Network Design Reference for Avaya Virtual Services Platform 4000

Release 3.1
NN46251-200
Issue 01.09
February 2014

Notice

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

Documentation disclaimer

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

Link disclaimer

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

Warranty

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

Licenses

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

Avaya grants you a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to you. “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 or an Instance of the Software only on a number of Designated Processors up to the number indicated in the order. Avaya may require the Designated Processor(s) to be identified in the order by type, serial number, feature key, Instance, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

Heritage Nortel Software

“Heritage Nortel Software” means the software that was acquired by Avaya as part of its purchase of the Nortel Enterprise Solutions Business in December 2009. The Heritage Nortel Software currently available for license from Avaya is the software contained within the list of Heritage Nortel Products located at http://support.avaya.com/LicenseInfo/ under the link “Heritage Nortel Products”, or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants 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, Hosted Service, or hardware provided by Avaya. All content on this site, the documentation, Hosted Service, and the Product provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software unless expressly authorized by Avaya. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

Third Party Components

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

Preventing Toll Fraud

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

Avaya Toll Fraud intervention

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

Trademarks

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

Avaya is a registered trademark of Avaya Inc.

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

Downloading Documentation

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

Contact Avaya Support

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

Contents

Chapter 1: Introduction

Purpose

This document provides information to help you build robust, efficient networks using the VSP 4000 platform.

You can use the examples and important design guidelines listed in this document for many features and protocols.

Related resources

Documentation

See the *Avaya Virtual Services Platform 4000 Documentation Roadmap*, NN46251–100 for a list of the documentation for this product.

Training

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

Avaya Mentor videos

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

About this task

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

- To find videos on the Avaya Support website, go to http://support.avaya.com, select the product name, and check the *videos* checkbox to see a list of available videos.
- 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.

Note:

Videos are not available for all products.

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.

Searching a documentation collection

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

Before you begin

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

Procedure

1. Extract the document collection zip file into a folder.
2. Navigate to the folder that contains the extracted files and open the file named <*product_name_release*>.pdx.
3. In the Search dialog box, select the option **In the index named <*product_name_release*>.pdx**.
4. Enter a search word or phrase.
5. Select any of the following to narrow your search:
   - Whole words only
   - Case-Sensitive
   - Include Bookmarks
   - Include Comments
6. Click **Search**.
   The search results show the number of documents and instances found. You can sort the search results by Relevance Ranking, Date Modified, Filename, or Location. The default is Relevance Ranking.

Chapter 2: New in this release

The following sections detail what is new in the *Network Design Reference for Avaya Virtual Services Platform 4000* (NN46251–200), for release 3.1.

Features

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

IP multicast over SPBM

In Release 3.1, Virtual Services Platform 4000 supports IP multicast over Shortest Path Bridging MAC (SPBM). IP multicast over SPBM greatly simplifies multicast deployment, with no need for any multicast routing protocols such as PIM.

With IP multicast over SPBM, Avaya Virtual Services Platform 4000 supports the following:

- Layer 2 Virtual Services Network with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network (Layer 2 VSN with IP multicast over SPBM).
- IP multicast routing support for Global Routing Table using SPBM in the core (IP Shortcuts with IP multicast over SPBM).
- Layer 3 Virtual Services Network with VRF based IP multicast routing support over SPBM in the core and IGMP on the access (Layer 3 VSN with IP multicast over SPBM).

For more information, see SPBM design guidelines on page 73.

Autogenerated CFM MEP and MIP levels

Release 3.1 simplifies Connectivity Fault Management (CFM) configuration with autogenerated CFM. With the simplified autogenerated CFM, you use the commands **`cfm spbm enable`** and **`cfm cmac enable`** and the device creates a default MD, MA, MEPs and MIPs.

You do not have to configure explicit MEPs and MIPs and associate multiple VLANs with MEPs and MIPs. Now you can use the autogenerated CFM feature that creates a MEP and MIP at a specified level for every SPBM Backbone VLAN (B-VLAN) or C-VLAN.

In Release 3.1, CFM also extends the debugging of Layer 2 networks to Customer VLANs (C-VLANs).

- For SPBM B-VLANs, you can use either autogenerated or explicitly-configured CFM MEPs.
- For C-VLANs, you can only use autogenerated CFM MEPs.

Important:

Previous explicit CFM configurations of MDs, MAs, and MEPs on SPBM B-VLANs continue to function in this release. However, if you want to enable the new autogenerated commands, you must first remove the existing MEP and MIP on the SPBM B-VLANs. VSP 4000 only supports one type of MEP or MIP for each SPBM B-VLAN.

If you choose to explicitly configure CFM, you must configure an MD, MA, and MEP ID. You do not have to configure an MD, MA, MIPs and MEPs if you configured autogenerated CFM, which enables the device to create default MD, MA, MEPs, and MIPs.

For more information on CFM, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 4000*, NN46251–510.

IGMP Layer 2 querier

Beginning with Release 3.1, you can use Virtual Services Platform 4000 as a Layer 2 querier for a Layer 2 multicast network that does not have a multicast router. For more information, see IGMP Layer 2 querier on page 113.

IGMP virtualization

From Release 3.1 onward you can configure IGMP with the Virtual Routing and Forwarding (VRF) Lite feature. For more information, see Multicast and VRF-lite on page 107.

Layer 2 ping output

Release 3.1 updates the output for the **l2 ping ip-address**. The **l2 ping ip-address** command output now displays `path attempted`, which refers to the number of Layer 2 ping attempts.

Layer 2 ping and Layer 2 traceroute

In Release 3.1, you can use the **l2 ping** and **l2 traceroute** commands with C-VLANs. These commands provide a simpler command syntax than the standard Loopback Message and Linktrace Messages.

Trace commands

Release 3.1 adds new trace commands to provide additional troubleshooting procedures.

The following trace commands provide trace information for SPBM IS-IS:

- **trace spbm isis level <0–4>**
- **show trace spbm isis**

The following trace commands provide trace information form CFM:

- **trace cfm level <0–4>**
- **show trace cfm**

Layer 2 tracemroute

Release 3.1 adds the **l2tracemroute** command to trace the multicast tree for a given multicast flow.

Transparent UNI

Release 3.1 support Transparent UNI feature. A Transparent UNI (T-UNI) maps an entire port or MLT to an I-SID. On Transparent UNI, ISID untagged and tagged(802.1q) packets are Layer 2–switched and the forwarding decision is based only on the MAC Addresses; VLAN ID is not used.

You can map multiple ports to a T-UNI I-SID. Multiple ports of the same switch or of different switches can be mapped to same Transparent UNI I-SID.

L3 protocol support

- **OSPF and RIP**:

  Release 3.1 introduces support for the configuration of the Open Shortest Path First (OSPF) and the Routing Information Protocol (RIP) on the Avaya VSP 4000.

  OSPF is an Interior Gateway Protocol (IGP) that distributes routing information between routers that belong to a single autonomous system (AS). Intended for use in large networks, OSPF is a link-state protocol that supports IP subnets, Type of Service (TOS)-based routing, and tagging of externally-derived routing information.

  In routed environments, routers communicate with one another to track available routes. Routers can dynamically learn about available routes using the RIP. The Avaya VSP 4000 software implements standard RIP to exchange IP route information with other routers.

  For more information, see the *Avaya Virtual Services Platform 4000 Configuration — OSPF and RIP*, NN46251–506.

- **BGP** :

  Release 3.1 also introduces support for the configuration of Border Gateway Protocol (BGP) on the Avaya VSP 4000.

  The following operations are supported by BGP:

  - ipv4
  - 4 byte AS
  - Peer groups
  - Redistribution

  For more information, see the *Avaya Virtual Services Platform 4000 Configuration — BGP*, NN46251–507.

Other changes

There are no other changes to this document for release 3.1.

Chapter 3: Network design fundamentals

To efficiently and cost-effectively use Avaya Virtual Services Platform 4000, you must properly design your network. Use the information in this section to help you properly design the network. To design networks, you must consider

- reliability and availability
- platform redundancy
- desired level of redundancy

A robust network depends on the interaction between system hardware and software. System software can be divided into different functions as shown in the following figure.

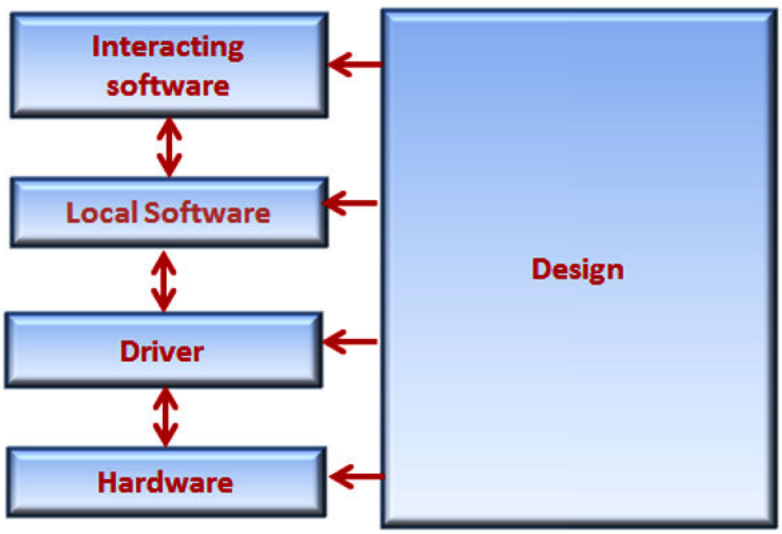

Figure 1: Hardware and software interaction

These levels are based on the software function. A driver is the lowest level of software that actually performs a function. Drivers reside on a single module and do not interact with other modules or external devices. Drivers are very stable.

Statically configured MultiLink Trunking (MLT) is a prime example of local software because it interacts with several modules within in the same device. No external interaction is needed, so you can easily test the function.

Interacting software is the most complex level of software because it depends on interaction with external devices. The Open Shortest Path First (OSPF) protocol is a good example of this software level. Interaction can occur between devices of the same type or with devices of other vendors than run a completely different implementation.

Based on network problem-tracking statistics, the following list is an approximate stability estimation model of a system that uses these components:

- Hardware and drivers represent a small portion of network problems.
- Local software represents a more significant share.
- Interacting software represents the vast majority of the reported issues.

Based on this model, network design attempts to off-load the interacting software level as much as possible to the other levels, especially to the hardware level. Avaya recommends that you follow these generic rules when you design networks:

1. Design networks as simply as possible.
2. Provide redundancy, but do not over-engineer your network.
3. Use a toolbox to design your network.
4. Design according to the product capabilities described in the latest release notes.
5. Follow the design rules provided in this document and also in the various configuration documents for the device.

Chapter 4: Hardware fundamentals and guidelines

This section provides general hardware guidelines to use this product in a network. Use the information in this section to help you during the hardware design and planning phase.

Platform considerations

This section provides VSP 4000 platform power and cooling considerations. You must properly power and cool your device, or nonoptimal operation can result.

Platform power

Virtual Services Platform 4000 power considerations

The Virtual Services Platform 4000 supports both AC and DC power supplies. One power supply is installed in the system.

The following redundant power supplies are supported. You can install the redundant power supply based on additional power requirements or to provide power redundancy.

Table 1: Redundant power supplies for the VSP 4000 platform

Part Number	Power supply (AC/DC)	Description
AL1905?08-E5*	300W AC redundant power supply	For use in the ERS 4626GTS, 4850GTS, VSP 4850GTS and WL8180, WL8180-16L wireless controllers. [EUED RoHS 5/6 compliant].
AL1905?21-E6*	Stackable 1000W AC POE+ power supply	For use in VSP 4850–GTS PWR+ and VSP 4450GSX-PWR+.
AL1905005-E5	Redundant 300W DC power supply	For use in the VSP 4850GTS-DC, ERS5698TFD, 5650TD, and 5632FD. (EUED RoHS 5/6 compliant). DC connector included.

*__Note__: The seventh character (?) of the switch order number must be replaced with the proper letter to indicate desired product nationalization. See the following for details:

“A”: No power cord included.

“B”: Includes European “Schuko” power cord common in Austria, Belgium, Finland, France, Germany, The Netherlands, Norway, and Sweden.

“C”: Includes power cord commonly used in the United Kingdom and Ireland.

“D”: Includes power cord commonly used in Japan.

“E”: Includes North American power cord.

“F”: Includes Australian power cord.

Device cooling

Virtual Services Platform 4000 — device cooling

The VSP 4000 platform has an in-built cooling module that is not removable. Each cooling module includes four fans providing cooling from front-to-back. There are three 12–volts fans to maintain optimal operating temperature inside the box. The fans are also speed controlled, based on the temperature in the box in order to minimize fan noise. Temperature sensors allow the fan speed controller to properly support the entire unit.

Caution:

Risk of electromagnetic interference:

This device is a Class A product. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users are required to take appropriate measures necessary to correct the interference at their own expense.

Avaya Virtual Services Platform 4000 Switch 4850GTS and 4850GTS-PWR+

The following table describes the regulatory AC power specifications for the Avaya Virtual Services Platform 4000 4850GTS and 4850GTS-PWR+ switches. The regulatory power specifications are based on the maximum rated capacity of the power supplies and are not based on typical power consumption which is lower.

Table 2: AC power specifications

	4850GTS	4850GTS-PWR+
Input Current:	5A/2.5A	12A/6A
Input Voltage (rms):	100 to 240VAC at 50 to 60 Hz	100 to 240VAC at 50 to 60 Hz

	4850GTS	4850GTS-PWR+
Power Consumption:	94.6W maximum	248W maximum
Thermal Rating:	323 BTU/Hr maximum	508 BTU/Hr maximum
Inrush Current:	40A maximum	70A maximum
Turn on Condition:	1 second maximum after application of AC power	1 second maximum after application of AC power
Important: 12 V output rise time, from 10 to 90 percent, must be the maximum of 50 ms and monotonic under all defined input and output conditions.		
Efficiency:	70 percent minimum	70 percent minimum

Avaya Virtual Services Platform power supply power specification

VSP 4850GTS-PWR+ supports two external field replaceable power supplies. You can install a secondary power supply to provide redundancy, load sharing, and add Power over Ethernet Plus (PoE+) power budget on PWR+ models.

VSP 4850GTS-PWR+ supports dual 54V 1000W Power over Ethernet Plus (PoE+) AC power supplies.

Figure 2: 1000W power supply

The Avaya VSP 4850GTS supports 300W AC power supplies.

Figure 3: 300W power supply

The 300W and 1000W AC power supplies use an IEC 60320 C16 AC power cord connector. The AC power cord is in close proximity to the hot air exhaust, and supports high operating temperatures

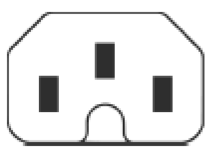

Figure 4: IEC 60320 C16 connector

Power over Ethernet Plus specifications

Model	Maximum PoE+ W	Average PoE+ W on 50 port models
Avaya VSP 4850GTS—PWR+ models	855W with one power supply 1855W with two power supplies	15.4W (802.3af) 17.8W (802.3.at) — 1 power supply 32.4W (802.3at) — 2 power supplies

- VSP 4850GTS-PWR+ can support 802.3af 15.4W on each port with one power supply installed. You can add a second power supply for redundancy.
- VSP 4850GTS-PWR+ can support 802.3at 32.4W on each port with two power supplies installed. PoE+ power reduces to an average of 17.8W on each port with one power supply.

Avaya DC power supply

The following table describes the DC power supply specifications for the Avaya Virtual Services Platform 4000.

Table 3: DC power supply specifications

Order Number	AL1905005-E5
Description	DC-DC-12V-300W
Output power	300W
Input voltage	48V DC
Input current	10A
Output voltage	12V DC
Output current	25A

Hardware compatibility

The following tables describe the Avaya Virtual Services Platform 4000 hardware.

Table 4: Hardware

VSP 4000 model	Description	Part number
VSP 4850GTS	• 48 10/100/1000 BaseTX RJ-45 ports • two SFP ports • two SFP+ ports • Base Software License • one field replaceable 300W PSU	EC4800A78-E6
	• Same content as EC4800A78-E6 with a EU power cord.	EC4800B78-E6
	• Same content as EC4800A78-E6 with a UK power cord.	EC4800C78-E6
	• Same content as EC4800A78-E6 with a JP power cord.	EC4800D78-E6

VSP 4000 model	Description	Part number
	• Same content as EC4800A78-E6 with a NA power cord.	EC4800E78-E6
	• Same content as EC4800A78-E6 with a EU power cord.	EC4800F78-E6
VSP 4850GTS-PWR+	• 48 10/100/1000 802.3at PoE+ • two SFP ports • two SFP+ ports • Base Software License • one field replaceable 1000W PSU	EC4800A88-E6
	• Same content as EC4800A88-E6 with a EU power cord.	EC4800B88-E6
	• Same content as EC4800A88-E6 with a UK power cord.	EC4800C88-E6
	• Same content as EC4800A88-E6 with a JP power cord.	EC4800D88-E6
	• Same content as EC4800A88-E6 with a NA power cord.	EC4800E88-E6
	• Same content a EC4800A88-E6 with a AU power cord.	EC4800F88-E6
VSP 4850GTS DC	• 48 10/100/1000 Base TX RJ-45 ports • two shared SFP ports • two 10GE SFP+ ports • one field replaceable 300W DC PSU	EC4800078-E6

Supported optical devices

Use optical devices to achieve high bit-rate communications and long transmission distances. The following section describes the supported optical devices on the VSP 4000 system.

Caution:

Avaya recommends that you use Avaya branded SFP and SFP+ transceivers as they undergo extensive qualification and testing. Avaya is not responsible for any problems that arise from using non-Avaya branded SFP and SFP+ transceivers.

