost Multi Path (ECMP), you can configure the VSP 7000 Series to use up to four equal cost paths to the same destination prefix. The Layer 3 (L3) switch can use multiple paths for traffic load sharing and in the event of network failure, achieve faster convergence to other active paths. When the L3 switch maximizes load sharing among equal cost paths, the system uses links more efficiently for IP traffic transmission. ECMP supports OSPF, RIP, and static routes.

ARP

With the Address Resolution Protocol (ARP) you can configure the VSP 7000 Series to dynamically learn Layer 2 Media Access Control (MAC) addresses, and to build an ARP cache table with corresponding Layer 3 IP addresses.

For network switches using the IP protocol to transmit packets to a network host, the network switch requires both a MAC (physical) address and an IP address for the network host. If a network switch knows only the IP address of a network host , the switch can use ARP to determine a network host MAC address and bind the 32-bit IP address to the 48-bit MAC address. A network switch can use ARP across a single network only, and the network hardware must support physical broadcasts.

IGMP Profiles

IGMP Profiles, also referred to as IGMP Selective Channel Block, gives you the control to block the streaming of specific channels on some ports.

In certain deployment scenarios, you might prefer to disallow the multicast streaming from specific group addresses to users on specific ports. With IGMP Profiles, you can configure the

IGMP membership of ports by selectively channel blocking IGMP reports received from users on that port, destined for the specific group address or addresses. The profile can be configured to block a single multicast address or range of addresses.

IGMP Profiles work regardless of whether the switch is in Layer 2 IGMP snooping mode or the full IGMP mode. The blocking of channels is implemented by blocking the ports from joining an IGMP group, and is applicable for IGMP v1, v2 and v3.

You can configure up to 240 channels for blocking.

You cannot use IGMP Profiles to snoop the multicast streams that are sent from a group to a port.

You can use IGMP Profiles for both MLT and LACP trunk interfaces. However, you cannot apply a profile directly to an MLT or LACP trunk. You must apply an IGMP profile to a member of the trunk.

When you apply an IGMP profile to a port which is a member of an MLT or LACP trunk, the system applies the profile to all ports of the MLT or LACP. When you dynamically add or remove a port from a MLT or LACP with a IGMP profile associated with it, the system adds or removes all ports from the IGMP profile.

You can use IGMP Profiles in standalone or stacking mode. In stacking mode, the configuration propagates from any unit to all other units in the stack.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration IP Routing* (NN47202–511).

IGMP Multicast Flood Control

The IGMP Multicast Flood Control feature can minimize IP multicast flooding in the network. IGMP Multicast Flood Control is also referred to as IGMP multicast filter mode. You can enable IGMP multicast filter mode to globally limit IP multicast traffic.

IGMP Multicast Flood Control can limit IP multicast traffic without affecting other control protocols. IGMP Multicast Flood Control is an enhancement over IGMP unknown multicast filtering, since using IGMP unknown multicast filtering inhibits other control protocols such as OSPF, RIP, VRRP, and IPv6 traffic.

IGMP Multicast Flood Control also eliminates the need to send queries when IGMP Snooping is enabled, which reduces problems in networks with active queriers. If you enable IGMP Multicast Flood Control, the IGMP multicast filter mode applies to all VLANs.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration IP Routing* (NN47202–511).

TDR

Time Domain Reflectometer (TDR) provides a diagnostic capability to test connected cables for defects (such as a short pin or pin open). TDR tests only apply to Ethernet copper ports, fiber ports cannot be tested. You can test one or multiple Ethernet copper ports at the same time using ACLI or EDM.

TDR tests run at 1 Gbps. When you test a cable with the TDR, the port speed changes to 1 Gbps during the test and is restored to the previous operating speed when the test is complete. If the cable is running a 100 Mbps link, the test results might be incomplete since TDR does not test all the pins on the connector. TDR tests can run simultaneous with currently operating 1 Gbps links.

Note:

The accuracy margin of cable length diagnosis is between three to five meters. TDR does not support skew, swap and polarity tests.

Other changes

See the following sections for information about changes and updates to existing information:

MDA warm swap

Media Dependent Adapter (MDA) warm swap is supported. You can warm insert, warm remove, or hot remove an MDA.

You can perform a warm insert by enabling the MDA once it is inserted in a unit. You must enable an MDA for the MDA ports to function. An MDA insertion event does not interrupt traffic within the switch or stack.

You can perform a warm removal by disabling the MDA before removal. A warm removal ensures MDA egress traffic routes to an operating switch port. Hot removal is also supported, however a hot removal can cause packet loss. Once an MDA is hot removed, packets scheduled to egress the MDA ports are discarded, and an unsafe MDA removal log message is generated.

For more information, see *Avaya Virtual Services Platform 7000 Series Getting Started* (NN47202–303).

Validated third party devices

Avaya recommends you to use Avaya SFP and SFP+ devices to provide maximum compatibility and support for the VSP 7000 Series. The following third party SFP+ devices are validated by Avaya to function with the VSP 7000 Series.

Vendor	Model	Details
Cisco	450–16141	Cisco SFP+ direct attach cable — 5 m
Cisco	450–16140	Cisco SFP+ direct attach cable — 3 m
Cisco	450–16142	Cisco SFP+ direct attach cable — 1 m
HP	J9281B	HP SFP+ direct atttach cable — 1 m
HP	J9283B	HP SFP+ direct attach cable — 3 m
HP	J9285B	HP SFP+ direct attach cable — 7 m
HP	J9286B	HP SFP+ direct attach copper cable — 10 m

If you use a verified third party device, you can use the `show interface gbic` command to display the vendor SEPROM data and the cable type.

Note:

Using Avaya devices is recommended to provide maximum compatibility and support for the VSP 7000 Series.

AUR license enhancement

Automatic Unit Replacement (AUR) is updated to enable the automatic update of a license for any replacement stack unit, including the base unit.

For more information, see *Avaya Virtual Services Platform 7000 Series Getting Started* (NN47202–303).

Industry standard CLI

Avaya Command Line Interface (ACLI) commands for Address Resolution Protocol (ARP), Spanning Tree Protocol (STP), and Virtual Local Area Network (VLAN) are modified to use industry standard CLI command syntax.

For more information, see *Avaya Virtual Services Platform 7000 Series Getting Started* (NN47202–303).

Dynamic change of management IP address

You can dynamically change the in-band management IP address using Telnet, SSH, SNMP, HTTP, and HTTPS.

Also, with the introduction of out-of-band switch or stack management, IP management sections now indicate in-band or out-of-band switch or stack management.

For more information, see *Avaya Virtual Services Platform 7000 Series Getting Started* (NN47202–303).

Disabled ACLI audit and remote ACLI command logging

You can disable the recording and logging of all ACLI commands issued on the switch (ACLI Audit). If the switch is configured for remote logging, issued ACLI commands are logged on the remote syslog server, even though they are not logged in switch NVRAM.