Small form factor (SFP) pluggable transceivers

SFPs are hot-swappable input and output enhancement components designed for use with Avaya products to allow gigabit Ethernet ports to link with other gigabit Ethernet ports over various media types.

You can use various SFP (1Gb/s) and SFP+ (10Gb/s) to attain different line rates and reaches. The following table describes the SFPs including the reach provided by various SFPs.

This table is informational only—not all Avaya Ethernet switching and routing products support all the SFPs listed here.

For more information about SFP and SFP+ transceivers, including technical specifications and installation instructions, see *Installing Transceivers and Optical components on the Avaya Virtual Services Platform 4000*, NN46251-301.

Important:

The attainable cable length can vary depending on the quality of the fiber optic cable used.

Table 5: Compatible SFPs

Hardware	Description	Part number
1000BASE-T (RJ-45) SFP	Gigabit Ethernet, RJ-45 connector Range up to 100 m	AA1419043-E6
1000BASE-SX (LC) DDI	850 nm, Gigabit Ethernet, duplex LC connector Range up to 500 m	AA1419048-E6
1000BASE-LX (LC) DDI	1310 nm, Gigabit Ethernet, duplex LC connector Range up to 10 km over single mode fiber (SMF); up to 550 m over multimode fiber (MMF)	AA1419049-E6
1000BASE-XD DDI	1310 nm, Gigabit Ethernet, duplex LC connector	AA1419050-E6
	1550 nm, Gigabit Ethernet, duplex LC connector	AA1419051-E6
1000BASE-ZX DDI	1550 nm, Gigabit Ethernet, duplex LC connector	AA1419052-E6
1000BASE-XD CWDM (LC)	1470 nm to 1610 nm Range up to 40 km	AA1419053-E6 to AA1419060–E6

Hardware	Description	Part number
1000BASE-ZX CWDM (LC)	1470 nm to 1610 nm Range up to 70 km	AA1419061-E6 to AA1419068-E6
1000BASE-BX10 DDI SFP	1310 nm, single fiber LC Range up to 10 km	AA1419069-E6
1000BASE-BX10 DDI SFP	1490 nm, single fiber LC Range up to 10 km	AA1419070-E6
1000BASE-EX DDI SFP	1550 nm Range up to 120 km	AA1419071-E6
1000BASE-BX40 bidirectional SFP	1310 nm, single fiber LC Range up to 40 km	AA1419076–E6
1000BASE-BX40 bidirectional SFP	1490 nm, single fiber LC Range up to 40 km	AA1419077–E6
100BASE-FX SFP	1310 nm, LC connector	AA1419074-E6

Small form factor (SFP+) pluggable plus transceivers

SFP+ transceivers are hot-swappable input and output enhancement components that allow 10 gigabit connections. All Avaya SFP+ transceivers use Lucent connectors (LC) to provide precision keying and low interface losses.

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

Table 6: Compatible SFP+s

Hardware	Description	Part number
10GBASE-LR/LW SFP+	1310 nm SMF Range up to 10 km	AA1403011-E6
10GBASE-ER/EW SFP+	1550 nm SMF Range up to 40 km	AA1403013-E6
10GBASE-SR/SW SFP+	850 nm Range up to 300 m	AA1403015-E6
10GBASE ZR/ZW SFP+	1550 nm 70km SMF	AA1403016-E6
10GBASE-ER CWDM SFP +	1470 to 1610 nm with a range up to 40 km	AA1403153-E6 to AA1403160-E6
10GBASE-LRM SFP+	220 m, 1260 to 1355 nm; 1310 nm nominal MMF	AA1403017-E6
10GBase-CX	4-pair twinaxial copper cable that plugs into the SFP+ socket and connects two 10 Gb ports. The supported lengths are 3 m, 5 m, and 10 m.	AA1403018-E6 to AA1403020-E6

Optical power considerations

When you connect the device to collocated equipment, ensure that enough optical attenuation exists to avoid overloading the receivers of each device. You must consider the minimum attenuation requirement based on the specifications of third-party equipment. For more information about minimum insertion losses for Avaya optical products, see *Installing Transceivers and Optical components on the Avaya Virtual Services Platform 4000*, NN46251-301.

Dispersion considerations for long reach

Precise engineering of transmission links is difficult; specifications and performance are often unknown, undocumented, or impractical to measure before equipment installation. Moreover, the skills required to perform rigorous link budget analysis are extensive. Fortunately, a simple, straightforward approach can assure robust link performance for most optical fiber systems in which you use Avaya switches and routers.

This method uses an optical power budget, the difference between transmitter power and receiver sensitivity, to determine whether the installed link can operate with low bit error ratio for extended periods. The power budget must accommodate the sum of link loss (that is, attenuation), dispersion, and system margin, described in the following paragraphs.

Link losses are the sum of cabled fiber loss, splices, and connectors, often with an allocation for additional connectors. Cabled fiber loss is wavelength and installation-dependant, and is typically in the range of 0.20 to 0.5 dB/km. See the cable plant owner or operator for specifications of the cable you use, particularly if the available system margin is unsatisfactory. Engineered links require precise knowledge of the cable plant.

For long, high bit rate systems, pulse distortion, caused by the transmitter laser spectrum interaction with fiber chromatic dispersion, reduces receiver sensitivity. Transceivers for long reach single mode fiber systems have an associated maximum dispersion power penalty (DPP_{max}) specification, which applies to G.652 (dispersion unshifted) single mode fiber and the rated transceiver reach. The actual power penalty that you must use is

DPP_{budget} = [link length(km) / transceiver max reach (km)] * DPP_{max}

For example, if an 80 km transceiver is specified as having DPP < 3 dB, and if the actual link length will be 40 km, DPP_{budget} is one-half the maximum, or 1.5 dB.

Link operating margins are sometimes allocated for impairments such as aging, thermal, or other environmental effects. Due to the potentially large number of factors that can degrade performance, you can usually rely on statistics to represent these factors as a single margin value, in dB, to cover all effects. Margin is life and design-dependent, but is typically 3.5 to 4.5 dB, minimum. Whether you require additional margin depends on the details, such as whether actual or specified transmitter power and receiver sensitivity are used. Avaya specifications represent worst-case values.

The sum of margin, dispersion power penalty, and passive cable plant losses must be less than the available power budget. Alternatively, if you calculate available power margin as the difference between available budget and the sum of losses and dispersion, the margin can be more or less than required, which determines whether additional consideration is needed. If the power budget is exceeded or margin is insufficient, you can either use a transceiver rated for longer distance operation, or calculate budget and losses using actual values rather than specified limit values. Either method can improve link budget by 4 to 5 dB or more.

10/100BASE-X and 1000BASE-TX reach

The following table lists maximum transmission distances for 10/100BASE-X and 1000BASE-TX Ethernet cables.

Table 7: Maximum cable distances

	10BASE-T	100BASE-TX	1000BASE-TX
IEEE standard	802.3 Clause 14	802.3 Clause 21	802.3 Clause 40
Date rate	10 Mb/s	100 Mb/s	1000 Mb/s
Cat 5 UTP distance	100 m	100 m	100 Ω, 4 pair: 100 m

10/100/1000BASE-TX Auto-Negotiation recommendations

Auto-Negotiation lets devices share a link and automatically configures both devices so that they take maximum advantage of their abilities. Auto-Negotiation uses a modified 10BASE-T link integrity test pulse sequence to determine device ability.

The Auto-Negotiation feature allows the devices to switch between the various operational modes in an ordered fashion and allows management to select a specific operational mode. The Auto-Negotiation feature also provides a parallel detection (also called autosensing) function to allow the recognition of 10BASE-T, 100BASE-TX, 100BASE-T4, and 1000BASE-TX compatible devices, even if they do not support Auto-Negotiation. In this case, only the link speed is sensed; not the duplex mode. Avaya recommends the Auto-Negotiation configuration as shown in the following table, where A and B are two Ethernet devices.

Table 8: Recommended Auto-Negotiation configuration on 10/100/1000BASE-TX ports

Port on A	Port on B	Remarks	Recommendations
Auto-Negotiation enabled	Auto-Negotiation enabled	Ports negotiate on highest supported mode on both sides.	Avaya recommends that you use this configuration if both ports support Auto-Negotiation mode.

Port on A	Port on B	Remarks	Recommendations
Full-duplex	Full-duplex	Both sides require the same mode.	Avaya recommends that you use this configuration if you require full-duplex, but the configuration does not support Auto-Negotiation.

Auto-Negotiation cannot detect the identities of neighbors or shut down misconnected ports. Upper-layer protocols perform these functions.

Auto MDIX

Automatic medium dependent interface crossover (Auto-MDIX) automatically detects the need for a straight-through or crossover cable connection and configures the connection appropriately. This removes the need for crossover cables to interconnect switches and ensures either type of cable can be used. The speed and duplex setting of an interface must be set to auto for Auto-MDIX to operate correctly.

CANA

Use Custom Auto-Negotiation Advertisement (CANA) to control the speed and duplex settings that the interface modules advertise during Auto-Negotiation sessions between Ethernet devices. Modules can only establish links using these advertised settings, rather than at the highest common supported operating mode and data rate.

Use CANA to provide smooth migration from 10/100 Mb/s to 1000 Mb/s on host and server connections. Using Auto-Negotiation only, the switch always uses the fastest possible data rates. In limited-uplink-bandwidth scenarios, CANA provides control over negotiated access speeds, and improves control over traffic load patterns.

You can use CANA only on 10/100/1000 Mb/s RJ-45 ports. To use CANA, you must enable Auto-Negotiation.

Important:

If a port belongs to a Multilink Trunking (MLT) group and you configure CANA on the port (that is, you configure an advertisement other than the default), then you must apply the same configuration to all other ports of the MLT group (if they support CANA).

If a 10/100/1000 Mbit/s port that supports CANA is in a MLT group that has 10/100BASE-TX ports, or any other port type that does not support CANA, then use CANA only if it does not conflict with MLT abilities.

Chapter 5: Optical routing design

The Avaya optical routing system uses coarse wavelength division multiplexing (CWDM) in a grid of eight optical wavelengths. Use the Avaya optical routing system to maximize bandwidth on a single optical fiber. This section provides optical routing system information that you can use to help design your network.

Optical routing system components

Small form factor pluggable (SFP) transceivers transmit optical signals from Gigabit Ethernet (GbE) ports to multiplexers in a passive optical shelf.

Multiplexers combine multiple wavelengths traveling on different fibers onto a single fiber. At the receiver end of the link, demultiplexers separate the wavelengths and route them to different fibers, which terminate at separate CWDM devices. The following figure shows multiplexer and demultiplexer operations.

Important:

For clarity, the following figure shows a single fiber link with signals traveling in one direction only. A duplex connection requires communication in the reverse direction as well.

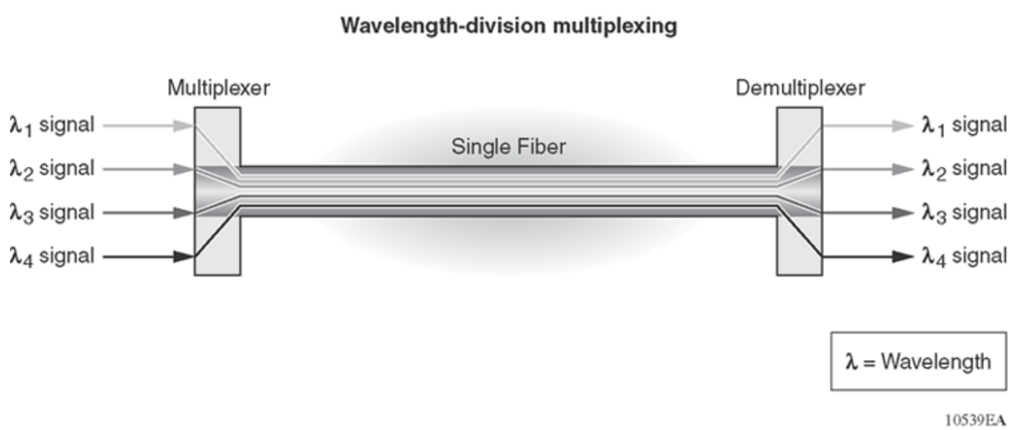

Figure 5: Wavelength division multiplexing

The Avaya optical routing system supports both ring and point-to-point configurations. The optical routing system includes the following parts:

- CWDM SFPs
- Optical add/drop multiplexers (OADM)

- Optical multiplexer/demultiplexers (OMUX)
- Optical shelf to house the multiplexers

OADMs drop or add a single wavelength from or to an optical fiber.

For the list of supported optical devices on the VSP 4000 platform for the current release, see Supported optical devices on page 24.

Chapter 6: Platform redundancy

This section includes recommendations to provide a fault tolerant platform.

Power redundancy

The Avaya VSP 4850GTS-PWR+ model supports dual 54V 1000W Power over Ethernet Plus (PoE+) AC power supplies. This model supports two external field replaceable power supplies. You can install a secondary power supply to provide redundancy, load sharing, and add Power over Ethernet Plus (PoE+) power budget on PWR+ models.

The 1000W AC power supplies use an IEC 60320 C16 AC power cord connector. The AC power cord is in close proximity to the hot air exhaust, and supports high operating temperatures.

Table 9: Power over Ethernet Plus specifications

Model	Maximum PoE+ W	Average PoE+ W on 50 port models
Avaya VSP 4850GTS — PWR+ models	• 855W with one power supply • 1855W with two power supplies	• 15.4W (802.3af) • 17.8W (802.3.at) — 1 power supply • 32.4W (802.3at) — 2 power supplies

The VSP 4850GTS-PWR+ can support 802.3af 15.4W on each port with one power supply installed. PoE+ power reduces to an average of 17.8W on each port with one power supply. It can support 802.3at 32.4W on each port with two power supplies installed.

Important:

Avaya recommends that power supplies use the same input voltage. Do not operate the chassis with power supplies under different input voltage conditions. The product can function but Avaya does not support this configuration.

Input/output port redundancy

You can protect I/O ports using a link aggregation mechanism. MultiLink Trunking (MLT), which is compatible with 802.3ad static, provides a load sharing and failover mechanism to protect against module, port, fiber, or complete link failures.

You can use MLT with Link Access Control Protocol (LACP) disabled or use LACP enabled by itself.

Configuration redundancy

You can define primary and backup configuration file paths. This configuration protects against system failures. For example, the primary path can point to system flash memory and the backup path to the external Compact Flash card.

Link redundancy

Provide physical and link layer redundancy to eliminate a single point of failure in the network. For more information, see Link redundancy on page 35.

Chapter 7: Link redundancy

You can build link redundancy into your network to

- help eliminate a single point of failure in your network (provide physical and link layer redundancy)
- prevent a service interruption caused by a faulty link (provide link layer redundancy)

This chapter explains the following design options that you can use to achieve link redundancy (provide alternate data paths) :

- physical layer redundancy
- Multilink Trunking
- 802.1ad-based link aggregation

Physical layer redundancy

To ensure that a faulty link does not cause a service interruption, you can provide physical layer redundancy in your network.

You can also configure the platform to detect link failures with, for example:

- Remote fault indication
- VLACP

Gigabit Ethernet and remote fault indication

The 802.3z Gigabit Ethernet (GbE) standard defines remote fault indication (RFI) as part of the Auto-Negotiation function.

Because RFI is part of the Auto-Negotiation function, if you disable Auto-Negotiation, you automatically disable RFI.

The stations on both ends of a fiber pair use RFI to inform one another after a problem occurs on one of the fibers.

Avaya recommends that you enable Auto-Negotiation on GbE links when the devices on both ends of a fiber link support Auto-Negotiation because, without RFI support, if one of two unidirectional fibers that form the connection between the two platforms fails, the transmitting side cannot determine that the link is broken in one direction (see the following figure).

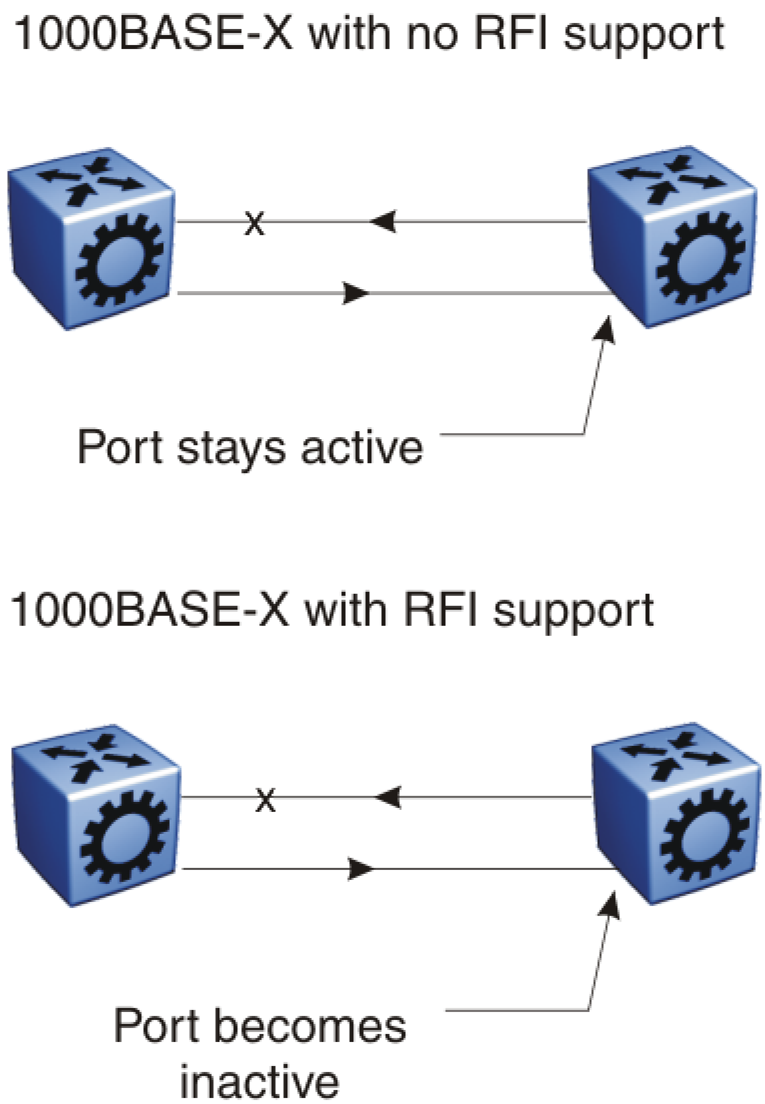

Figure 6: 1000BASE-X RFI

End-to-end fault detection and VLACP

Because remote fault indication (RFI) terminates at the next Ethernet hop, the device that uses only RFI cannot determine failures on an end-to-end basis over multiple hops.