SNMP trap support and enhancements

SNMP trap enhancements provide the option to enable or disable notifications for objects on specific interfaces, as well as globally. Most of the notifications are enabled by default, while some are disabled by default as their application is not enabled. You can modify them according to your requirements, using SNMP trap notification control, globally and per interface.

You can configure traps using the SNMPv1 or SNMPv2c or SNMPv3 format. If you do not identify the SNMP version, the system formats the traps in the SNMPv1 format. A community string can be entered if the system requires one.

The switch supports both industry-standard SNMP traps, as well as private Avaya enterprise traps.

You can enable or disable SNMP traps for the following features:

- Rapid Spanning Tree Protocol (RSTP)
- DHCP Snooping
- Dynamic ARP Inspection (DAI)
- IP Source Guard (IPSG)
- Unicast Storm Control
- Split Multi-Link Trunk (SMLT)
- Open Shortest Path First (OSPF)

- MAC Security
- Virtual Router Redundancy Protocol (VRRP)

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Security* (NN47202–501).

MLT/DMLT/LAG dynamic VLAN changes

Enhancements have been made to Link Aggregation Groups (LAG) to provide consistent operation of Multi-Link Trunk (MLT), Distributed Multi-Link Trunk (DMLT), and LAGs so that you can make VLAN changes on trunks without disabling the trunk first.

The switch allows you to move a LAG member into a VLAN and all ports that have LACP enabled with the same LACP key are moved. This behavior is similar to MLT and DMLT.

If you attempt to remove all VLANs from an active MLT/DMLT/LAG, the system outputs a message warning you of possible loss of connectivity to the switch, and requests a confirmation to continue. If you remove all MLT/DMLT/LAG ports from all VLANs, the trunk is disabled.

When you add a port to a new STG, you can use STG port membership in auto mode, so that STP is automatically enabled on that port to prevent loops.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration - Layer 2* (NN47202–502).

Show flash

You can use the **`show flash`** ACLI command to display flash memory capacity and current flash allocation, and file usage information about dual images and backup configurations. The command output displays the actual file sizes and space allocated to them.

For more information, see *Avaya Virtual Services Platform 7000 Series Getting Started* (NN47202–303).

Flow Control

You can configure the Flow Control for ports to symmetric when the switch is operating in Lossless mode.

Use the following ACLI interface configuration command for 10 Gb ports that do not support autonegotiation: **`flowcontrol symmetric`**

Use the following ACLI interface configuration command for 1 Gb ports with autonegotiation: **`auto-negotiation-advertisements 1000-full pause-frame`**

Note:

Flow Control symmetric port mode is only supported in Lossless QoS buffer mode. You can configure Flow Control as disabled or asymmetric when operating in a non-lossless mode.

For more information, see *Avaya Virtual Services Platform 7000 Series Configuration Quality of Service* (NN47202–504).

Chapter 3: Important notices

This section provides important software and hardware related notices.

Warnings and important notices

The following sections provides warning notices and important notices for the VSP 7000 Series.

Agent upgrade notice

Warning:

DATA LOSS CAN OCCUR — Do not upgrade directly from Release 10.0 to Release 10.2 or later.

If the switch is running Release 10.0, you must upgrade to Release 10.1 before upgrading to 10.2 or later software. Upgrading from Release 10.0 to the current release can cause accidental erasure of the agent image on the switch. If the primary agent image is erased, during the next reboot the switch attempts to boot from the secondary agent image.

Fabric Interconnect cables notice

Important:

You must orient each cable so that the alignment slot on the FI cable connector is correctly aligned with the switch. The FI cable alignment slot must be facing upwards. For more information, see the following figures.

Caution:

Risk of equipment damage

Incorrect FI cable insertion can cause physical damage to the VSP 7000 Series switch. For more information, see the following figures.

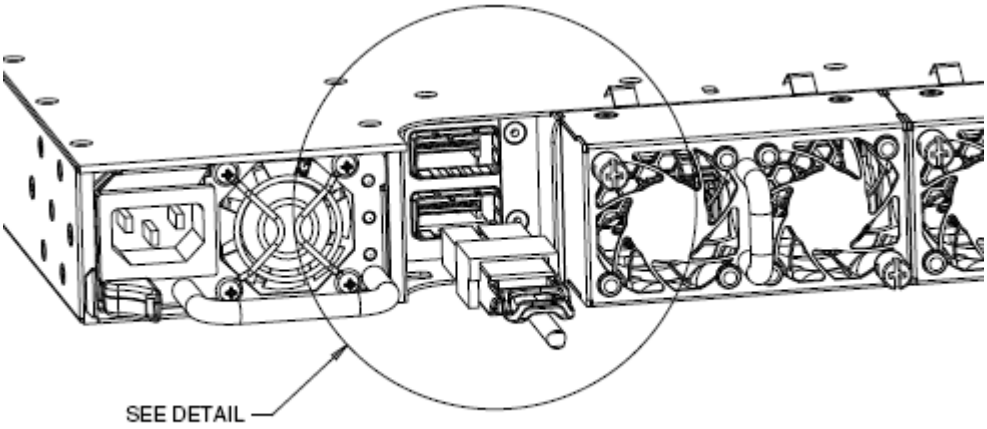

Figure 1: Installing Fabric Interconnect cables

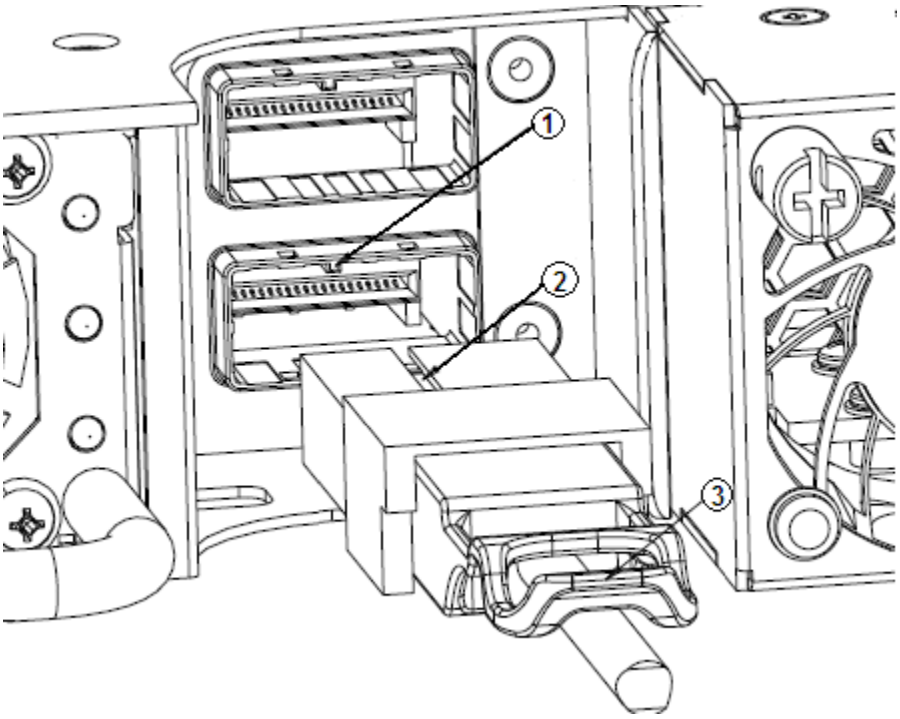

Callout	Description
1	FI port alignment tab
2	FI cable alignment slot (Insert cable with slot facing UP and aligned with tab on the port)

Callout	Description
3	FI cable connector pull tab (Ensure that the connector pull tab is facing UP)

Figure 2: Installing Fabric Interconnect cables detail

Important:

Remove the FI cables before changing between stacking and rear-port modes, or before fully defaulting a switch to avoid network loops.

Fiber Fabric Interconnect notice

Important:

Rear port mode is not supported over Fiber Fabric Interconnect (FFI) cables in the initial standard Release 10.2.0.

FFI cables in a FI mesh configuration might fail to establish a link if one end of the FFI cable connection is reset or power cycled.

You must physically remove and reinsert the FFI cable to resolve the link. A release patch is being developed to resolve this issue.

Note:

Rear port mode supports copper FI cables.

FI stacking supports FFI cables.

Media Dependent Adaptor notice

Important:

Inserting the MDA might require a larger than anticipated amount of force to fully seat the MDA into the MDA slot. To ensure that the MDA is fully inserted, securely install the switch chassis in an equipment rack before installing the MDA.

Caution:

Risk of equipment damage

If the MDA is not fully seated, do not use the thumb screws in an attempt to pull in the MDA. This can deform the front metal surround of the MDA.

Note:

The VSP–7008XT- MDA only supports full duplex mode of operation. Half duplex is not supported on 10GBASE-T ports.

Power Supplies recommendation

Avaya recommends the installation of two VSP 7000 power supplies to ensure minimum disruptions due to power outages.

File names for this release

The following table describes the Avaya Virtual Services Platform 7000 Series Release 10.2 software files. File sizes are approximate.

Module or file type	Description	File name	File size (bytes)
Standard runtime image for Release 10.2.	Agent software image for the Avaya Virtual Services Platform 7000 Series	7000_102006.img — non-secure image 7000_102007s.img — secure image	11129508 11371940
Diagnostics software for Release 10.2.	Diagnostics software for the Avaya Virtual Services Platform 7000 Series	7000_10106_diags.bin — diagnostics	3866700
MIB definition files for Release 10.2.	MIB definition files	Virtual_Services_Platform_70xx_MIBs_10.2.0..zip	1213448
EDM Help files for Release 10.2.	EDM help files	vsp7000v1020_HELP_EDM.zip	4548442

Software feature license file information

Features not specifically included in a particular release, those not fully tested and deemed technology features, are available under a demonstration license feature. The demonstration license provides a 30 day trial for a technology feature.

Once a technology feature becomes a licensed feature, you can obtain full use of the feature with an Advanced or Premier license.

When you create a license file to enable licensed features on an Avaya Virtual Services Platform 7000 Series switch with the Avaya Electronic Licensing Portal, you must specify a file name.

Follow the instructions on the License Certificate within the License Kit.

Rules for generating a license file:

You must apply the following rules when you generate and name the license file:

- A maximum of 63 alphanumeric characters are permitted.
- Only lower case letters are permitted.

- Spaces or special characters are not permitted.
- Underscore (_) is permitted.
- The dot (.) and a three-character file extension are required.

An example of a file name is **abcdefghijk_1234567890.lic**.

The format of the file that you upload to the license generation tool, and that contains the list of MAC addresses, must be as follows:

- ASCII file format.
- One MAC address per line.
- No other characters, spaces, or special characters are permitted.
- The MAC address must be in a hexadecimal, capitalized format, with each pair of characters separated by colon; for example, **XX:XX:XX:XX:XX:XX**
- The file must contain the correct MAC addresses. Any incorrect MAC addresses will result in the licensed features not working on designated units.
- The number of MAC addresses must not exceed the number of MAC addresses allowed for the License Authorization Code entered for a particular file.

For more information about feature licensing, see *Avaya Virtual Services Platform 7000 Series Fundamentals*, NN47202–101.

Software and hardware capabilities

The following table lists supported software and hardware scaling capabilities for the Avaya Virtual Services Platform 7000 Series Software Release 10.2.

The information in this table supersedes information contained in other technical documentation for VSP 7000 Series.

Feature	Maximum number supported
General	
Fabric Interconnect Stack bandwidth (8 units)	5,120 Gbps (full duplex)
Fabric Interconnect Stack (number of units).	8
Fabric Interconnect Mesh bandwidth (32 units)	20,480 Gbps (full duplex)
Fabric Interconnect Mesh (number of units).	200
MDA supported on each VSP 7000	1

Feature	Maximum number supported
Layer 2	
Avaya Spanning Tree Groups	8
DHCP Snooping table entries	1,024
MAC addresses	131,071 (32,767 with SMLT)
Multi-Link Trunking (MLT), Distributed Multi-Link Trunking (DMLT), and Link Aggregation (LAG) groups	32
MLT Links or ports per MLT, DMLT, or LAG	8
MLT Maximum MAC Learning rate on an MLT trunk	2000 new MAC addresses per second
Spanning Tree Group instances (802.1s)	8
Static MAC addresses	1,024
VLAN Concurrent VLANs	1024
VLAN Protocol-based VLANs	7
VLAN Supported VLAN IDs	1–4094 • 0 and 1095 reserved • 4001 reserved by STP • 4002–4008 reserved by multiple STP groups
Layer 3	
IP interfaces (VLANs or Brouter ports)	256
ARP Entries total (local, static, and dynamic)	4,096
ARP Entries — local (IP interfaces for each switch or stack)	256
ARP Entries — static	256
IPv4 Routes total (local, static, and dynamic)	4,096
IPv4 Local Routes	256
IPv4 Static Routes	512
Dynamic Routing interfaces (RIP and OSPF)	64
OSPF Areas	4 (3 areas plus area 0)
OSPF Adjacencies	64
OSPF Link State Advertisements (LSAs)	10,000
OSPF Virtual Links	16