However, you can use Virtual Link Aggregation Control Protocol (VLACP) to provide an end-to-end failure detection mechanism. You can configure VLACP on a port and enable it over single links or Multilink trunks (MLT).

You can use VLACP with MLT to enhance its capabilities and provide quick failure detection. With VLACP, the device can detect far-end failures, which permits MLT to failover properly when end-to-end connectivity is not guaranteed for some links in an aggregation group.

To minimize network outages, you can also use VLACP to switch traffic around entire network devices before Layer 3 protocols detect a network failure.

VLACP is an extension of the Link Aggregation Control Protocol (LACP) but LACP and VLACP are independent features.

VLACP does not perform link aggregation, it detects end-to-end link failures.

VLACP periodically checks the end-to-end health of a point-to-point connection and it uses the hello mechanism of LACP to periodically send hello packets to ensure end-to-end communication.

If VLACP does not receive hello packets it transitions to a failure state, which indicates a service provider failure, and the port is disabled. The system sends VLACP trap messages to the management stations if the VLACP state changes. If the failure is local, the system generates only port linkdown or port linkup traps.

VLACP works for port-to-port communications only where a guarantee exists for a logical port-to-port match through the service provider.

VLACP does not work for port-to-multiport communications where no guarantee exists for a point-to-point match through the service provider.

Example:

When the enterprise networks connect the aggregated Ethernet trunk groups through a service provider network connection, far-end failures cannot be signaled with Ethernet-based functions that operate end-to-end through the service provider network. The Multilink trunk (between enterprise switches S1 and S2) extends through the service provider network.

The following figure shows an MLT that operates with VLACP. VLACP can operate end-to-end, but you can also use it in a point-to-point link.

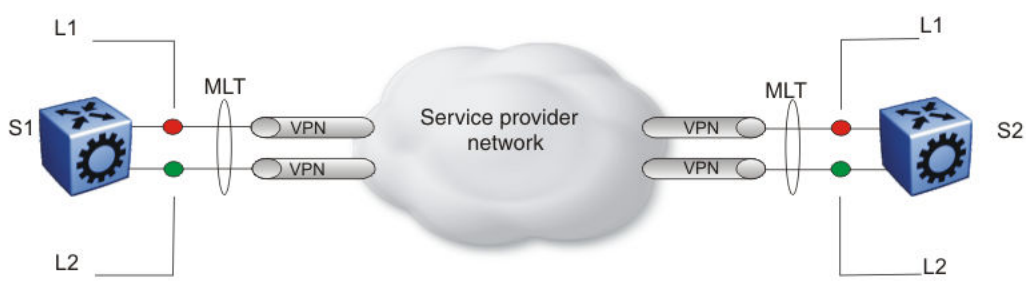

Figure 7: Problem description (1 of 2)

In the following figure, if the L2 link on S1 (S1/L2) fails, the link-down failure is not propagated over the service provider network to S2 and S2 continues to send traffic over the failed S2/ L2 link.

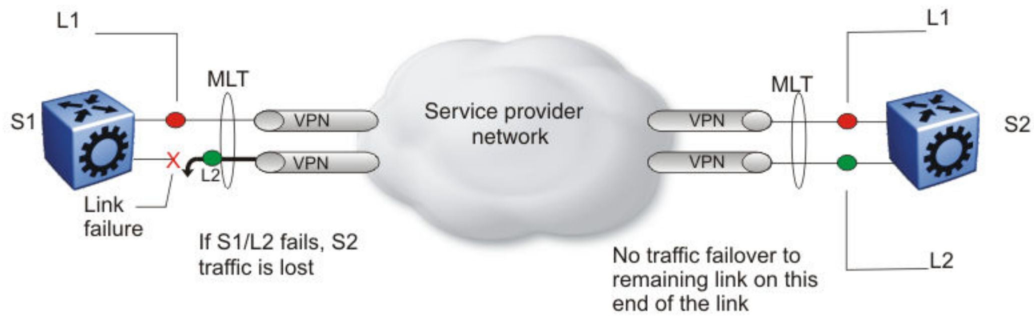

Figure 8: Problem description (2 of 2)

However, if you use VLACP to detect far-end failures and allow MLT to failover when end-to-end connectivity is not guaranteed for links in an aggregation group, then VLACP prevents the failure scenario in the preceding figure.

Avaya recommends that you use the following guidelines for VLACP implementation:

- Do not use VLACP on configured LACP MLTs because LACP provides the same functionality as VLACP for link failure. Virtual Services Platform 4000 does not support VLACP and LACP on the same link.
- Use the following best practice standard settings for VLACP :
    - a short timer—no less than 500 milliseconds (ms)
    - a time-out scale of 5

    Tip:

    The Avaya Virtual Services Platform 4000 supports both faster timers and lower time-out scales, but if VLACP flapping occurs, increase the short timer and the time-out scale to their recommended values: 500 and 5, respectively. Although the software configuration supports VLACP short timers of less than 30 ms, the platform does not support using values less than 30 ms in practice. The shortest (fastest) supported VLACP timer is 30 ms with a timeout of 3, which achieves sub-100 ms failover.
- Do not configure VLACP timers to less than 100 ms if you plan to use a Layer 3 core with Equal Cost Multipath (ECMP).

    Note:

    This recommendation assumes a combination of basic Layer 2 and Layer 3 with Open Shortest Path First (OSPF). If you have more complex configurations, you can require higher timer values.
- Ensure that the VLACP configuration at the port level is consistent, that both sides of the point-to-point connection are either enabled or disabled. If a VLACP-enabled port does not receive a VLACP protocol data unit (PDU), it enters the disabled state. However, occasions exist when a VLACP-enabled port does not receive a VLACP PDU but remains

in the forwarding state. You can avoid this situation with consistent port level VLACP configuration.

- Configure VLACP on an individual port basis.

  The port can be either an individual port or an MLT member. Each VLACP-enabled port periodically sends VLACP PDUs. This action allows the exchange of VLACP PDUs from an end-to-end perspective. If a particular link does not receive VLACP PDUs, the platform shuts the link down after the expiry timeout occurs (timeout scale x periodic time). As a result of this action the ports stay in a disabled state.

Multilink Trunking

Use MLT to provide link layer redundancy. You can use MLT to provide alternate paths around failed links. When you configure MLT links, consider the following information:

- The device supports 24 MLT aggregation groups.
- Up to 8 ports can belong to a single MLT group.

MLT and LACP groups and port speed

Ensure that all ports that belong to the same MLT or LACP group use the same port speed, for example, 1 Gb/s, even if you use Auto-Negotiation. The software does not enforce this requirement. Avaya recommends that you use Custom Auto-Negotiation Advertisement (CANA) to ensure proper speed negotiation in mixed-port type scenarios.

To maintain Link Aggregation Group (LAG) stability during failover, use CANA: configure the advertised speed to be the same for all LACP links. For 10/100/1000 ports, ensure that CANA uses one particular setting, for example, 1000-full or 100-full. Otherwise, a remote device can restart Auto-Negotiation and the link can use a different capability.

Each port must use only one speed and duplex mode; all links in Up state are guaranteed to have the same capabilities. If you do not use Auto-Negotiation and CANA, you must use the same speed and duplex mode settings on all ports of the MLT.

Platform-to-platform MLT link recommendations

Avaya recommends that you connect physical connections in platform-to-platform MLT and link aggregation links in a specific order. To connect an MLT link between two platforms, connect the lower number port on one platform with the lower number port on the other platform. For example, to establish an MLT platform-to-platform link between ports 1/1 and 1/4 on platform A with ports 1/1 and 1/4 on platform B, do the following:

- Connect port 1/1 on platform A to port 1/1 on platform B
- Connect port 1/4 on platform A to port 1/4 on platform B

In Virtual Services Platform 4000, brouter ports do not support MLT. You cannot use brouter ports to connect two platforms with an MLT. An alternative is to use a VLAN. This configuration option provides a routed VLAN with a single logical port or MLT. For more information on MLT configuration, see *Avaya Virtual Services Platform 4000 Configuration — Link Aggregation and MLT*, NN46251–503.

MLT and spanning tree protocols

The implementation of 802.1w (Rapid Spanning Tree Protocol—RSTP) and 802.1s (Multiple Spanning Tree Protocol—MSTP), provides a path cost calculation method. The following table provides the path costs associated with each interface type:

Table 10: Path cost for RSTP or MSTP mode

Link speed	Recommended path cost
Less than or equal 100 Kb/s	200 000 000
1 Mb/s	20 000 000
10 Mb/s	2 000 000
100 Mb/s	200 000
1 Gb/s	20 000
10 Gb/s	2000
100 Gb/s	200
1 Tb/s	20
10 Tb/s	2

802.3ad-based link aggregation

Link aggregation provides link layer redundancy. Use IEEE 802.3ad-based link aggregation (IEEE 802.3 2002 clause 43) to aggregate one or more links together to form LAGs to allow a MAC client to treat the LAG as if it were a single link. Use link aggregation to increase aggregate throughput of the interconnection between devices and provide link redundancy. LACP can dynamically add or remove LAG ports, depending on their availability and states.

Although IEEE 802.3ad-based link aggregation and MLT provide similar services, MLT is statically defined. By contrast, IEEE 802.3ad-based link aggregation is dynamic and provides additional functionality.

LACP and MLT

When you configure standards-based link aggregation, you must enable the aggregatable parameter. This configuration creates a one-to-one mapping between the LACP aggregator and the specified MLT.

A newly-created MLT or LAG adopts the VLAN membership of its member ports after the first port attaches to the aggregator associated with this LAG. After a port detaches from an aggregator, the port is deleted from the associated LAG port member list. After the last port member is deleted from the LAG, the LAG is deleted from all VLANs.

After you configure the MLT as aggregatable, you cannot add or delete ports or VLANs manually.

To enable tagging on ports that belong to a LAG, first disable LACP on the port, enable tagging on the port, and then enable LACP.

Important:

Enabling IS-IS is not supported on LACP based MLT.

LACP and spanning tree interaction

Only the physical link state or the LACP peer status affects the operation of LACP. After a link goes up and down, the LACP module receives notification. The spanning tree forwarding state does not affect the operation of the LACP module. LACP data units (LACPDU) can be sent even if the port is in spanning tree blocking state.

Configuration changes (such as speed, duplex mode, and so on) made to a LAG member port do not apply to all the member ports of the MLT. Instead, the changed port is removed from the LAG, and the corresponding aggregator and user is alerted.

In contrast to MLT, IEEE 802.3ad-based link aggregation does not require the system to replicate BPDUs over all ports in the trunk group.

LACP and minimum link

The minimum link function defines the minimum number of active links required for a LAG to remain in the forwarding state. You cannot configure the minimum link on Virtual Services Platform 4000. The minimum link value is always 1.

If the number of active links in a LAG is 0, the entire LAG is declared down and Virtual Services Platform 4000 informs the remote end of the LAG state by using an LACPDU.

Link aggregation group rules

Link aggregation is compatible with RSTP and MSTP. LAGs operate using the following rules:

- All ports in a LAG must operate in full-duplex mode.
- All ports in a LAG must use the same data rate.
- All ports in a LAG must be in the same VLANs.
- LAGs form using LACP.
- The platform supports a maximum of 128 LAGs.
- Each LAG supports a maximum of eight active links.

For LACP fundamentals and configuration procedures, see *Avaya Virtual Services Platform 4000 Configuration — Link Aggregation and MLT* NN46251-503.

Chapter 8: Layer 2 loop prevention

This section provides information about how to use bandwidth and network resources efficiently, and to prevent Layer 2 data loops. Use the information in this section to use loop prevention mechanisms.

Loop prevention and detection

In certain network designs, loops can form. For example, loops can form if you have incorrect configuration or cabling.

There are two solutions to detect loops: Loop Detect, and Simple Loop Prevention Protocol (SLPP). Both solutions perform the following functions:

- detect the loop
- automatically stop the loop
- determine on which port the loop is occurring
- shut down the port on which the loop is occurring

Avaya recommends the following loop prevention and recovery features in order of preference:

1. SLPP
2. Loop Detect with ARP-Detect activated

For more information about SLPP and loop detection, see *Avaya Virtual Services Platform 4000 Configuration — VLANs and Spanning Tree*, NN46251-500.

SLPP

Use SLPP to protect the network against Layer 2 loops. If you configure and enable SLPP, the switch sends a test packet to the VLAN. A loop is detected if the switch or a peer aggregation switch on the same VLAN receives the original packet. If the switch detects a loop, the switch disables the port. After the port is disabled, you must enable the port manually, or use port auto enable to reenable the port after a predefined interval.

Loops can be introduced into the network in many ways. One way is through the loss of a multilink trunk configuration caused by user error or malfunction. This scenario does not introduce a broadcast storm, but because all MAC addresses are learned through the looping ports, Layer 2 MAC learning is significantly impacted. Spanning tree protocols cannot always detect such a configuration issue, whereas SLPP reacts and disables the malfunctioning links, which minimizes the impact on the network.

SLPP configuration considerations and recommendations

SLPP uses an individual VLAN hello packet mechanism to detect network loops. Sending hello packets on an individual VLAN basis allows SLPP to detect VLAN-based network loops for untagged and tagged IEEE 802.1Q VLAN link configurations. You determine to which VLANs a switch sends SLPP test packets to. All port members of the SLPP-enabled VLAN replicate the packets.

Use the information in this section to understand the considerations and recommendations to configure SLPP in your network:

- You must enable SLPP packet receive on each port to detect a loop.
- SLPP test packets (SLPP-PDU) are forwarded for each VLAN.
- SLPP-PDUs are automatically forwarded on VLAN ports configured for SLPP.
- The SLPP-PDU destination MAC address is the switch MAC address, with the multicast bit set; the source MAC address is the switch MAC address.

  Note:

  VSP 4000 SLPP design is different from that of ERS 8800 SLPP. On the ERS 8800, the source MAC address is the switch *VLAN* MAC address.
- The SLPP-PDU is sent out as a multicast packet and is constrained to the VLAN on which it is sent.
- If an MLT port receives an SLPP-PDU, the port goes down.
- The originating CP receives the SLPP-PDU. All other switches regard the SLPP-PDU as a normal multicast packet, and forward it to the VLAN.
- SLPP is port-based; therefore, a port is disabled if it receives SLPP-PDU on one or more VLANs on a tagged port. For example, if the SLPP packet receive threshold is 5, a port is shut down if it receives 5 SLPP-PDU from one or more VLANs on a tagged port.
- The switch does not act on SLPP packets other than on the SLPP packets that it transmits.
- For square and full-mesh configurations that use a routed core, create a separate core VLAN. Enable SLPP on the core VLAN and the square or full-mesh links between switch clusters. This configuration detects loops created in the core, and loops at the edge do not affect core ports.
- You can tune network failure behavior by selecting the number of SLPP packets that must be received before a switch takes action.

Avaya recommends the values in the following table.

Table 11: SLPP recommended values

Parameter	Configuration
Primary switch	
Packet Rx threshold	5
Transmission interval	500 milliseconds (ms) (default)
Secondary switch	
Packet Rx threshold	50

Parameter	Configuration
Transmission interval	500 ms (default)

Loop Detect

Use the Loop Detect feature at the edge of a network to prevent loops. This feature detects whether the same MAC address appears on different ports. Loop Detect can disable a VLAN or a port. The Loop Detect feature can also disable a group of ports if it detects the same MAC address on two different ports five times in a configurable amount of time.

On a individual port basis, the Loop Detect feature detects MAC addresses that are looping from one port to other ports. After the switch detects a loop, it disables the port on which the MAC addresses were learned. Additionally, if Loop Detect finds a MAC address loop, it disables the MAC address for that VLAN.

ARP Detect

The ARP-Detect feature is an enhancement over Loop Detect to account for ARP packets on IP configured interfaces. For network loops that involve ARP frames on routed interfaces, Loop-Detect does not detect the network loop condition due to how ARP frames are copied to the CPU. Use ARP-Detect on Layer 3 interfaces. The ARP-Detect feature supports only the vlan-block and port-down options.

VLACP

This feature provides an end-to-end failure detection mechanism that prevents potential problems caused by misconfigurations in a switch cluster design.

Configure VLACP on an individual port basis. The system forwards traffic only across the uplinks when VLACP is operating correctly. You must configure the ports on each end of the link for VLACP. VLACP takes the point-to-point hello mechanism of LACP and uses it to periodically send PDU packets to ensure end-to-end reachability and provide failure detection, across a Layer 2 domain. If one end of the link does not receive the VLACP PDUs, it logically disables that port and no traffic passes. This action ensures that even if no link exists on the port at the other end, and if it is not processing VLACP PDUs correctly, no traffic is sent. This function alleviates potential black hole situations by only sending traffic to ports that are functioning properly.

You can reduce VLACP timers to 400 milliseconds between two Avaya Virtual Services Platform 4000 systems. The timer provides approximately one second failure detection and switchover. When you configure VLACP, you must configure both ends of the link with the same multicast MAC address and timers. Most products in the Avaya Ethernet switch and Ethernet routing switch line use the same timers, with the exception of the FastPeriodicTimer, which is 200ms on the Ethernet Routing Switch 8800, VSP 9000, and VSP 4000; and 500ms on all other switches.

SLPP example scenarios

The following examples illustrate some situations where layer 2 loops can occur and how SLPP prevents loops in those cases.

Scenario 1: VSP 4000 as an edge router

Scenario 1 demonstrates a triangular setup with ERS 8800 switches as IST peers, and VSP 4000 on the edge. From VSP 4000, there are four links that are part of the same MLT, with SLPP enabled on the VSP 4000 ports. Because the MLT ports are misconfigured, loops can occur. For example, port 1/1 on VSP 4000 can be part of the MLT, but on the ERS port, 2/1 is not part of the MLT, although they are on the same VLAN.

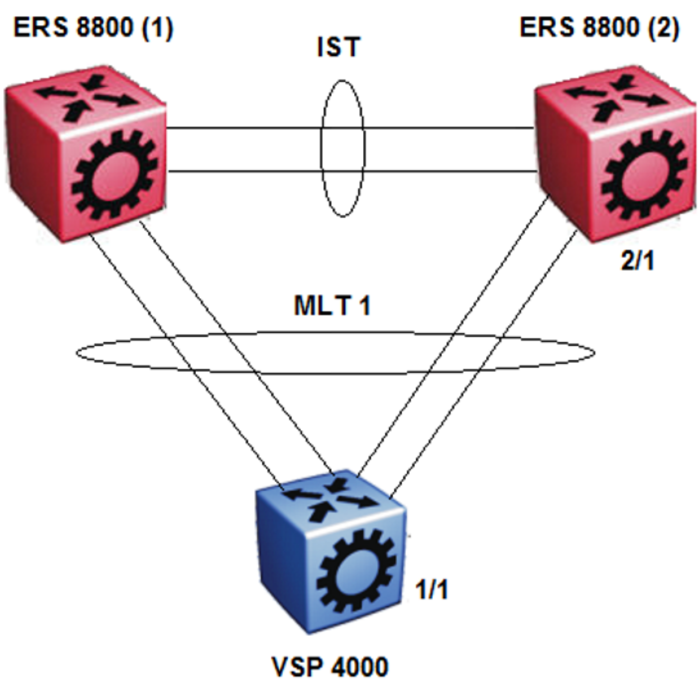

Figure 9: VSP 4000 as an edge router

SLPP PDUs are generated by VSP 4000. If there is a loop, then the SLPP PDUs return to port 1/1. After the threshold value is reached, SLPP brings the ports down.

Scenario 2: VSP 4000 as an edge router but with an additional link to the ERS 8800

Scenario 2 is similar to scenario 1 except that there is an additional link from ERS 8800 to VSP 4000 that is not part of MLT 1. The additional link is a member of the SLPP-enabled VLAN, and does not have to be directly connected from ERS 8800 to VSP 4000, but can be connected from other DUTs interconnected to the ERS 8800 and VSP 4000.

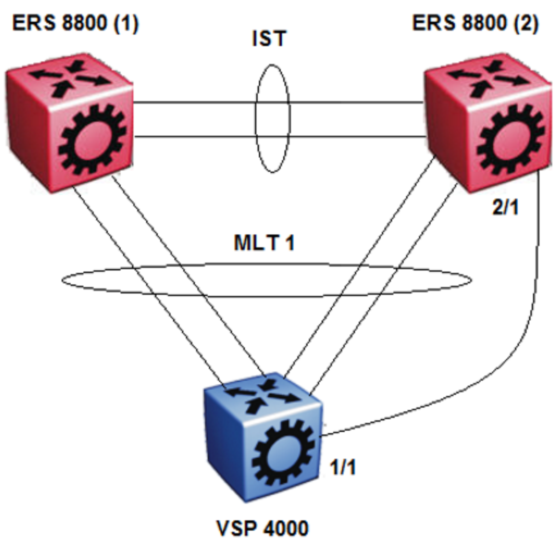

Figure 10: VSP 4000 as an edge router and with an additional link with ERS 8800

The SLPP PDUs generated by VSP 4000 return to the same DUT through the additional link. After the threshold value set on the SLPP enabled ports is reached, the ports go down.

Scenario 3: VSP 4000 as a BEB connected to an edge router

In scenario 3, VSP 4000 acts as a BEB and is connected to a BayStack device. SLPP is enabled on the UNI ports of the VSP 4000. Because the MLT ports are misconfigured, loops can occur. For example, port 10 on BayStack is part of the MLT, but on VSP 4000 port 1/1 is not part of the MLT, but both devices are on the same VLAN.

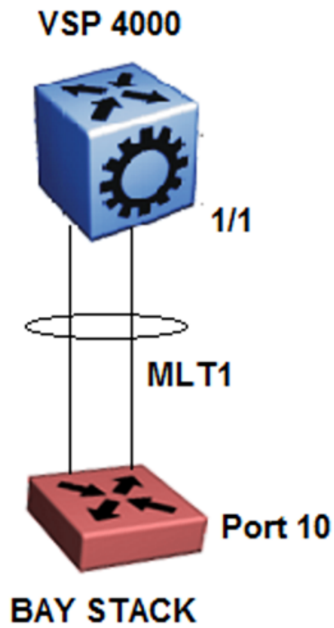

Figure 11: VSP 4000 as a BEB connected to an edge router

In this scenario, either SLPP or RSTP/MSTP can bring the ports down.

Scenario 4: Two VSP 4000 switches acting as BEBs

In scenario 4, there are two VSP 4000 devices that act as BEBs and are connected to each other through MLT, with two BayStack devices connected to each of the BEBs. The interface that connects the VSP 4000 interfaces is an ISIS interface with STP disabled. SLPP is enabled in the UNI ports of the VSP 4000. Because the link between the VSP 4000 uses ISIS interfaces, and STP is disabled on these interfaces, STP may not be able to detect the loop.

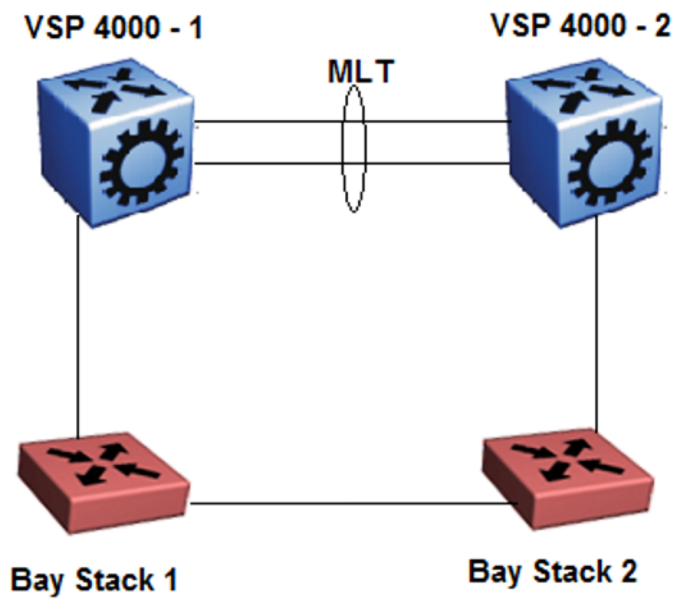

Figure 12: Two VSP 4000 switches acting as BEBs

The SLPP PDUs generated by the VSP 4000-1 return to itself through VSP 4000–2, Bay Stack 2, and Bay Stack 1. After reaching the threshold value, the SLPP brings the port down, therefore eliminating the loop.

Chapter 9: Spanning tree

Spanning tree prevents loops in switched networks. Virtual Services Platform 4000 supports Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP). This section describes issues to consider when you configure spanning tree protocols.

For more information about spanning tree protocols, see *Avaya Virtual Services Platform 4000 Configuration — VLANs and Spanning Tree*, NN46251-500.

Spanning tree and protection against isolated VLANs

Virtual Local Area Network (VLAN) isolation disrupts packet forwarding. Figure 13: VLAN isolation on page 52 shows the problem. Two VLANs (V1 and V2) connect four devices, and both VLANs are in the same STG. V2 includes three of the four devices, whereas V1 includes all four devices. After a spanning tree protocol detects a loop, it blocks the link with the highest link cost. In this case, the 100 Mbit/s link is blocked, which isolates a device in V2. To avoid this problem, either configure V2 on all four devices or use MSTP with a different Multiple Spanning Tree Instance (MSTI) for each VLAN.

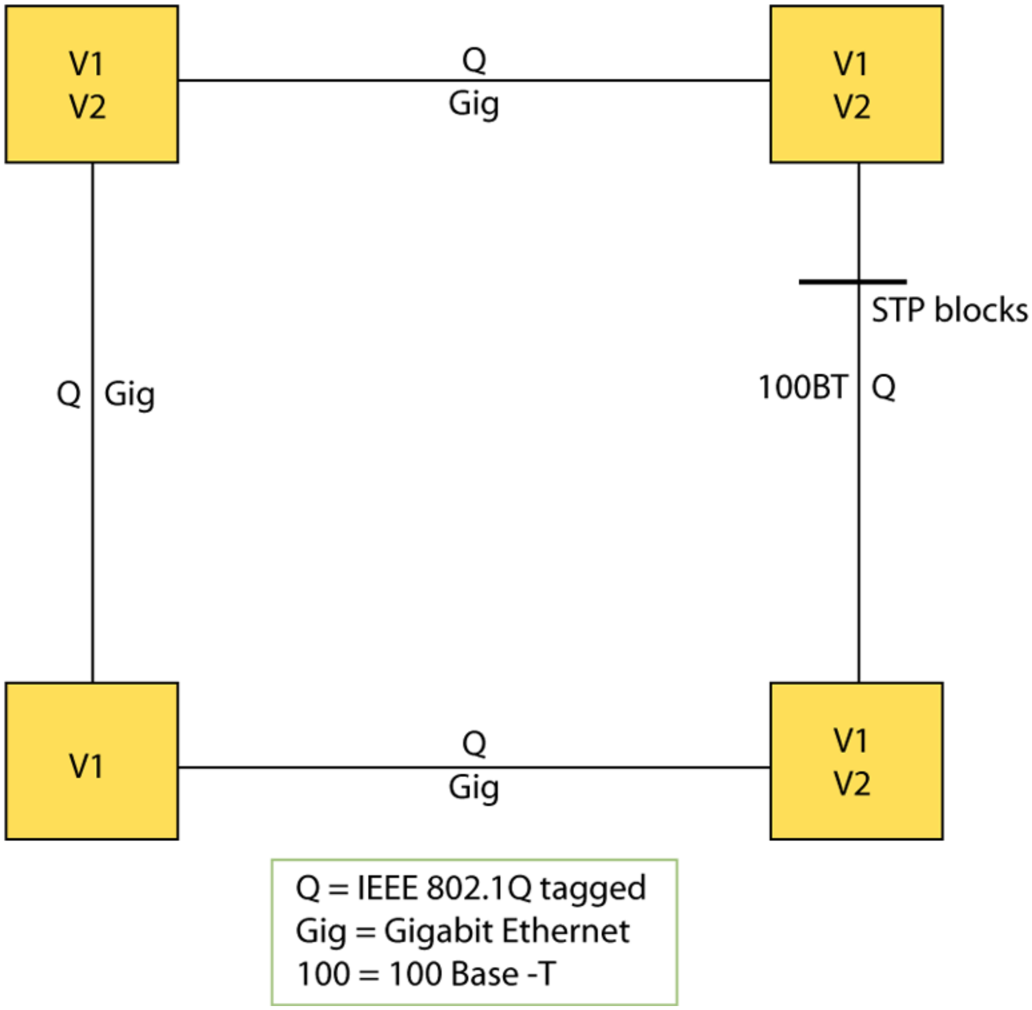

Figure 13: VLAN isolation

MSTP and RSTP considerations

The Spanning Tree Protocol (STP) provides loop protection and recovery, but it is slow to respond to a topology change in the network (for example, a dysfunctional link in a network). The RSTP (IEEE 802.1w) and MSTP (IEEE 802.1s) protocols reduce the recovery time after a network failure. RSTP and MSTP also maintain a backward compatibility with IEEE 802.1D. Typically, the recovery time of RSTP and MSTP is less than 1 second. RSTP and MSTP also reduce the amount of flooding in the network by enhancing the way that Topology Change Notification (TCN) packets are generated.

Use MSTP to configure MSTIs on the same switch. Each MSTI can include one or more VLANs.

In MSTP mode you can configure up to 64 instances. Instance 0 or Common and Internal Spanning Tree (CIST) is the default group, which includes default VLAN 1. Instances 1 to 63 are MSTIs.

RSTP and MSTP provide a global spanning tree parameter, called **`version`**, for backward compatibility with legacy STP. You can configure **`version`** to either STP-compatible mode, RSTP mode, or MSTP mode:

- An STP-compatible port transmits and receives only STP Bridge Protocol Data Units (BPDU). An RSTP or MSTP BPDU that the port receives in this mode is discarded.
- An RSTP or MSTP port transmits and receives only RSTP or MSTP BPDUs. If an RSTP or MSTP port receives an STP BPDU, it becomes an STP port. You must manually intervene to configure this port for RSTP or MSTP mode again. This process is called Port Protocol Migration.

You must be aware of the following recommendations before you implement MSTP or RSTP:

- The default mode is MSTP. A special boot configuration flag identifies the mode.
- You can lose your configuration if you change the spanning tree mode from MSTP to RSTP and the configuration file contains VLANs configured with MSTI greater than 0. RSTP only supports VLANs configured with the default instance 0.
- For best interoperability results, contact your Avaya representative.

Chapter 10: Layer 3 network design

This section describes Layer 3 design considerations that you need to understand to properly design an efficient and robust network.

VRF Lite

The Virtual Services Platform 4000 supports the Virtual Router Forwarding (VRF) Lite feature, which supports many virtual routers, each with its own routing domain. VRF Lite virtualizes the routing tables to form independent routing domains, which eliminates the need for multiple physical routers.

To use VRF Lite, you must use the Premier Software License.

For more information about VRF Lite, see *Avaya Virtual Services Platform 4000 Configuration — IP Routing*, NN46251-505.

VRF Lite route redistribution

Using VRF Lite, Virtual Services Platform 4000 can function as many routers; each VRF routing engine works independently. Normally, no route leak occurs between different VRFs. Use the route redistribution option to facilitate the redistribution of routes. VRFs can redistribute Open Shortest Path First (OSPF), Routing Information Protocol (RIP), Border Gateway Protocol (BGP), direct, and static routes.

If you enable route redistribution between two VRFs, ensure that the IP addresses do not overlap. The software does not enforce this requirement.

VRF Lite capability and functionality

On a VRF instance, VRF Lite supports the following protocols: IP, Internet Control Message Protocol (ICMP), Address Resolution Protocol (ARP), static routes, default routes, RIP, OSPF, external BGP (eBGP), route policies, Virtual Router Redundancy Protocol (VRRP), and the Dynamic Host Configuration Protocol/BootStrap Protocol relay agent.

The device uses VRF Lite to perform the following actions:

- partition traffic and data, and represent an independent router in the network
- provide virtual routers that are transparent to end-users
- support overlapping IP address spaces in separate VRFs
- support addresses that are not restricted to the assigned address space given by host Internet Service Providers (ISP)
- support eBGP

VRF Lite architecture examples

VRF Lite enables a router to act as many routers. This provides virtual traffic separation for each user and provides security. For example, you can use VRF Lite to

- provide different departments within a company with site-to-site connectivity as well as Internet access
- provide centralized and shared access to data centers.

The following figure shows how VRF Lite can emulate VPNs.

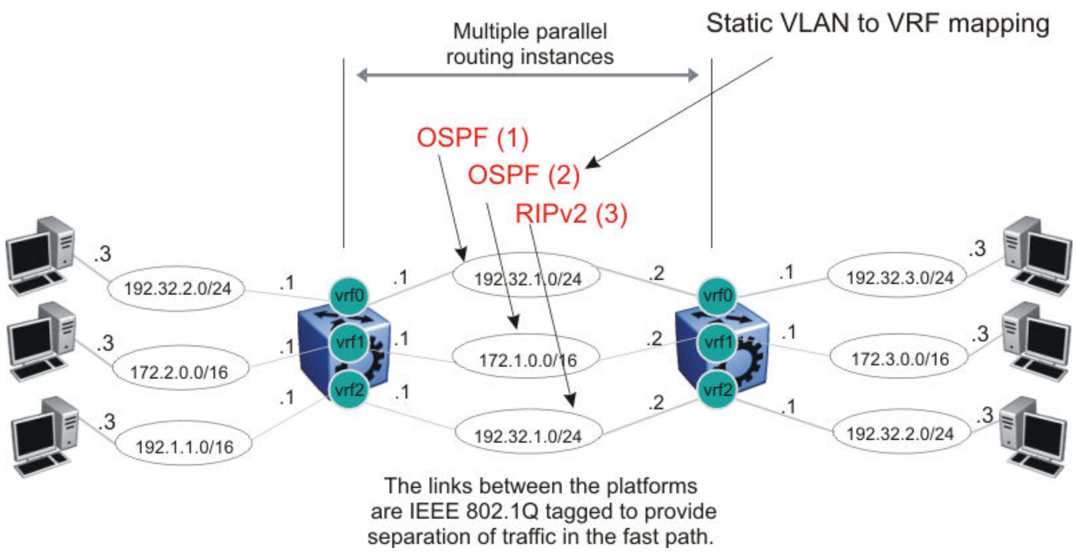

Figure 14: VRF Lite example

The following figure shows how VRFs can interconnect through an external firewall.