Feature	Maximum number supported
OSPF Host Routes	32
ECMP (Max concurrent equal cost paths)	4
ECMP (Max next hop entries)	4,096
VRRP instances	255 IDs (64 active)
Management Routes	4
UDP Forwarding Entries	128
DHCP Relay Entries	256
DHCP Relay Forward Paths	512
Multicast	
IGMP Allow-flood multicast addresses	4096 (The maximum number of allow-flood multicast entries is the aggregate of the number of devices in each VLAN receiving multicast streams.)
IGMP Allow-flood multicast addresses per VLAN	128
IGMP multicast groups	1024
Quality of Service	
Egress queues	Configurable 1–8
Egress queues (Lossless Mode)	2
QoS rules	Precedence levels (slices); 10 Max QoS policies per port: 8 Max Filters per precedence: 128 (Precedence 1–4) Max Filters per precedence: 256 (Precedence 5–10) Max Meters per precedence: 128 Max Counters per precedence: 64 (Precedence 1–4) Max Counters per precedence: 128 (Precedence 5–10) Range Check Entries: 32 Traffic-profile Entries: 1024
QoS Traffic Profile Criteria — Layer 2 **Note:** Traffic Profiles provide the combined benefits of ACLs, Filters, and Classifiers.	• Source MAC address/mask • Destination MAC address/mask • VLAN ID range • VLAN tag • EtherType

Feature	Maximum number supported
	• Packet type • 802.1p priority values
QoS Traffic Profile Criteria — IPv4	• IPv4 source address/mask • IPv4 destination address/mask • IPv4 address type • IPv4 protocol type • IPv4 DSCP value • IPv4 source TCP port range • IPv4 source UDP port range • IPv4 destination TCP port range • IPv4 destination UDP port range • IPv4 flags • TCPv4 control flags • IPv4 options
QoS Traffic Profile Criteria — IPv6	• IPv6 source address/mask • IPv6 destination address/mask • IPv6 address type • IPv6 flow identifier • IPv6 next-header • IPv6 DSCP value • IPv6 source TCP port range • IPv6 source UDP port range • IPv6 destination TCP port range • IPv6 destination UDP port range
QoS Traffic Profile Criteria — System	• unknown IP multicast • known IP multicast • unknown non-IP multicast • known non-IP multicast • non-IP packet • unknown unicast packet
SMLT	
SMLT operational mode	Standalone or Stacked

Feature	Maximum number supported
SMLT configuration	Triangle or Square
SMLT: MLT uplinks	31
SMLT: SLT uplinks	128
SMLT: SMLT/LACP uplinks	5
SMLT: SLT/LACP links	12
SMLT: SLPP VLANs	20
SMLT: IST using LACP	Not supported in this release
SMLT: IST/LACP	Not supported in this release
SMLT: Static IP Routes supported across IST	Supported
SMLT: Static IP Routes over SLT/MLT links	Supported
SMLT: Dynamic IP Routing over SLT/MLT links	Not supported in this release
SMLT: Dynamic Routing Protocol over SMLT (IST)	Not supported in this release
SMLT: IGMP over SMLT	Not supported in this release
SMLT: SPB over SMLT	Not supported in this release
SMLT: SMLT/IST over rear ports in rear-port mode	Not supported in this release
SPB	
SPB operational mode	Standalone
SPB Customer VLANs (C-VLANs) per node	500
SPB ISIDs per node	1024
SPB Switched UNIs	4096
SPB nodes per region (HW records)	500
SPB nodes per region	200
SPB (IS-IS) adjacencies per node	24
Miscellaneous	
HTTP Server IPv4	3 sessions
HTTP Server IPv6	3 sessions
IPFix number of sampled flows	100,000
LLDP Neighbors	800
LLDP Neighbors per port	16

Feature	Maximum number supported
Port Mirroring instances	4
RMON alarms	800
RMON Ethernet history	249
RMON Ethernet statistics	110
RMON events	800
Telnet Client IPv6	4 sessions
Telnet Server IPv6	4 sessions

Supported browsers

Virtual Services Platform 7000 supports the following browsers to access the Enterprise Device Manager (EDM):

- Microsoft Internet Explorer 7.x
- Microsoft Internet Explorer 8.x
- Mozilla Firefox 3.6.x
- Mozilla Firefox 12.x
- Mozilla Firefox 14.x

Note:

Due to an issue in Firefox versions greater than 3.6.x, you might not be able to import SSL certificates using IPv6. As a workaround, you can use the hostname (with host IPv6 address resolved by DNS or editing the local hosts file), or use Microsoft Internet Explorer 8.x.

Upgrading switch software using ACLI

Use this procedure to specify the download target image and change the software version running on the switch.

About this task

You can update either of the following:

- the active software image
- the non-active software image

Warning:

DATA LOSS CAN OCCUR — Do not upgrade directly from Release 10.0 to Release 10.2 or later.

If the switch is running Release 10.0, you must upgrade to Release 10.1 before upgrading to 10.2 or later software. Upgrading from Release 10.0 to the current release can cause accidental erasure of the agent image on the switch. If the primary agent image is erased, during the next reboot the switch attempts to boot from the secondary agent image.

The software image download process occurs automatically within a stack if different software is present. This process deletes the contents of the flash memory and replaces it with the specified software image.

Tip:

To track the progress of the download process, you can observe the switch front panel LEDs.

Depending on network conditions, the download process may take up to 10 minutes.

Important:

Do not interrupt the download process.

You can update the runtime image (agent code) on the switch while the switch is operational. If you specify the no-reset option, the new software is updated on FLASH, but is not running on the switch. If you do not specify the no-reset option, once the download of switch software is complete, the switch or Fabric Interconnect Stack resets and restarts with the new image.

Procedure

1. Log on to the ACLI Privileged Executive mode.
2. At the command prompt, enter the following command:

```
download [address <A.B.C.D> | <WORD>] {primary | secondary}
{image <image_name> | image-if-newer <image_name> | diag
<image_name>} [no-reset] [<usb> | <tftp> | <sftp>]
```

Variable definitions

The following table describes parameters to help you use the **download** command to upgrade agent software.

Variable	Value
address<*A.B.C.D*> | <*WORD*>	The IPv4 or IPv6 address of the TFTP server on which the agent image is hosted. • A.B.C.D — Specifies the IP address in IPv4 format. • WORD — Specifies the IP address in IPv6 format. The address parameter is optional and, if omitted, the switch defaults to the TFTP server specified by the **tftpserver** command unless software download is to take place using a USB mass storage device.
primary | secondary	Specifies the image to download: primary or secondary.
image<*image_name*>	Specifies the name of the software image file to be downloaded from the TFTP server.
image—if—newer <*image_name*>	Specifies the name of the software image to be downloaded from the TFTP server if newer than the currently running image.
diag <*image_name*>	Specifies the name of the diagnostic image to be downloaded from the TFTP server.
no-reset	Stops the switch from resetting after completion of the software download.
<*usb*> | <*tftp*> | <*sftp*>	Specifies the software download source. • USB — download is from a USB mass storage device. • TFTP — download is from a TFTP server. • SFTP — download is from a SFTP server.
Note: The image, image-if-newer, and diag parameters are mutually exclusive and you can execute only one at a time.	

Upgrading switch software using EDM

Use the following procedure to change the software version running on the switch using Enterprise Device Manager (EDM).

Warning:

DATA LOSS CAN OCCUR — Do not upgrade directly from Release 10.0 to Release 10.2 or later.