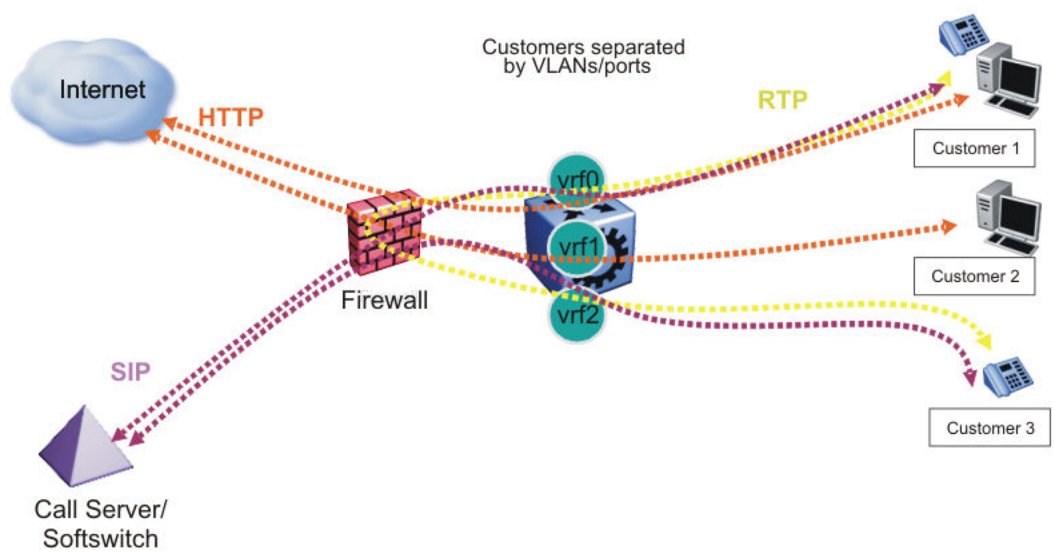

Figure 15: Inter-VRF forwarding based on external firewall

Although customer data separation into Layer 3 virtual routing domains is usually a requirement, sometimes customers must access a common network infrastructure. For example, they want to access the Internet, data storage, VoIP-PSTN, or call signaling services. To interconnect VRF instances, you can use an external firewall that supports virtualization, or use inter-VRF forwarding for specific services. Using the inter-VRF solution, you can use routing policies and static routes to inject IP subnets from one VRF instance to another, and filters to restrict access to certain protocols.

The following figure shows inter-VRF forwarding. In this solution, you can use routing policies to leak IP subnets from one VRF to another. You can use filters to restrict access to certain protocols. This configuration enables hub-and-spoke network designs, for example, for VoIP gateways.

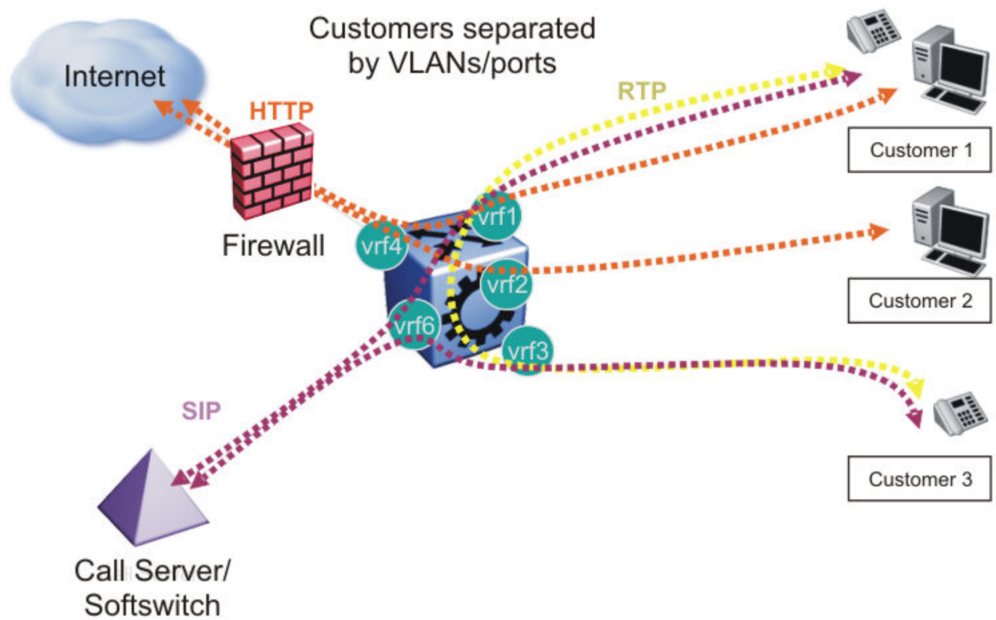

Figure 16: Inter VRF communication, internal inter-VRF forwarding

Virtual Router Redundancy Protocol

The Virtual Router Redundancy Protocol (VRRP) provides a backup router that takes over if a router fails, which is important if you must provide redundancy mechanisms.

VRRP guidelines

VRRP provides another layer of resiliency to your network design by providing default gateway redundancy for end users. If a VRRP-enabled router that connects to the default gateway fails, failover to the VRRP backup router ensures no interruption for end users who attempt to route from their local subnet.

Only the VRRP Master router forwards traffic for a given subnet. The backup VRRP router does not route traffic destined for the default gateway.

To allow both VRRP switches to route traffic, Virtual Services Platform 4000 has an extension to VRRP, the BackupMaster, that creates an active-active environment for routing. If you enable BackupMaster on the backup router, the backup router no longer switches traffic to the VRRP Master. Instead the BackupMaster routes all traffic received on the BackupMaster IP interface according to the switch routing table.

Figure 17: VRRP with BackupMaster

Avaya recommends that you stagger VRRP instances on a network or subnet basis. The following figure shows the VRRP Masters and BackupMasters for two subnets. For more information about how to configure VRRP using ACLI and Enterprise Device Manager (EDM), see *Avaya Virtual Services Platform 4000 Configuration — IP Routing*, NN46251-505.

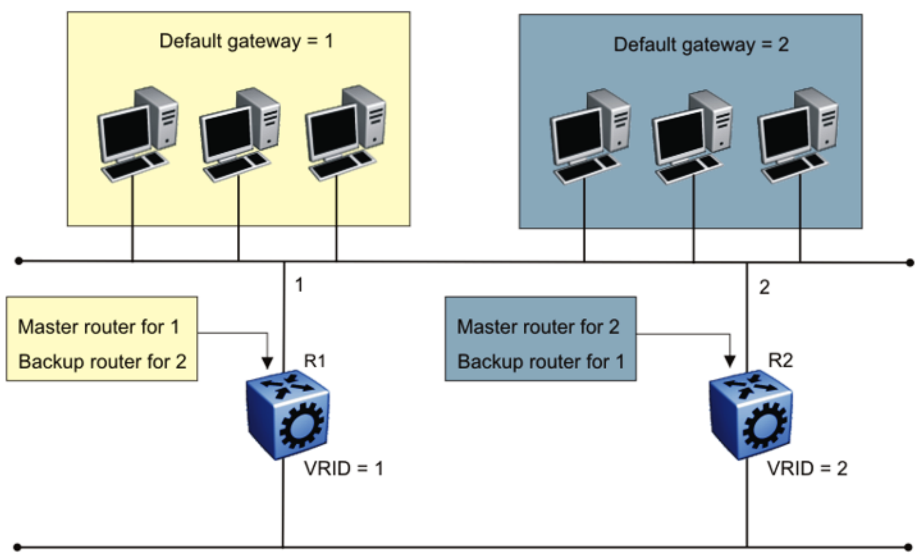

Figure 18: VRRP network configuration

The VRRP BackupMaster uses the VRRP standardized backup switch state machine. Thus, VRRP BackupMaster is compatible with standard VRRP.

Avaya recommends that you use the following best practices to implement VRRP:

- Do not configure the virtual address as a physical interface that is used on the routing switches. Instead, use a third address, for example:
    - Interface IP address of VLAN A on Switch 1 = x.x.x.2
    - Interface IP address of VLAN A on Switch 2 = x.x.x.3
    - Virtual IP address of VLAN A = x.x.x.1

  Note:

  Avaya does not support a VRRP virtual IP address to be the same as the local physical address of the device.
- Configure the VRRP hold down timer long enough that the Interior Gateway Protocol (IGP) routing protocol has time to converge and update the routing table. In some cases,

configuring the VRRP hold down timer to a minimum of 1.5 times the IGP convergence time is sufficient. For OSPF, Avaya recommends that you use a value of 90 seconds if you use the default OSPF timers.

- Implement VRRP BackupMaster for an active-active configuration (BackupMaster works across multiple switches that participate in the same VRRP domain.).
- Configure VRRP priority as 200 to configure VRRP Master.
- Stagger VRRP Masters between switches in the core to balance the load between switches.
- If you implement VRRP Fast, you create additional control traffic on the network and also create a greater load on the CPU. To reduce the convergence time of VRRP, the VRRP Fast feature allows the modification of VRRP timers to achieve subsecond failover of VRRP. Without VRRP Fast, normal convergence time is approximately 3 seconds.
- Do not use VRRP BackupMaster and critical IP at the same time. Use one or the other.
- When you implement VRRP on multiple VLANs between the same switches, Avaya recommends that you configure a unique VRID on each VLAN.

VRRP and spanning tree

Virtual Services Platform 4000 can use one of two spanning tree protocols. These include the Rapid Spanning Tree Protocol (RSTP) and the Multiple Spanning Tree Protocol (MSTP).

VRRP protects clients and servers from link or aggregation switch failures. Configure the network to limit the amount of time a link is down during VRRP convergence. The following figure shows two possible configurations of VRRP and spanning tree; configuration A is optimal and configuration B is not.

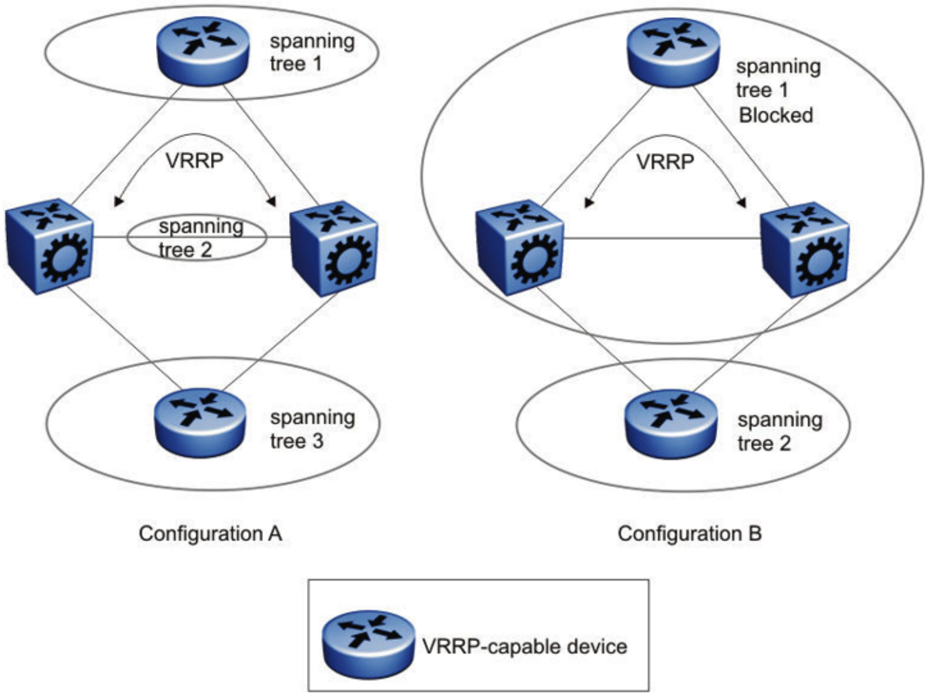

Figure 19: VRRP and STG configurations

In this figure, configuration A is optimal because VRRP convergence occurs within 2 to 3 seconds. In configuration A, three spanning tree instances exist and VRRP runs on the link between the two routers. Spanning tree instance 2 exists on the link between the two routers, which separates the link between the two routers from the spanning tree instances found on the other devices. All uplinks are active.

In configuration B, VRRP convergence takes between 30 and 45 seconds because it depends on spanning tree convergence. After initial convergence, spanning tree blocks one link (an uplink), so only one uplink is used. If an error occurs on the uplink, spanning tree reconverges, which can take up to 45 seconds. After spanning tree reconvergence, VRRP can take a few more seconds to failover.

VRRP and ICMP redirect messages

You can use VRRP and ICMP together. However, doing so can provide nonoptimal network performance.

Consider the network shown in the following figure. Traffic from the client on subnet 30.30.30.0, destined for the 10.10.10.0 subnet, is sent to routing switch 1 (VRRP Master). Routing switch 1 forwards this traffic on the same subnet to routing switch 2 where it is routed to the destination. With ICMP redirect enabled, for each packet received, routing switch 1 sends an ICMP redirect message to the client to inform it of a shorter path to the destination through routing switch 2.

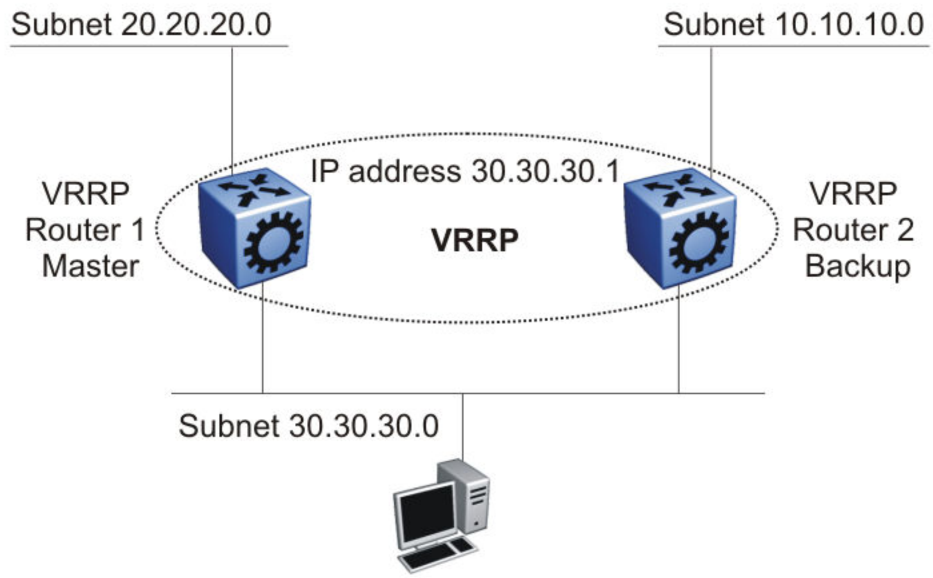

Figure 20: ICMP redirect messages

If network clients do not recognize ICMP redirect messages, disable ICMP redirect messages on Avaya Virtual Services Platform 4000 to avoid excessive ICMP redirect messages. Avaya recommends the network designs shown in the following figures.

Ensure that the routing path to the destination through both routing switches has the same metric to the destination. One hop goes from 30.30.30.0 to 10.10.10.0 through routing switch 1 and routing switch 2.

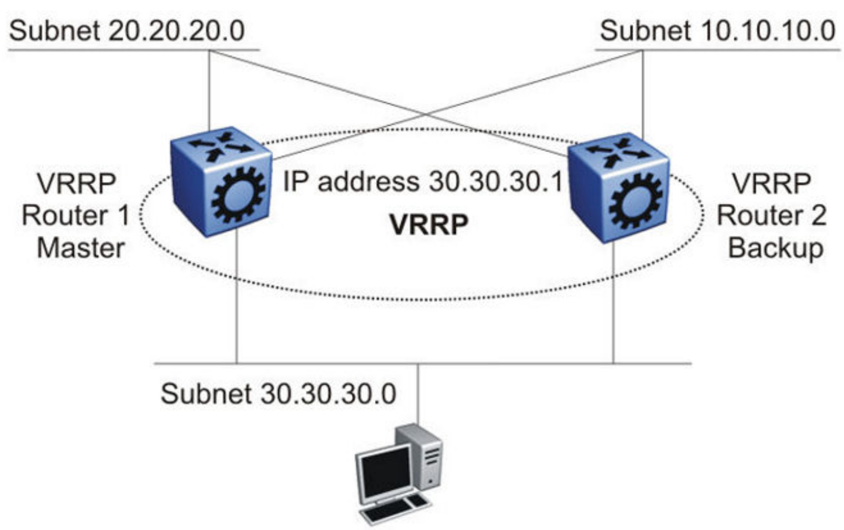

Figure 21: Avoiding excessive ICMP redirect messages without SMLT

Open Shortest Path First

Use OSPF to ensure that the switch can communicate with other OSPF-speaking routers. This section describes some general design considerations and presents a number of design scenarios for OSPF.

For more information about OSPF concepts and configuration, see *Avaya Virtual Services Platform 4000 Configuration — OSPF and RIP*, NN46251-506.

OSPF LSA limits

To determine OSPF link state advertisement (LSA) limits:

1. Use the command **`show ip ospf area`** to determine the LSA_CNT and to obtain the number of LSAs for a given area.
2. Use the following formula to determine the number of areas. Ensure the total is less than 16,000 (16K):

   $$\sum \mathrm{Adj}_N * \mathrm{LSA\_CNT}_N < \mathbf{16}\mathrm{k}$$

   N = 1 to the number of areas for each switch

   Adj_N = number of adjacencies for each Area N

   $\mathrm{LSA\_CNT}_N$ = number of LSAs for each Area N

For example, assume that a switch has a configuration of three areas with a total of 18 adjacencies and 1000 routes. This includes:

- 3 adjacencies with an LSA_CNT of 500 (Area 1)
- 10 adjacencies with an LSA_CNT of 1000 (Area 2)
- 5 adjacencies with an LSA_CNT of 200 (Area 3)

Calculate the number as follows:

3*500+10*1000+5*200=12.5K < 16K

This configuration ensures that the switch operates within accepted scalability limits.