If the switch is running Release 10.0, you must upgrade to Release 10.1 before upgrading to 10.2 or later software. Upgrading from Release 10.0 to the current release can cause accidental erasure of the agent image on the switch. If the primary agent image is erased, during the next reboot the switch attempts to boot from the secondary agent image.

Procedure

1. From the navigation tree, double-click **Edit**.
2. In the Edit tree, click **File System**.
3. In the work area, click the **Config/Image/Diag file** tab.
4. Configure the parameters required to perform the image download.
5. On the toolbar, click **Apply**.

Result

The software download occurs automatically once you click Apply. This process erases the contents of the flash memory and replaces it with the new software image.

Important:

Do not interrupt the download. Depending on network conditions, this process can take up to 10 minutes.

When the download is complete, the switch automatically resets and the new software image initiates a self test.

Important:

During the download process, the management functionality of the switch is locked. Normal switching operations continue to function until the switch resets.

Variable definitions

The following table describes updating the binary configuration, image, and diagnostic files.

Variable	Value
TftpServerInetAddressType	Specifies the IP address type of the TFTP or SFTP server. Values include IPv4 or IPv6.
TftpServerInetAddress	Specifies the IP address of the TFTP or SFTP server.
BinaryConfigFilename	Specifies the binary configuration file currently associated with the switch. This field only applies to binary configuration files.
BinaryConfigUnitNumber	Specifies the unit number portion of the configuration file to be used for the standalone unit configuration. Values range from 0 to 8. If 0, the unit number is ignored. This field only applies to binary configuration files.
ImageFileName	Specifies the name of the image file currently associated with the switch. You can change this field to the filename of the software image to be downloaded.
FwFileName(Diagnostics)	Specifies the name of the diagnostic file currently associated with the switch. You can change this field to the filename of the software image to be downloaded.
Usb TargetUnit	Specifies the unit number for USB, or the transfer type to use during the upload or download operation. Values include: • 1 to 8 — USB on unit 1 to 8 • 9 — USB on a standalone unit • 0 — TFTP server • 10 — SFTP server
Image	Specifies if the image to download is the primary or secondary image.

Variable	Value
Action	Specifies the action to perform during the file transfer. Values include: • dnldConfig — Downloads the configuration file from a TFTP or SFTP server • upldConfig — Uploads the configuration file to a TFTP or SFTP server • dnldConfigFromUsb — Downloads the configuration file from a USB storage device. • upldConfigToUsb — Uploads the configuration file to a USB storage device. • dnldImg — Downloads the agent image file from a TFTP or SFTP server. • dnldImgIfNewer — Only downloads if newer than current image. • dnldImgNoReset — Downloads the agent image and does not reset the switch. • dnldImgFromUsb — Downloads the agent image from a USB storage device. • dnldFw — Downloads the diagnostic image from a TFTP or SFTP server. • dnldFwNoReset — Downloads the diagnostic image and does not reset the switch. • dnldFwFromUsb — Downloads the diagnostic image from a USB storage device. • dnldImgFromSftp — Downloads the agent image from a SFTP server. • dnldFwFromSftp — Downloads the diagnostic image from a SFTP server. • dnldConfigFromSftp — Downloads the configuration file from a SFTP server. • upldonfigToSftp — Uploads the configuration file to a SFTP server.

Variable	Value
	• dnldImgFromSftpNoReset — Downloads the agent image from a SFTP server and does not reset the switch. • dnldFwFromSftpNoReset — Downloads the diagnostic image from a SFTP server and does not reset the switch.
Status	Indicates the status of the last action since the last switch reboot. Values include: • other — No action has taken place. • inProgress — The selected action is currently in process. • success — The selected action completed successfully. • fail — The selected action failed.

Supported standards, MIBs, and RFCs

This section lists the standards, MIBs, and RFCs supported by the Avaya Virtual Services Platform 7000 Series.

Standards

The following IEEE Standards contain information that applies to the Avaya Virtual Services Platform 7000 Series.

- IEEE 802.1 — Port VLAN, Port and Protocol VLANs, VLAN Name, Protocol Entity
- IEEE 802.1AB — Layer Link Discovery Protocol
- IEEE 802.1aq — Shortest Path Bridging
- IEEE 802.1ax — Link Aggregation Control Protocol
- IEEE 802.1D — Standard for Spanning Tree Protocol
- IEEE 802.1p — Prioritizing
- IEEE 802.1Q — VLAN Tagging
- IEEE 802.1s — Multiple Spanning Tree Protocol
- IEEE 802.1v — VLAN Classification by Protocol and Port

- IEEE 802.1w — Rapid Spanning Tree Protocol
- IEEE 802.3 — Ethernet
- IEEE 802.3ab — Gigabit Ethernet over Copper
- IEEE 802.3ad — Link Aggregation
- IEEE 802.3ae – 10 Gbps Ethernet
- IEEE 802.3aq — Ethernet over multimode fiber
- IEEE 802.3x — Flow Control
- IEEE 802.3z — Gigabit Ethernet over Fiber-Optic

RFCs and MIBs

For more information about networking concepts, protocols, and topologies, consult the following RFCs and associated MIBs:

- RFC 768 (UDP)
- RFC 791 (IP)
- RFC 792 (ICMP)
- RFC 793 (TCP)
- RFC 826 (ARP)
- RFC 854 (Telnet)
- RFC 894 (IP over Ethernet)
- RFC 951 (BootP)
- RFC 1058 (RIP v1)
- RFC 1112 (IGMPv1)
- RFC 1157 (SNMP)
- RFC 1213 (MIB-II)
- RFC 1271 (RMON)
- RFC 1305 (NTP v3)
- RFC 1350 (TFTP)
- RFC 1493 (Bridge MIB)
- RFC 1583 (OSPF v2)
- RFC 1757 (RMON)
- RFC 1769 (SNTP)

- RFC 1850 (OSPF v2 MIB)
- RFC 1886 (DNS Extensions for IPv6)
- RFC 1905 (SNMP)
- RFC 1906 (SNMP Transport Mappings)
- RFC 1907 (SNMP MIB)
- RFC 1945 (HTTP v1.0)
- RFC 1981 (Patch MTU Discovery for IPv6)
- RFC 2011 (SNMPv2 IP MIB)
- RFC 2012 (SNMPv2 TCP MIB)
- RFC 2013 (SNMPv2 UDP MIB)
- RFC 2131 (BootP/DHCP Relay Agent)
- RFC 2236 (IGMPv2)
- RFC 2328 (OSPF v2)
- RFC 2453 (RIP v2)
- RFC 2460 (IPv6)
- RFC 2461 (Neighbor Discovery for IPv6)
- RFC 2462 (Auto-Configuration for Link Local Addresses)
- RFC 2464 (Transmission of IPv6 packets over Ethernet networks)
- RFC 2474 (DiffServ)
- RFC 2475 (DiffServ)
- RFC 2665 (Ethernet MIB)
- RFC 2674 (Q-BRIDGE-MIB)
- RFC 2737 (Entity MIBv2)
- RFC 2819 (RMON MIB)
- RFC 2863 (Interfaces Group MIB)
- RFC 2865 (RADIUS)
- RFC 2866 (RADIUS Accounting)
- RFC 2933 (IGMP MIB)
- RFC 3046 (DHCP Relay Agent information option)
- RFC 3162 (RADIUS and IPv6)
- RFC 3246 (Expedited Forwarding Behavior)
- RFC 3410 (SNMPv3)