OSPF design guidelines

Follow these additional OSPF guidelines:

- OSPF timers must be consistent across the entire network.
- Use OSPF area summarization to reduce routing table sizes.
- Use OSPF passive interfaces to reduce the number of active neighbor adjacencies.
- Use OSPF active interfaces only on intended route paths.

  Configure wiring closet subnets as OSPF passive interfaces unless they form a legitimate routing path for other routes.

- Minimize the number of OSPF areas for each switch to avoid excessive shortest path calculations.

  The switch executes the Djikstra algorithm for each area separately.
- Ensure that the OSPF dead interval is at least four times the OSPF hello interval
- Use MD5 authentication on untrusted OSPF links.
- Use stub or NSSA areas as much as possible to reduce CPU overhead.

OSPF and CPU utilization

After you create an OSPF area route summary on an area border router (ABR), the summary route can attract traffic to the ABR for which the router does not have a specific destination route. Enabling ICMP unreachable message generation on the switch can result in a high CPU utilization rate.

To avoid high CPU utilization, Avaya recommends that you use a black hole static route configuration. The black hole static route is a route (equal to the OSPF summary route) with a next-hop of 255.255.255.255. This configuration ensures that all traffic that does not have a specific next-hop destination route is dropped.

OSPF network design examples

You can use OSPF routing in the core of a network. For more information, see Layer 1, 2, and 3 design examples on page 139.

The following figure describes a simple implementation of an OSPF network: enabling OSPF on two switches (S1 and S2) that are in the same subnet in one OSPF area.

Figure 22: Example 1: OSPF on one subnet in one area

The routers in the preceding figure use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1, and the OSPF port uses an IP address of 192.168.10.1.
- S2 has an OSPF router ID of 1.1.1.2, and the OSPF port uses an IP address of 192.168.10.2.

The general method to configure OSPF on each routing switch is:

1. Enable OSPF globally.
2. Enable IP forwarding on the switch.

3. Configure the IP address, subnet mask, and VLAN ID for the port.
4. Disable RIP on the port, if you do not need it.
5. Enable OSPF for the port.

After you configure S2, the two switches elect a designated router (DR) and a backup designated router (BDR). They exchange hello packets to synchronize their link state databases (LSDB).

The following figure shows a configuration in which OSPF operates on three switches. OSPF performs routing on two subnets in one OSPF area. In this example, S1 directly connects to S2, and S3 directly connects to S2, but traffic between S1 and S3 is indirect, and passes through S2.

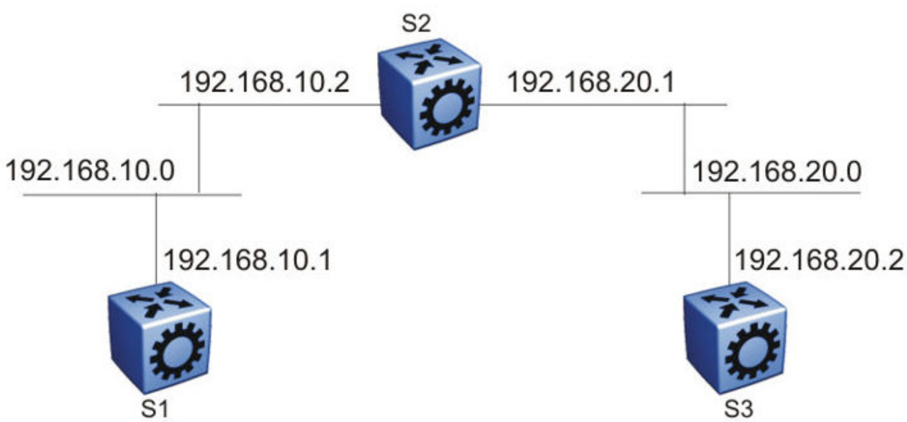

Figure 23: Example 2: OSPF on two subnets in one area

The routers in example 2 use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1, and the OSPF port uses an IP address of 192.168.10.1.
- S2 has an OSPF router ID of 1.1.1.2, and two OSPF ports use IP addresses of 192.168.10.2 and 192.168.20.1.
- S3 has an OSPF router ID of 1.1.1.3, and the OSPF port uses an IP address of 192.168.20.2.

The general method to configure OSPF on each routing switch is:

1. Enable OSPF globally.
2. Insert IP addresses, subnet masks, and VLAN IDs for the OSPF ports on S1 and S3, and for the two OSPF ports on S2. The two ports on S2 enable routing and establish the IP addresses related to the two networks.
3. Enable OSPF for each OSPF port allocated with an IP address.

After you configure all three switches for OSPF, they elect a DR and BDR for each subnet and exchange hello packets to synchronize their LSDBs.

The following figure shows an example where OSPF operates on two subnets in two OSPF areas. S2 becomes the ABR for both networks.

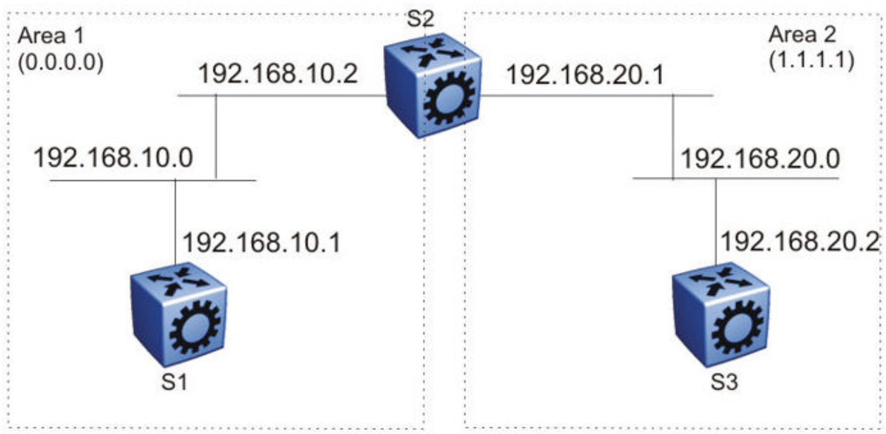

Figure 24: Example 3: OSPF on two subnets in two areas

The routers in scenario 3 use the following configuration:

- S1 has an OSPF router ID of 1.1.1.1. The OSPF port uses an IP address of 192.168.10.1, which is in OSPF area 1.
- S2 has an OSPF router ID of 1.1.1.2. One port uses an IP address of 192.168.10.2, which is in OSPF area 1. The second OSPF port on S2 uses an IP address of 192.168.20.1, which is in OSPF area 2.
- S3 has an OSPF router ID of 1.1.1.3. The OSPF port uses an IP address of 192.168.20.2, which is in OSPF area 2.

The general method to configure OSPF for this three-switch network is:

1. On all three switches, enable OSPF globally.
2. Configure OSPF on one network.

   On S1, insert the IP address, subnet mask, and VLAN ID for the OSPF port. Enable OSPF on the port. On S2, insert the IP address, subnet mask, and VLAN ID for the OSPF port in area 1, and enable OSPF on the port. Both routable ports belong to the same network. Therefore, by default, both ports are in the same area.
3. Configure three OSPF areas for the network.
4. Configure OSPF on two additional ports in a second subnet.

   Configure additional ports and verify that IP forwarding is enabled for each switch to ensure that routing can occur. On S2, insert the IP address, subnet mask, and VLAN ID for the OSPF port in area 2, and enable OSPF on the port. On S3, insert the IP address, subnet mask, and VLAN ID for the OSPF port, and enable OSPF on the port.

The three switches exchange hello packets.

In an environment with a mix of non-Avaya and Avaya switches and routers, you may need to manually modify the OSPF parameter RtrDeadInterval to 40 seconds.

Border Gateway Protocol

Use the Border Gateway Protocol (BGP) to ensure that the switch can communicate with other BGP-speaking routers on the Internet backbone. BGP is an exterior gateway protocol that exchanges network reachability information with other BGP systems in the same or other autonomous systems (AS). This network reachability information includes information about the AS list that the reachability information traverses. By using this information, you can prune routing loops and enforce policy decisions at the AS level.

BGP performs routing between two sets of routers that operate in different autonomous systems. An AS can use two kinds of BGP: Interior BGP (IBGP), which refers to the protocol that BGP routers use within an autonomous system, and Exterior BGP (EBGP), which refers to the protocol that BGP routers use across two different autonomous systems. BGP information is redistributed to Interior Gateway Protocols (IGP) that run in the autonomous system.

BGP version 4 (BGPv4) supports classless inter-domain routing (CIDR). BGPv4 advertises the IP prefix and eliminates the concept of network class within BGP. BGP4 can aggregate routes and AS paths. BGP aggregation does not occur when routes have different Multi-Exit Discriminators (MED) or next-hops.

BGP Equal-Cost Multipath (ECMP) allows a BGP speaker to perform route balancing within an AS by using multiple equal-cost routes submitted to the routing table by OSPF or RIP. BGP performs load balancing on an individual packet basis.

To control route propagation and filtering, RFC1772 and RFC2270 recommends that multihomed, nontransit Autonomous Systems not run BGPv4. To address the load sharing and reliability requirements of a multihomed user, use BGP between them.

For more information about BGP concepts and configuration, see *Avaya Virtual Services Platform 4000 Configuration — BGP Services*, NN46251-507.

BGP implementation guidelines

To successfully implement BGP in a Virtual Services Platform 4000 network, follow these guidelines:

- BGP does not operate with an IP router in nonforwarding (host-only) mode. Ensure that the routers with which you want BGP to operate are in forwarding mode.
- If you use BGP for a multi-homed AS (one that contains more than a single exit point), Avaya recommends that you use OSPF for the IGP, and BGP for the sole exterior gateway protocol. Otherwise, use intra-AS IBGP routing.
- If OSPF is the IGP, use the default OSPF tag construction. The use of EGP or the modification of the OSPF tags makes network administration and proper configuration of BGP path attributes difficult.

- For routers that support both BGP and OSPF, you must configure the OSPF router ID and the BGP identifier to the same IP address. The BGP router ID automatically uses the OSPF router ID.
- In configurations where BGP speakers reside on routers that have multiple network connections over multiple IP interfaces (the typical case for IBGP speakers), consider using the address of the circuitless (virtual) IP interface as the local peer address. In this configuration, you ensure that BGP is reachable as long as an active circuit exists on the router.
- By default, BGP speakers do not advertise or inject routes into their IGP. You must configure route policies to enable route advertisement.
- Coordinate routing policies among all BGP speakers within an AS so that every BGP border router within an AS constructs the same path attributes for an external path.
- Configure accept and announce policies on all IBGP connections to accept and propagate all routes. Make consistent routing policy decisions on external BGP connections.
- Use the max-prefix parameter to limit the number of routes BGP imports from a peer. Use a configuration of 0 to accept an unlimited number of prefixes.
- You cannot enable or disable the MED selection process. BGP aggregation does not occur when routes have different MEDs or next-hops.

BGP and OSPF interaction

RFC1745 defines the interaction between BGP and OSPF when OSPF is the IGP within an autonomous system. For routers that run both protocols, the OSPF router ID and the BGP ID must be the same IP address. You must configure a BGP route policy to allow BGP advertisement of OSPF routes.

Interaction between BGPv4 and OSPF includes the ability to advertise supernets to support CIDR. BGPv4 supports interdomain supernet advertisements; OSPF can carry supernet advertisements within a routing domain.

BGP and other vendor interoperability

BGP interoperability is compatible between Virtual Services Platform 4000, Cisco 6500 Software Release IOS 11.3, and Juniper M20 Software Release 5.3R2.4.

For more information about BGP, see *Avaya Virtual Services Platform 4000 Configuration — BGP Services*, NN46251-507.

BGP and Internet peering

By using BGP, you can perform Internet peering directly between Virtual Services Platform 4000 and another edge router. In such a scenario, you can use each Virtual Services Platform 4000 for aggregation and peer it with a Layer 3 edge router, as shown in the following figure.

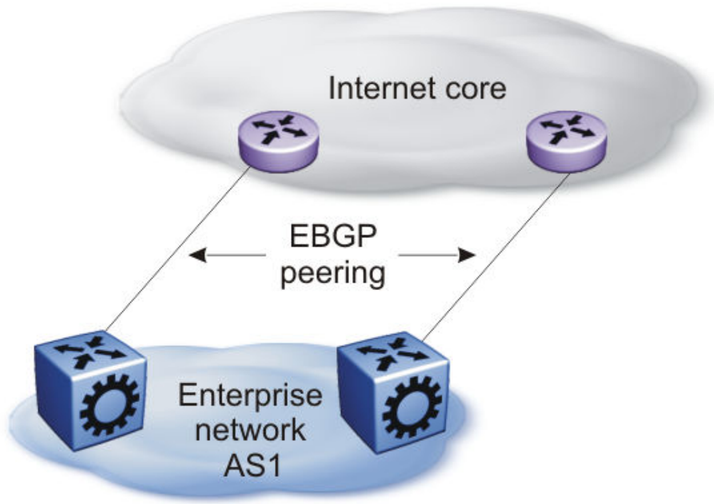

Figure 25: BGP and Internet peering

In cases where the Internet connection is single-homed, to reduce the size of the routing table, Avaya recommends that you advertise Internet routes as the default route to the IGP. The VSP 4000 supports a total of 16000 eBGP routes. The maximum FIB size for IPv4 routes is also 16000.

Routing domain interconnection with BGP

You can implement BGP so that autonomous routing domains, such as OSPF routing domains, connect. This connection allows the two different networks to begin communicating quickly over a common infrastructure, thus providing additional time to plan the IGP merger. Such a scenario is particularly effective when you need to merge two OSPF area 0.0.0.0s (see the following figure).

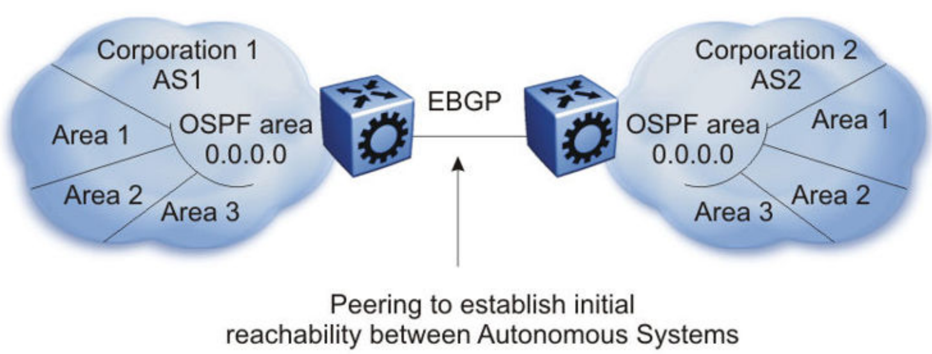

Figure 26: Routing domain interconnection with BGP

BGP and edge aggregation

You can perform edge aggregation with multiple point of presence or edge concentrations. Virtual Services Platform 4000 supports 12 pairs (peering services). You can use BGP to inject dynamic routes rather than using static routes or RIP (see the following figure).

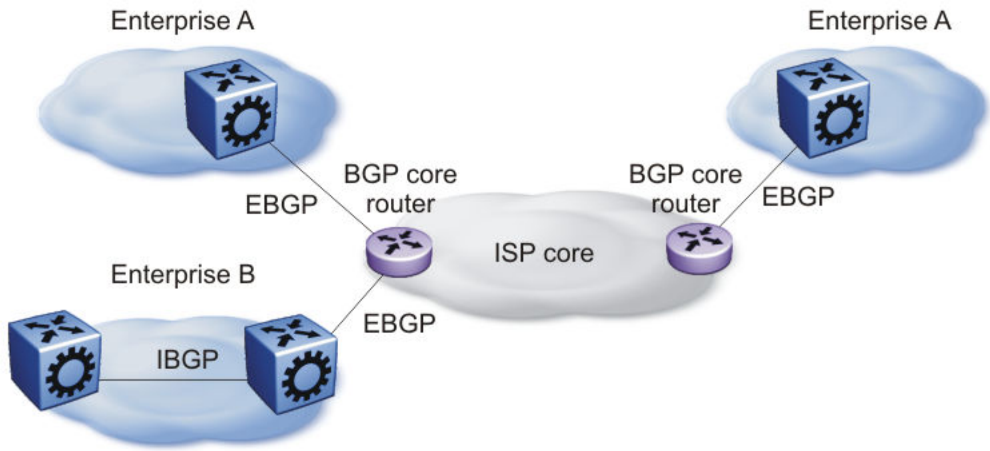

Figure 27: BGP and edge aggregation

BGP and ISP segmentation

You can use the platform as a peering point between different regions or ASs that belong to the same ISP. In such cases, you can define a region as an OSPF area, an AS, or a part of an AS.

You can divide the AS into multiple regions that each run different IGPs. Interconnect regions logically by using a full IBGP mesh. Each region then injects its IGP routes into IBGP and also injects a default route inside the region. For destinations that do not belong to the region, each region defaults to the BGP border router.

Use the community parameter to differentiate between regions. To provide Internet connectivity, this scenario requires you to make your Internet connections part of the central IBGP mesh (see the following figure).

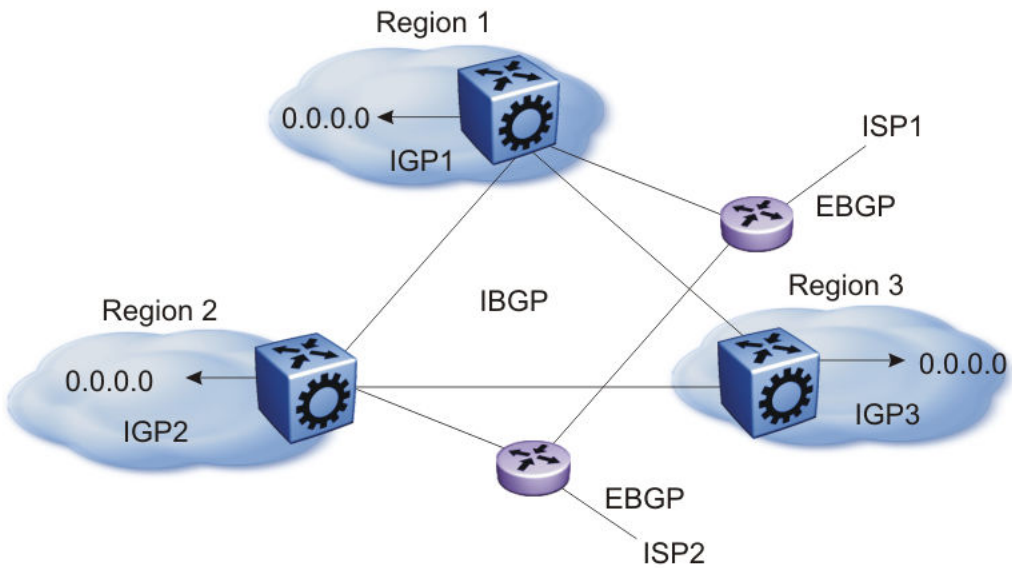

Figure 28: Multiple regions separated by IBGP

In the preceding figure, consider the following

- The AS is divided into three regions that each run different and independent IGPs.
- Regions logically interconnect by using a full-mesh IBGP, which also provides Internet connectivity.
- Internal nonBGP routers in each region default to the BGP border router, which contains all routes.
- If the destination belongs to another region, the traffic is directed to that region; otherwise, the traffic is sent to the Internet connections according to BGP policies.

To configure multiple policies between regions, represent each region as a separate AS. Implement EBGP between ASs, and implement IBGP within each AS. In such instances, each AS injects its IGP routes into BGP where they are propagated to all other regions and the Internet.

The following figure shows the use of EBGP to join several ASs.

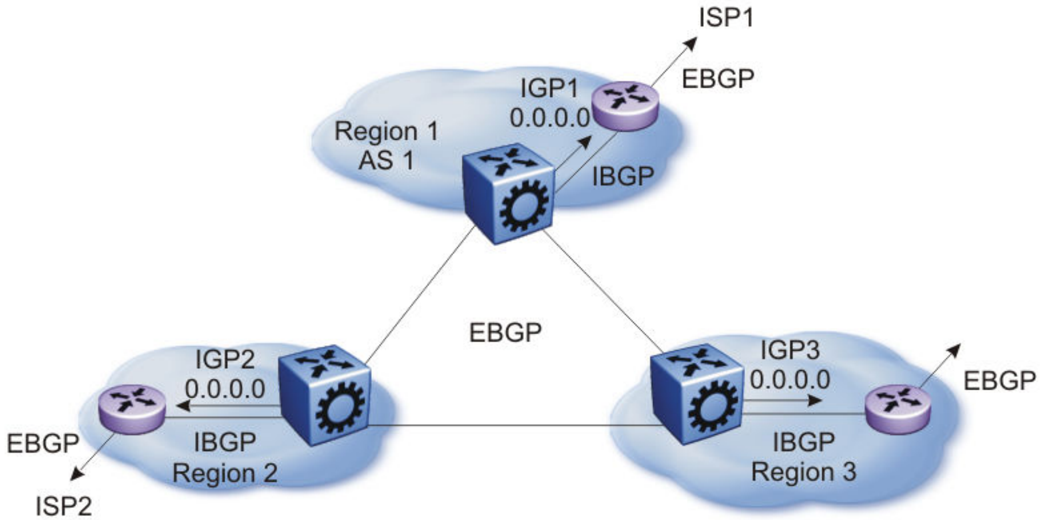

Figure 29: Multiple regions separated by EBGP

You can obtain AS numbers from the Inter-Network Information Center (NIC) or use private AS numbers. If you use private AS numbers, be sure to design your Internet connectivity carefully. For example, you can introduce a central, well-known AS to provide interconnections between all private ASs and the Internet. Before it propagates the BGP updates, this central AS strips the private AS numbers to prevent them from leaking to providers.

The following figure illustrates a design scenario in which you use multiple OSPF regions to peer with the Internet.

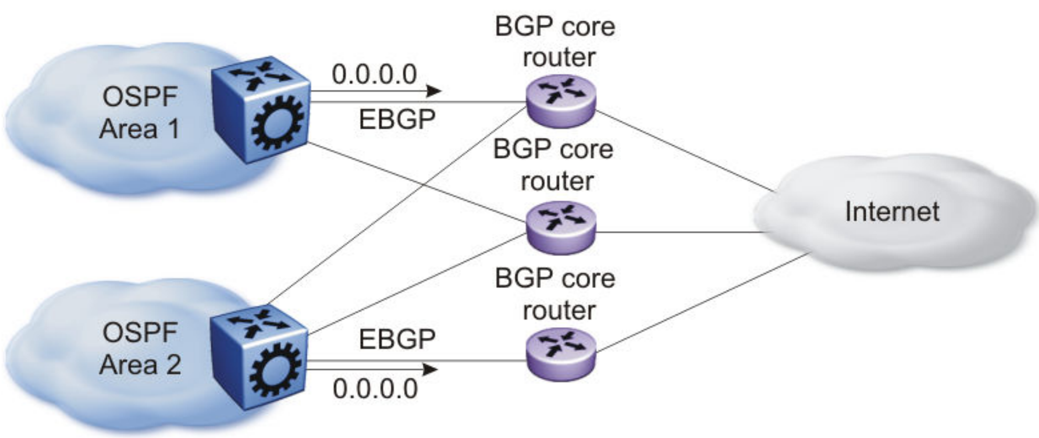

Figure 30: Multiple OSPF regions peering with the Internet

IP routed interface scaling

Virtual Services Platform 4000 supports up to 256 IP routed interfaces.

When you configure a large number of IP routed interfaces, use passive interfaces on most of the configured interfaces. You can make very few interfaces active.

Chapter 11: SPBM design guidelines

Shortest Path Bridging MAC (SPBM) is a next generation virtualization technology that revolutionizes the design, deployment, and operations of enterprise edge campus core networks and data centers. The benefits of the technology are clearly evident in its ability to provide massive scalability while at the same time reducing the complexity of the network. SPBM makes network virtualization a much easier paradigm to deploy within the enterprise environment than other technologies.

The following sections provide design guidelines that illustrate the operational simplicity of SPBM. It also lists best practices to configure SPBM in your network. For more information about SPBM fundamental concepts, command structure and basic configuration, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 4000*, NN46251–510.

802.1aq standard

Virtual Services Platform 4000 supports the IEEE 802.1aq standard of SPBM. SPBM makes network virtualization easy to deploy within the enterprise environment by reducing the complexity of the network while at the same time providing greater scalability. This technology provides all the features and benefits required by carrier-grade deployments to the enterprise market without the complexity of alternative technologies traditionally used in carrier deployments, for example, Multiprotocol Label Switching (MPLS). SPBM integrates into a single control plane all the functions that MPLS requires multiple layers and protocols to support.

IS-IS

SPBM eliminates the need for multiple overlay protocols in the core of the network by reducing the core to a single Ethernet-based, link-state protocol (IS-IS). IS-IS provides virtualization services, both Layer 2 and Layer 3, using a pure Ethernet technology base. SPBM also uses IS-IS to discover and advertise the network topology, which enables it to compute the shortest path to all nodes in the SPBM network.

Spanning Tree is a topology protocol that prevents loops but does not scale very well. Because SPBM uses IS-IS, which has its own mechanisms to prevent loops, SPBM does not have to use Spanning Tree to provide a loop free Layer 2 domain.

SPBM uses the IS-IS shortest path trees to populate forwarding tables for the individual backbone MAC (B-MAC) addresses of each participating node. Depending on the topology, SPBM supports as many Equal Cost Multipath trees as there are backbone VLAN IDs (B-VIDs) provisioned (with a maximum of 16 B-VIDs allowed by the standard and 2 allowed in this in release) per IS-IS instance. IS-IS interfaces operate in point-to-point mode only, which means that for any given port or MLT interface where IS-IS has been enabled, there can be only one IS-IS adjacency across that interface.

B-MAC

An SPBM backbone includes Backbone Edge Bridges (BEB) and Backbone Core Bridges (BCB). A BEB performs the same functionality as a BCB, but it also terminates one or more Virtual Service Networks (VSNs). A BCB does not terminate any VSNs and is unaware of the VSN traffic it transports. A BCB simply knows how to reach any other BEB in the SPBM backbone.

To forward customer traffic across the service provider backbone, the BEB for the VSN encapsulates the customer Ethernet packet received at the edge into a Backbone MAC header using the 802.1ah MAC-in-MAC encapsulation. This encapsulation hides the Customer MAC (C-MAC) address in a Backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header for delivery from end to end. Intermediate BCB nodes within the SPBM backbone perform packet forwarding using BMAC-DA alone. When the packet reaches the intended egress BEB, the B-MAC header is removed and the original customer packet is forwarded.

I-SID

SPBM introduces a service instance identifier called I-SID. SPBM uses I-SIDs to separate services from the infrastructure. After you create an SPBM infrastructure, you can add additional services (such as VLAN extensions or VRF extensions) by provisioning the endpoints only. The SPBM endpoints are BEBs, which mark the boundary between the core MAC-in-MAC SPBM domain and the edge customer 802.1Q domain. I-SIDs are provisioned on the BEBs to be associated with a particular service instance. In the SPBM core, the bridges are BCBs. BCBs forward encapsulated traffic based on the BMAC-DA.

The SPBM B-MAC header includes an I-SID. The length of the I–SID is 32 bits with a 24–bit ID. I-SIDs identify and transmit virtualized traffic in an encapsulated SPBM frame. These I-SIDs are used in a VSN for VLANs or VRFs across the MAC-in-MAC backbone:

- For a Layer 2 VSN, the I-SID is associated with a customer VLAN, which is then virtualized across the backbone. Layer 2 VSNs offer an any-any LAN service type. Layer 2 VSNs associate one VLAN per I-SID.
- For a Layer 2 VSN with multicast, the BEB associates a data I-SID with the multicast stream and a scope I-SID that defines the scope as a Layer 2 VSN. A multicast stream with a Layer 2 VSN scope can only transmit a multicast stream for the same Layer 2 VSN.
- For a Layer 3 VSN, the I-SID is associated with a customer VRF, which is also virtualized across the backbone. Layer 3 VSNs are always full-mesh topologies. Layer 3 VSNs associate one VRF per I-SID.
- For a Layer 3 VSN with multicast, the BEB associates a data I-SID with the multicast stream and a scope I-SID that defines the scope as a Layer 3 VSN. A multicast stream with a Layer 3 VSN scope can only transmit a multicast stream for the same Layer 3 VSN.
- For IP Shortcuts with multicast, the BEB associates a data I-SID with the multicast stream and defines the scope as Layer 3 GRT. A multicast stream with a scope of Layer 3 GRT can only transmit a multicast stream for the Layer 3 GRT.

Encapsulating customer MAC addresses in backbone MAC addresses greatly improves network scalability (no end-user C-MAC learning required in the core) and also significantly improves network robustness (loops have no effect on the backbone infrastructure).

The following figure shows the components of a basic SPBM architecture.

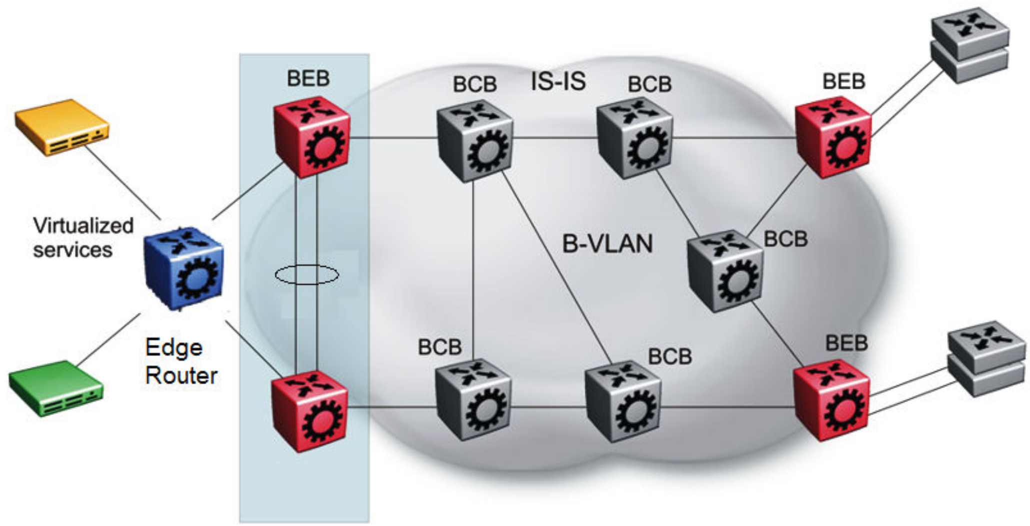

Figure 31: SPBM basic architecture

VLANs without member ports

The ERS 8800 manages VLANs without member ports differently than the VSP 9000 and VSP 4000.

- The ERS 8800 always designates the VLAN as operationally up if there is an attached I-SID.
- The VSP 9000 and VSP 4000 designate the VLAN as operationally up only if there is a matching I-SID in the SPBM network. For more information, see the following sections.

ERS 8800 implementation

If a VLAN has an IP address and is attached to an I-SID, the ERS 8800 designates that VLAN as operationally up whether it has a member port or not. When the VLAN is operationally up, the IP address of the VLAN will be in the routing table.

The ERS 8800 design behaves this way because the VLAN might be acting as an NNI in cases of Layer 2 Inter-VSN routing. If the VLAN was acting as a UNI interface, it would require a member port.

VSP 9000 and VSP 4000 implementation

If a VLAN is attached to an I-SID there must be another instance of that same I-SID in the SPBM network.

- If there is another instance of that I-SID, the device designates that VLAN as operationally up regardless of whether it has a member port or not.

  When the VLAN is operationally up, the IP address of the VLAN will be in the routing table.

- If there is NO matching instance of the I-SID in the SPBM network, then that VLAN does not have any reachable members so it is not acting as an NNI interface.

  The VLAN is also not acting as a UNI interface because it does not have a member port.

  Therefore, the device does not designate the VLAN as operationally up because the VLAN is not acting as a UNI or an NNI interface.

If the device is just acting as a BCB with two VLANs configured and two I-SIDs, there must be a UNI side with the corresponding I-SID existing in the network.

If the device is acting as both BEB and BCB, then there must be a member port in that VLAN in order to push out the UNI traffic.

Provisioning

This section summarizes how to provision SPBM. For information on specific configuration commands, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 4000*, NN46251–510.

Infrastructure provisioning

Provisioning an SPBM core is as simple as enabling SPBM and IS-IS globally, and on all the IS-IS core Ethernet links on all the BCB and BEB nodes. The IS-IS protocol operates at Layer 2 so it does not need IP addresses configured on the links to form IS-IS adjacencies with neighboring switches (like OSPF does). You do not need to configure IP addresses on any of the core links. The encapsulation of customer MAC addresses in backbone MAC addresses greatly improves network scalability.

No flooding or learning of end-user MACs occurs in the backbone. This SPBM provisioning significantly improves network robustness, as customer-introduced network loops have no effect on the backbone infrastructure.

Service provisioning

Provision I-SIDs on a BEB to associate that BEB with a particular service instance. After you map the customer VLAN or VRF into an I-SID, any BEB that has the same I-SID configured can participate in the same Layer 2 or Layer 3 VSN. This same simplicity extends to provisioning the services to run above the SPBM backbone:

- To create a Layer 2 VSN, associate an I-SID number with an edge VLAN.
- To create a Layer 3 VSN, associate an I-SID number with a VRF and configure the desired IS-IS IP route redistribution within the newly created Layer 3 VSN.

Note:

No service provisioning is needed on the core BCB SPBM switches. This provides a robust carrier grade architecture where configuration on the core switches never needs to be touched when adding new services.

IP multicast over SPBM

Provisioning IP multicast over SPBM is as simple as enabling SPBM multicast on the BEBs. You do not need to enable IP multicast over SPBM on the BCBs.

For Layer 2 VSN using IP multicast over SPBM, configure IGMP snooping on the VLAN that represents the Layer 2 VSN.

For Layer 3 VSNs using IP multicast over SPBM, configure the Layer 3 VSN as a multicast VSN, and then enable IP multicast over SPBM on each VLAN within the VRF to which IP multicast senders and receivers attach.

For IP Shortcuts using IP multicast over SPBM, enable IP multicast over SPBM on each of the VLANs within the GRT that need to support IP multicast traffic.

Implementation options

The SPBM architecture is architecturally simple and easy to provision, but it is not just for simple networks. SPBM supports multiple implementation options within the same network to meet the demands of the most complex network configurations. The following figure shows how SPBM supports multiple campus networks as well as multiple data centers.