- RFC 3411 (SNMP Frameworks)
- RFC 3412 (SNMP Message Processing)
- RFC 3413 (SNMPv3 Applications)
- RFC 3414 (SNMPv3 USM)
- RFC 3415 (SNMPv3 VACM)
- RFC 3416 (SNMP)
- RFC 3417 (SNMP Transport Mappings)
- RFC 3418 (SNMP MIB)
- RFC 3584 (Coexistence of SNMPv1/v2/v3)
- RFC 3768 (VRRP)
- RFC 3917 (IPFix)
- RFC 3954 (Netflow Services Export v9)
- RFC 3993 (DHCP Subscriber-ID suboption)
- RFC 4007 (Scoped Address Architecture)
- RFC 4022 (TCP MIB)
- RFC 4113 (UDP MIB)
- RFC 4250 (SSH Protocol assigned numbers)
- RFC 4251 (SSH Protocol architecture)
- RFC 4252 (SSH Authentication Protocol)
- RFC 4253 (SSH Transport Layer Protocol)
- RFC 4254 (SSH Connection Protocol)
- RFC 4291 (IPv6 addressing architecture)
- RFC 4293 (IPv6)
- RFC 4432 (SSH RSA)
- RFC 4443 (ICMPv6)
- RFC 4541 (Considerations for IGMP and MLD snooping switches)
- RFC 4604 (IGMPv3)

Chapter 4: Resolved issues

Use the information in this section to learn more about issues that have been resolved.

Resolved issues for Release 10.2

The following table lists the issues resolved in the current software release.

Reference number	Description
wi00996632	**10m Fabric Interconnect :** When the 10m Fiber Fabric Interconnect cable (AL77018005–E6) is used with the VSP 7000 packet errors are no longer seen on the connection if the ambient temperature is approaching 50 degrees Celsius.
wi00980557	**Broadcast, Multicast Rate Limiting :** When running the show rate-limit command, the multicast rate displayed is now shown correctly for the past 24 hours.
wi00978187	**EDM, Help Files :** When the Fabric Interconnect Stack is running from a Temporary Base Unit (TBU), Enterprise Device Manager (EDM) is now able to access Help Files if these are configured.
wi00946937	**EDM, TACACS password :** If ACLI password type is set to TACACS, when you connect to the switch using Enterprise Device Manager (EDM), the switch will now correctly prompt for a username or password.
wi00980352	**EDM, TCP Listeners :** When using Enterprise Device Manager (EDM), the switch now correctly displays TCP Listeners for IPv4.
wi00952515	**EDM, VLAN range :** When using Enterprise Device Manager (EDM), it is now possible to specify a VLAN range when trying to add a new L2 Classifier Element.
wi00973969	**RADIUS, Telnet Disconnect, Accounting Messages :** When a telnet session to the switch is closed and RADIUS accounting is enabled, the switch now correctly generate a single radius message.

Reference number	Description
wi00940709	**RMON, ACLI :** Display of RMON event Community and LastTimeSent information is now available when using ACLI.
wi00929094	**Stack Loopback test (External) :** The External Stack Loopback must be run from the top FI-ports. A message regarding this function is now displayed on the console.
wi00994847	**Stack Loopback test (Internal) :** When performing an internal Fabric Interconnect Stack loopback test, a failure message is no longer incorrectly displayed.
wi00974439	**USB :** Saving the binary config no longer fails when saving to a USB device which is not connected to the Base Unit (BU) or Temporary Base Unit (TBU) of a Fabric Interconnect Stack with a large number of units
wi00997689	**Lossless Mode, Flow Control, SFP+**: When operating in Lossless mode and flow control is enabled on SFP+ ports, if the SFP+ device is unplugged, and then plugged back in, the port no longer reverts to flow control disabled.
wi00984580	**USB**: When using older USB devices, on first insertion the USB device is now recognized by the VSP 7000.
wi00960906	**MLT**: If running 100% line rate traffic on all MLT links, the `show mlt utilization` command no longer incorrectly displays an MLT utilization of approximately 92% for each MLT link.
wi00973130	**MSTP**: When a Fabric Interconnect Stack is configured with MSTP mode of operation and a large number of VLANs with redundant MLT links, resting a non Base Unit no longer causes a broadcast storm if the network contains many redundant loops.

Chapter 5: Known issues and limitations

Use the information in this section to learn more about known issues and limitations.

Where appropriate, use workarounds provided for the known issues.

Known issues

The following table lists and describes known issues and limitations for Avaya Virtual Services Platform 7000 Series Software Release 10.2. Where available and appropriate, workarounds are provided.

Reference number	Description
wi00998949, wi01042576	**OS, Agent upgrades**: **DO NOT UPGRADE DIRECTLY FROM RELEASE 10.0 TO CURRENT RELEASE**. **Warning:** If you upgrade from Release 10.0 directly to Release 10.2, the new agent image buffer is exceeded and primary agent image is erased. **Workaround**: You must upgrade from Release 10.0 to 10.1, and then to 10.2. If you accidently upgrade the agent code directly from 10.0 to 10.2 or later, the primary agent image erases and the switch will attempt to boot from the secondary image. You can then perform an upgrade from 10.0 to 10.1, and then to 10.2.
wi00995011	**7008XLS MDA, Resource Counters**: When the ports on the 7008XLS MDA are oversubscribed due to known Unicast traffic, the “Dropped on no resources” counter does not increment.
wi00972139	**AAUR/DAUR, Release 10.0**: If you add a VSP 7000 switch running a software release prior to 10.1 to an operational Fabric Interconnect Stack, the unit will not join the stack. The UP/Down LEDs will remain amber on the 10.0 switch to indicate that the unit is unable to correctly join the Fabric Interconnect Stack. **Workaround**: You need to upgrade the agent code software on a VSP 7000 to release 10.1 or later before adding the unit to an operational Fabric Interconnect Stack.

Reference number	Description
wi00978314	**EDM, 32 ports**: When using Enterprise Device Manager (EDM) with the VSP 7000, some work areas might incorrectly indicated 32 ports are present, even if no MDA is or has been inserted in the switch.
wi00949421	**EDM, Chrome**: If you use Google Chrome to access the switch via Enterprise Device Manager (EDM), then you might not be able to login to the switch if local username and passwords are configured. **Workaround**: The supported browsers for managing a VSP 7000 switch are IE or Firefox. It is recommended to use one of the supported browsers.
wi00972061	**EDM, Fan Status LEDs**: When using Enterprise Device Manager (EDM), the Device physical view does not display the FAN status LED colors: only grey is displayed for the fan status. **Workaround**: Check the FAN status LEDs on the units.
wi00933142	**EDM, MLT BPDU setting**: When using Enterprise Device Manager (EDM), it is not possible to specify the MLT BPDU send or receive mode settings. **Workaround**: You must configure MLT BPDU send or receive mode using ACLI.
wi01029280	**EDM, TACACS password**: If ACLI password type is set to TACACS, Enterprise Device Manager (EDM) is disabled by default. **Workaround**: If EDM is required, manually enable the web server using ACLI. Log on to EDM using local password authentication.
wi00930939	**Netmask**: Modifying the netmask without an IP address might result in connectivity loss. **Workaround**: When modifying the netmask use an IP address in the command: `ip add A.B.C.D mask A.B.C.D`.
wi01010266	**Out of Band mgmt port**: Errors display on the console when the OOB mgmt port receives oversized packets. **Workaround**: Do not use jumbo frames with the OOB mgmt port.
wi01029850	**L2Ping, CFM**: Due to system timing on the VSP 7000, the roundtrip times (minimum, maximum, and average) displayed from L2ping show higher values than other platforms, such as the VSP 9000 and ERS 8800.
wi00971757	**Lossless Mode**: When the switch is operating in Lossless mode and flowcontrol is disabled on a port, the dropped on no resources counter stays at zero on egress port, even if the port becomes over-subscribed.

Reference number	Description
	Workaround: You must enable flow control on all ports if the switch is operating in Lossless mode.
wi00974573	**LACP**: LAGs might show duplicates on a Temporary Base Unit when you perform `show lacp agg`. **Workaround**: LAGs are not duplicated, this is a display issue only.
wi01048943	**Port mirroring**: ManytoOneRxTx port mirroring instance does not work for unknown unicast, multicast, and broadcast traffic.
wi01032538	**Port mirroring**: If you configure all four port mirroring instances, only the first two are functional.
wi01031322	**Rear-port mode, ISIS:** By default, all Fabric Interconnect ports operating in rear-port mode use the same LACP key. If you enable ISIS on a rear-port, ISIS is enabled on all rear-ports. **Workaround:** Working as designed. If you require ISIS to be enabled only on some rear-ports (Fabric Interconnect), you can group specific ports into LAGs according to LACP keys, and apply ISIS to the desired LAG.
wi01031500	**Rear-port mode, ISIS:** The command `show isis interface` displays all rear-ports, irrespective of state, because all rear-ports are members of the same default LAG. **Workaround:** Working as designed. If you require ISIS to be enabled only on some rear-ports (Fabric Interconnect), you can group specific ports into LAGs according to LACP keys, and apply ISIS to the desired LAG.
wi01027055	**Rear-port mode, SPBM**: By default, all Fabric Interconnect ports operating in rear-port mode use the same LACP key. If you modify a rear-port metric, such as the SPBM-L1–Metric, the modification applies to all ports which are members of the same LAG. **Workaround:** Working as designed. If you require different parameters on specific rear-ports (Fabric Interconnect), you can group specific ports into LAGs according to LACP keys..
wi00934503	**Traceroute**: If you have two VSP 7000 switches directly connected to each other, performing a traceroute from one switch to the other will fail; even though the switches can be reached using the ping command.
wi01045294	**Stress testing L3**: Stack might break in stress scenario with 4k ospf/rip routes, 1k vlans (256 L3 vlans) and 128k mac entries.

Reference number	Description
wi01046994	**EDM, LLDP**: LLDP tx-tlv local-mgmt-address is disabled for all virtual MDA ports if disabled on one port. **Workaround**: Re-enable the TLV using ACLI.
wi00979441	**Automation**: Intermittently configuration objects might change unexpectedly during a large number of random resets or power cycles. Changes are minor and not expected in customer configurations.
wi00971049	**Automation**: Intermittent, Error tCliAudit errors writing to flash. Flash program too long written to System Log. Issue is seen in automation setups, not likely in customer networks.
wi00974728	**VLACP, Traps**: Inconsistent logging of VLACP traps can occur if enabled.
wi01028347	**Rear-port mode, SMLT**: If you configure IST/SMLT on rear-ports the settings apply, causing high CPU usage and can lead to a system crash. **Important:** Do not configure IST or SMLT in rear-port mode. IST and SMLT is not supported on rear-ports.
wi01040581	**Rear-port mode, full default**: If you fully default a switch operating in rear-port mode, the rear-port mode is disabled. Upon reboot, the fully defaulted unit attempts to join a FI stack, causing other connected units operating in rear-port mode to crash. **Workaround**: Before performing a full default on a unit operating in rear-port mode, disconnect the FI cables on the unit. Reconnect the FI cable after the unit is re-configured as required.
wi01039526	**SPBM, ERS 8800 Interoperability**: When connecting a VSP 7000 to an ERS 8800 running SPBM, IS-IS adjacencies are not formed with the ERS 8000, unless the ERS 8800 is running Release 7.1.3 or later.
wi01043735	**SPBM, Port mirroring**: If you mirror a port with IS-IS enabled, the mac-in-mac 802.1ah header is stripped from all SPBM encapsulated packets.
wi01053545	**LACP/SMLT**: If you reboot a non-base unit in a stack with LAGs configured, packet loss or flood can occur for approximately 40 seconds.
wi01059397	**FI ports**: If you change a switch configuration from FI stacking to FI rear-port mode, or vise versa, without removing the FI cables, there is a high probability of causing a loop across the FI ports.

Reference number	Description
	Workaround: Before performing an FI port configuration change, disconnect the FI cables on the unit. Reconnect the FI cable after the unit is re-configured as required.
wi01060600	**SMLT/IST**: If the Base Unit is powered off, the non-base unit traffic might not recover if only one IST link is configured.
wi01049340	**QoS resources**: If a Release 10.1 unit with all QoS precedences used is upgraded to Release 10.2, the QoS policies will disable. This failure occurs because SPBM requires QoS precedence 9 even if SPBM is not configured. **Workaround**: Reduce QoS precedence usage before upgrading to Release 10.2.
wi01050660	**Rear-port mode**: If a switch operating in rear-port mode is reset, the rear port link might not restore after a reset. **Workaround**: Remove and re-insert the FI cables once the unit is operational.
wi01066446	**Rear-port mode**: If you change between standard rear-port mode and SPB rear-port mode the switch requires a reboot and partial configuration reset. Standard rear-port mode does not support SPB.
wi01045294	**Stress testing L3**: Stack break might occur In a large configuration with 4k OSFP/RIP routes, 1k VLANs (256 L3 VLANs), and 128k MAC entries.
wi01059600	**DHCP Relay**: DHCP Relay is not functioning on the Brouter port.
wi01009886	**Stress testing**: If you run a clear ARP command with a stream of 100 IPs at 5000 pps in a square SMLT and VRRP configuration, up to 15 seconds traffic loss might occur.
wi01023541	**Lossless PFC mode**: In Lossless-PFC mode, regardless of flow control settings the port sends PFC frames on oversubscription. **Workaround**: You cannot configure the flow control modes for Lossless-PFC. Symm, Asymm, or Disabled flow control modes apply to Lossless (PAUSE) mode only.
wi00888731	**Boot loader stops**: Intermittent. The diagnostic/boot/agent loader might stop on boot after a switch or stack reset or upgrade. The unit might stop loading Agent software and appear hung. **Workaround**: Press a key on the console interface, the unit will continue to load.
wi0101165	**EDM**: If you boot the unit in rear port mode the configuration does not partially default. **Workaround**: Use ACLI to enable rear port mode.

Reference number	Description
wi01011829	**VRRP**: Intermittent. Error message might occur when enabling or disabling VRRP. **Workaround**: Reset the unit.
wi01018227	**EDM**: If an active 10 Gb copper interface is enabled using EDM, the port status is amber although the port is up.
wi01018538	**RMON**: RMON event time is uptime regardless of other settings for the clock.
wi01019793	**IPv6, TFTP**: Binary configuration cannot be retrieved for a stack if using IPv6 management and TFTP server. **Workaround**: Use IPv4 based TFTP server to retrieve the binary configuration.
wi01026033	**OOB MGMT**: Autotopology is not functional on the Out-of-band management interface
wi01028730	**IST**: IST might bounce when a non-base unit rejoins the stack if VLACP short is enabled on IST ports. **Workaround**: Use Long Timers instead of Short Timers.
wi01034248	**Rear port mode**: Port 40 linked to port 36 in rear port spb mode can cause inconsistency regarding port state. **Workaround**: This is an invalid configuration.
wi01042491	**VLAN**: Adding one port to 1000 VLAN might take an extended period of time. **Workaround**: None, but adding multiple ports to a large number of VLANs takes time to complete.
wi01043365	**SMLT, EDM**: Intermittent. Cannot disable or enable IST using EDM (Inconsistent value). **Workaround**: Use ACLI to configure IST if the command did not execute from EDM.
wi01046311	**USB**: USB boot loading does not function with QoS lossless enabled. **Workaround**: Use ACLI to configure lossless mode.
wi01046994	**EDM, LLDP**: LLDP tx-tlv local-mgmt-address disables on all MDA ports if disabled on one port. **Workaround**: Use ACLI to re-enable the TLVs disabled on the MDA port.
wi01048843	**Rear port mode**: Binary configuration of a unit with rear port mode enabled cannot be retrieved on a defaulted unit.
wi01049115	**Rear port mode, LACP**: LACP mode is turned off for rear ports after loading a configuration, when ports are removed from a VLAN. **Workaround**: Use ACLI to re-enable LACP on the ports.
wi01049203	**EDM off-box**: SNMP agent intermittently times out while configuring or retrieving outputs, IP connectivity is up.

Reference number	Description
	Workaround: Increasing the number of SNMP retries and timeout timer can improve this issue.
wi01049509	**SMLT LACP**: SMLT LACP bounces at 14k MACs wien using LACP Short Timer. **Workaround**: Avaya recommends using Long Timers on LACP ports running SMLT.
wi01050158	**EDM off-box**: Timeout when uploading a configuration to TFTP server. The file is uploaded correctly. **Workaround**: Increasing the number of SNMP retries and timeout timer can improve this issue.
wi01050558	**EDM off-box**: The default timer cannot be set for VRRP hold-down-timer. **Workaround**: Use EDM on-box or ACLI to configure this parameter.
wi01061172	**EDM off-box**: Using Element Manager authenticated with SNMPv3 cannot create additional SNMPv3 users and causes error messages. **Workaround**: Use EDM on-box or ACLI to configure.
wi01064985	**EDM off-box**: Packet per second (PPS) rate limit value cannot be defaulted (0), timeout. **Workaround**: Use EDM on-box or ACLI to configure.
wi01050783	**SLPP**: LACP SLT links are disabled without loops after reboot with aggressive values 5 and 50. **Workaround**: Configure the threshold values at least 5 times the number of VLANs and use long timers.
wi01050967	**EDM**: IP ARP tab cannot display the brouter static ARPs if the ARP table contains multiple entries. **Workaround**: Use ACLI to display the table.
wi01052477	**EDM**: EDM multiple port configuration mode cannot configure the speed on multiple MDA ports. **Workaround**: Modify a single port at a time or use ACLI to modify multiple ports.
wi01017515	**EDM**: The asset ID string that follows after a < character does not display in EDM. **Workaround**: Use ACLI if required
wi01050306	**EDM**: Cannot configure rear ports using EDM. **Workaround**: Use ACLI to configure rear ports.
wi01057163	**SNMP**: Timeout when auto saving configuration to NVRAM and performing successive SNMP operations, such as deleting 500 VLANs. **Workaround**: Increasing the number of SNMP retries and timeout timer can improve this issue.

Reference number	Description
wi01057995	**Show port** : Enhancement info is incomplete when issuing show port over SSH with a terminal length of 0. **Workaround**: Configure the terminal length to 40 and try again.
wi01059621	**RADIUS**: RADIUS is not supported on OOB management port.
wi01060852	**RADIUS**: You cannot use EDM to change the RADIUS password. **Workaround**: Use ACLI to change the password.
wi01061771	**Display**: 50 m and 100 m fiber rear port connections show as 0.0m length with show stack-cable command.
wi01067876	**CFM**: CFM does not function after a reset when ISIS is configured on lower number of ports with IPFIX.
wi01068140	**VLAN**: Unclear error message if attempting to use VLAN 4001 with SPBM. **Workaround**: The VLAN ID is in use for another feature (STP). Working as designed.
wi01016012	**Route-map**: A route-map with a filename of “detail” cannot filter when displaying route-maps. **Workaround**: Change the route-map name.
wi01062498	**SMLT**: An error does not display when peer IP is the same as the local VLAN IP.
wi01068125	**SMLT**: When performing an SMLT failover to an ERS 8800, up to 4 seconds of packet loss might occur.
wi01068432	**SPBM**: SPBM nickname starting with 3.33 causes Multicast traffic to drop. **Workaround**: Do not use SPBM nicknames that begin with the string 3.33. Other nicknames are not affected.
wi01050082	**VCC**: You cannot disable tagging on ports if VLAN configuration control (VCC) is set to automatic. **Workaround**: You can disable tagging on ports if VCC is set to flexible or autopvid.
wi01074797	**SPBM**: If you modify UNI SPB ports from trusted to untrusted or unrestricted, traffic flow might interrupt on the associated CVLANs and switched UNI. **Workaround**: You must configure SPB ports as trusted interfaces (default). If a port QoS setting is changed to untrusted or unrestricted, you must configure QoS as trusted and reboot the switch for traffic to recover.