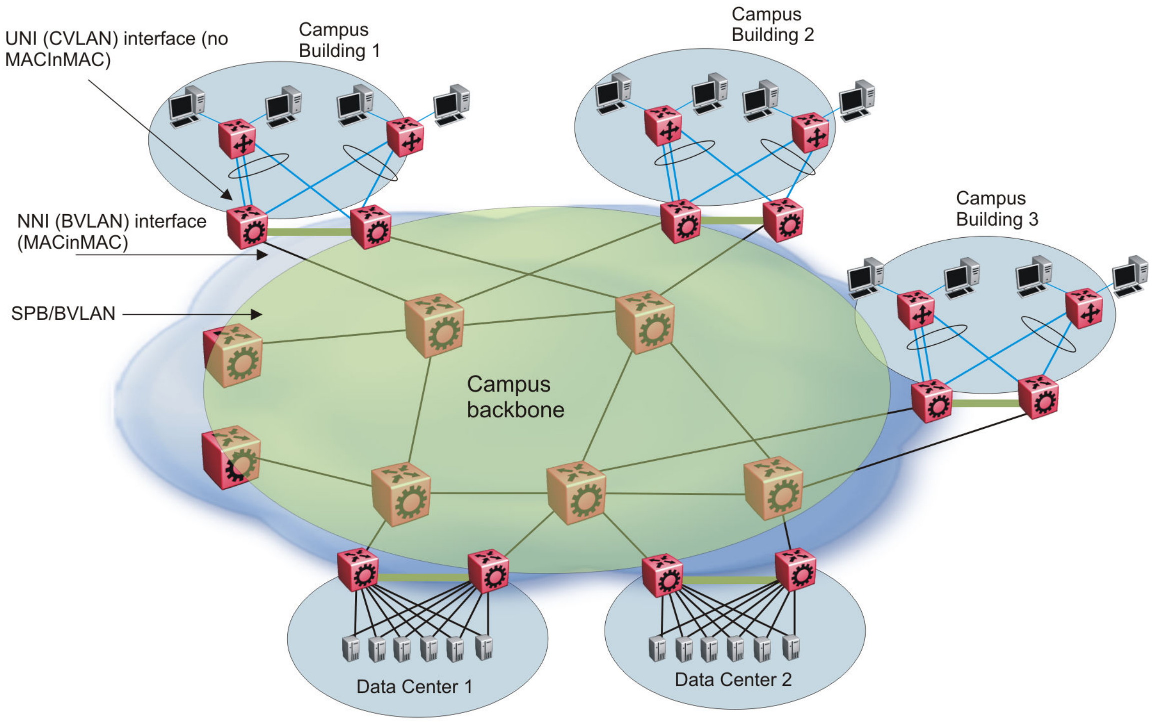

Figure 32: SPBM support for campus and data center architecture

Within the SPBM architecture, you can implement multiple options. The following figure shows all the options that SPBM supports.

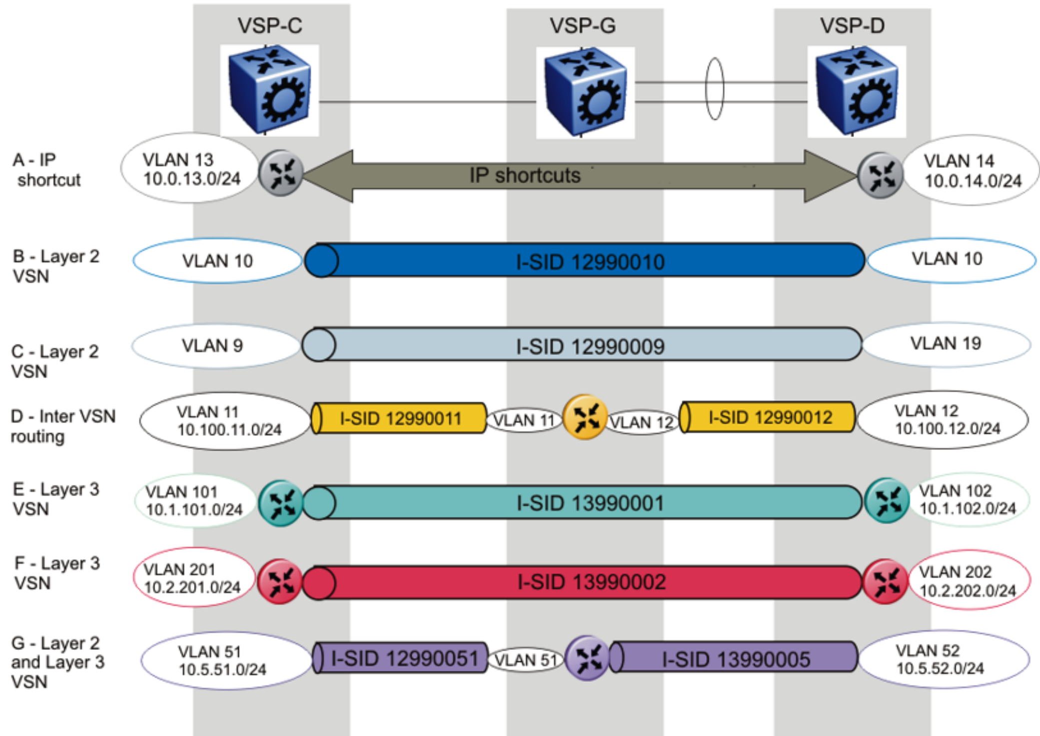

Figure 33: SPBM implementation options

A—IP shortcut

IP shortcuts forward standard IP packets over IS-IS. This option enables you to forward IP over the SPBM core, which is a simpler method than traditional IP routing or MPLS. SPBM nodes propagate Layer 3 reachability as leaf information in the IS-IS link state packets (LSP) using Extended IP reachability TLV 135, which contains routing information such as neighbors and locally configured subnets. SPBM nodes that receive the reachability information use this information to populate the routes to the announcing nodes. All TLVs announced in the IS-IS LSPs are grafted onto the shortest path tree (SPT) as leaf nodes.

An IP route lookup is only required once where the source BEB uses the routing table to identify the BEB closest to the destination subnet. All other nodes perform standard Ethernet switching based on the existing SPT. This scenario allows for end to end IP-over-Ethernet forwarding without the need for Address Resolution Protocol (ARP), flooding, or reverse learning. Because BCB SPBM nodes only forward on the MAC addresses that comprise the B-MAC header, and because unknown TLVs in IS-IS are relayed to the next hop but ignored locally, SPBM BCB nodes need not be aware of IP subnets to forward IP traffic. Only BEBs generate and receive Extended IP reachability TLV to build the routing table; BCBs just relay the TLV to the next hop based on the SPT. In fact, the Extended IP reachabilty TLV is ignored on BCBs.

With IP shortcuts there are only two IP routing hops (ingress BEB and egress BEB) as the SPBM backbone acts as a virtualized switching backplane.

IP shortcuts do not require I-SID configuration. However, you must enable IP on IS-IS, and configure the IS-IS source address to match a circuitless or loopback IP address.

In Figure 33: SPBM implementation options on page 79, node VSP-G acts as a BCB for the service, and has no IP configuration.

B—Layer 2 VSN

A Layer 2 Virtual Services Network (VSN) bridges customer VLANs (C-VLANs) over the SPBM core infrastructure. A Layer 2 VSN associates a C-VLAN with an I-SID, which is then virtualized across the backbone. All VLANs in the network that share the same I-SID can participate in the same VSN. If you want IS-IS to distribute traffic across two equal cost paths, then you need two backbone VLANs (B-VLAN) with a primary B-VLAN and a secondary B-VLAN. Otherwise, you need only a single B-VLAN.

One of the key advantages of the SPBM Layer 2 VSN is that network virtualization provisioning is achieved by configuring the edge of the network (BEBs) only. The intrusive core provisioning that other Layer 2 virtualization technologies require is not needed when new connectivity services are added to the SPBM network. For example, when new virtual server instances are created and need their own VLAN instances, they are provisioned at the network edge only and do not need to be configured throughout the rest of the network infrastructure.

Based on its I-SID scalability, this solution can scale much higher than any 802.1Q tagging based solution. Also, due to the fact that there is no need for Spanning Tree in the core, this solution does not need any core link provisioning for normal operation. Redundant connectivity between the C-VLAN domain and the SPBM infrastructure can be achieved by operating two SPBM switches in switch clustering (SMLT) mode. This allows the dual homing of any traditional link aggregation capable device into an SPBM network

Note:

In the current release, VSP 4000 does not support Inter-Switch Trunking (IST), and hence cannot be provisioned as S-MLT peer nodes.

In Figure 33: SPBM implementation options on page 79, nodes VSP-C and VSP-D act as BEBs for the VSN. Only these nodes have a MAC table or forwarding database for C-VLAN 10.

C—Layer 2 VSN with VLAN translation

Layer 2 VSNs with VLAN translation are basically the same as the Layer 2 VSNs, except that the BEBs on either end of the SPBM network belong to different VLANs. This option enables you to connect one VLAN to another VLAN. In Figure 33: SPBM implementation options on page 79, VLAN 9 connects to VLAN 19. The mechanism that connects them is that they use the same I-SID (12990009).

D— Inter-VSN routing

Inter-VSN routing allows routing between Layer 2 VLANs with different I-SIDs. You can use Inter-VSN routing to redistribute routes between Layer 2 VLANs. This option allows effective networking of multiple VSNs. Where Layer 2 VSN with VLAN translation enabled you to interconnect VLANs, this option takes that concept one step further and allows you to interconnect VSNs. This option also provides the ability to route IP traffic on Layer 2-VSNs that enter on NNI interfaces, which is especially useful for Layer 2 edge solutions.

As seen in Figure 33: SPBM implementation options on page 79, routing between VLANs 11 and 12 occurs on the SPBM core switch VSP-G shown in the middle of the figure. With Inter-VSN routing enabled, VSP-G transmits traffic between VLAN 11 (I-SID 12990011) and VLAN

12 (I-SID 12990012) on the VRF instance configured. Note that for these VSNs, node VSP-G acts as a BEB.

E—Layer 3 VSN

Layer 3 VSNs are very similar to Layer 2 VSNs. The difference between the two is that Layer 2 VSNs associate I-SIDs with VLANs. Layer 3 VSNs associate I-SIDs with VRFs. With the Layer 3 VSN option, all VRFs in the network that share the same I-SID can participate in the same VSN by advertising their reachable IP routes into IS-IS and installing IP routes learned from IS-IS. Suitable IP redistribution policies need to be defined to determine what IP routes a BEB will advertise to IS-IS.

As seen in Figure 33: SPBM implementation options on page 79, the green VRF on VSP-C is configured to advertise its local or direct IP routes into IS-IS within I-SID 13990001. The VRF on node VSP-D, which is also a member of the same I-SID, installs these IP routes in its VRF IP routing table with a next-hop B-MAC address of VSP-C. Therefore, when the VRF on node VSP-D needs to IP route traffic to the IP subnet off VSP-C, it performs a lookup in its IP routing table and applies a MAC-in- MAC encapsulation with B-MAC DA of VSP-C. The SPBM core ensures delivery to the egress BEB VSP-C where the encapsulation is removed and the packet is IP routed onwards.

Note:

Like the IP shortcut service, there are only two IP routing hops (ingress BEB and egress BEB) as the SPBM backbone acts as a virtualized switching backplane.

F—Layer 3 VSN

Figure 33: SPBM implementation options on page 79 shows two VRFs (green and red) to illustrate that the BEBs can associate ISIDs with multiple VRFs. The Layer 3 VSN option provides IP connectivity over SPBM for all of your VRFs.

G—Layer 2 VSN and Layer 3 VSN

Figure 33: SPBM implementation options on page 79 shows both a Layer 2 VSN and a Layer 3 VSN to show that you can configure both options on the same BEBs. This topology is simply made up of a number of BEBs that terminate VSNs of both types. What this example shows is the flexibility to extend one or more edge VLANs (using one or more Layer 2 VSNs) to use a default gateway that is deeper into the SPBM core. From here, traffic can then be IP routed onwards as either non-virtualized with IP shortcuts or, as shown in this example, with a virtualized Layer 3 VSN. Note that in this example the central node VSP-G is now also acting as BEB for both service types as it now maintains both a MAC table for the Layer 2 VSN it terminates, and an ARP cache and IP routing table for the Layer 3 VSN it also terminates.

Multiple tenants using different SPBM services

The following figure shows multiple tenants using different services within an SPBM metro network. In this network, you can use some or all of the SPBM implementation options to meet the needs of the community while maintaining the security of information within VLAN members.

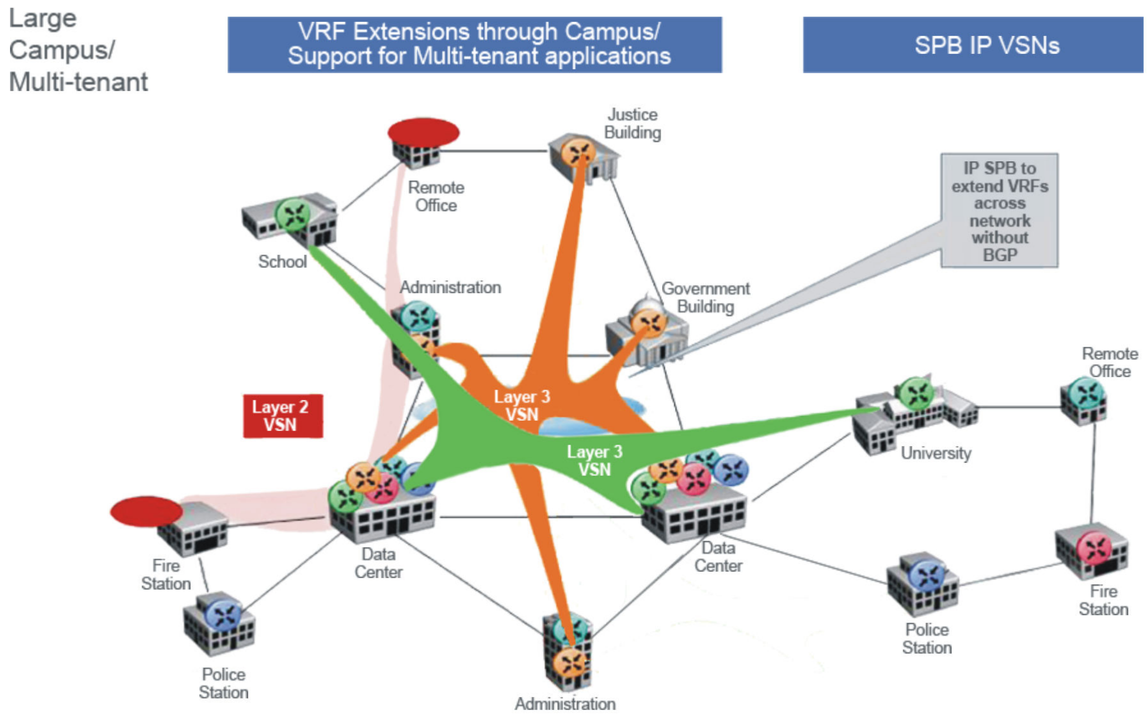

Figure 34: Multi-tenant SPBM metro network

To illustrate the versatility and robustness of SPBM even further, the following figure shows a logical view of multiple tenants in a ring topology. In this architecture, each tenant has their own domain where some users have VLAN requirements and are using Layer 2 VSNs and others have VRF requirements and are using Layer 3 VSNs. In all three domains, they can share data center resources across the SPBM network.

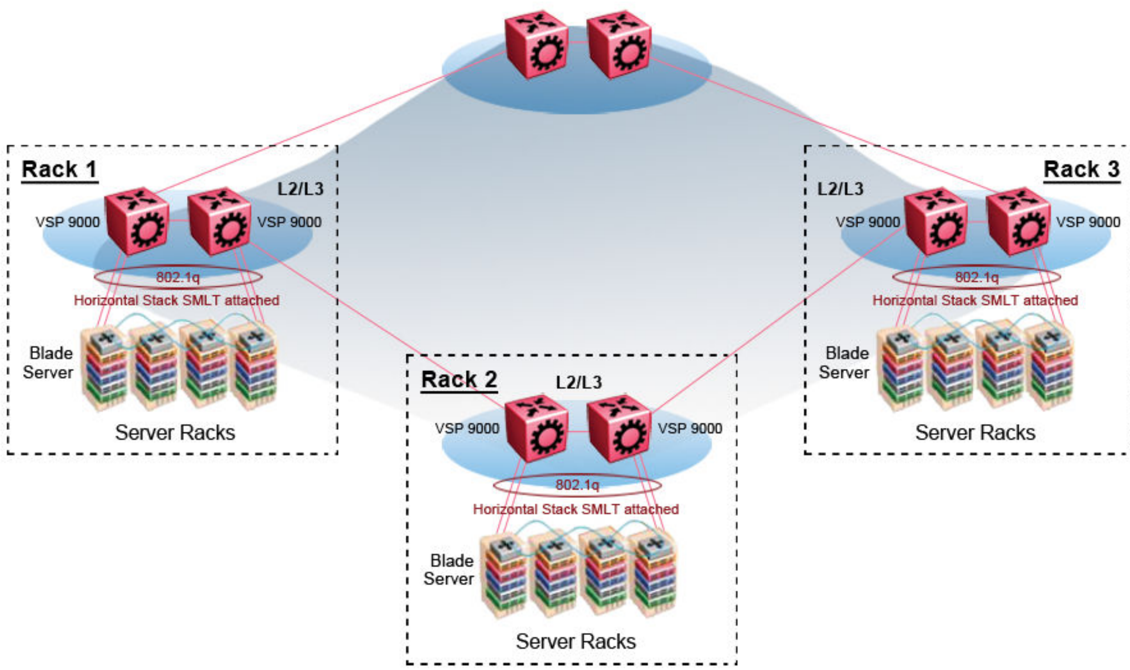

Figure 35: SPBM ring topology with shared data centers

Reference architectures

SPBM has a straightforward architecture that simply forwards encapsulated C-MACs across the backbone. Because the B-MAC header stays the same across the network, there is no need to swap a label or do a route lookup at each node. This architecture allows the frame to follow the most efficient forwarding path from end to end.

The following reference architectures illustrate SPBM with multiple VSP and ERS systems in a network. For information on solution-specific architectures like for Video Surveillance or Data Center implementation using the VSP 4000, see Solution specific reference architectures on page 96.

The following figure shows the MAC-in-MAC SPBM domain with BEBs on the boundary and BCBs in the core.

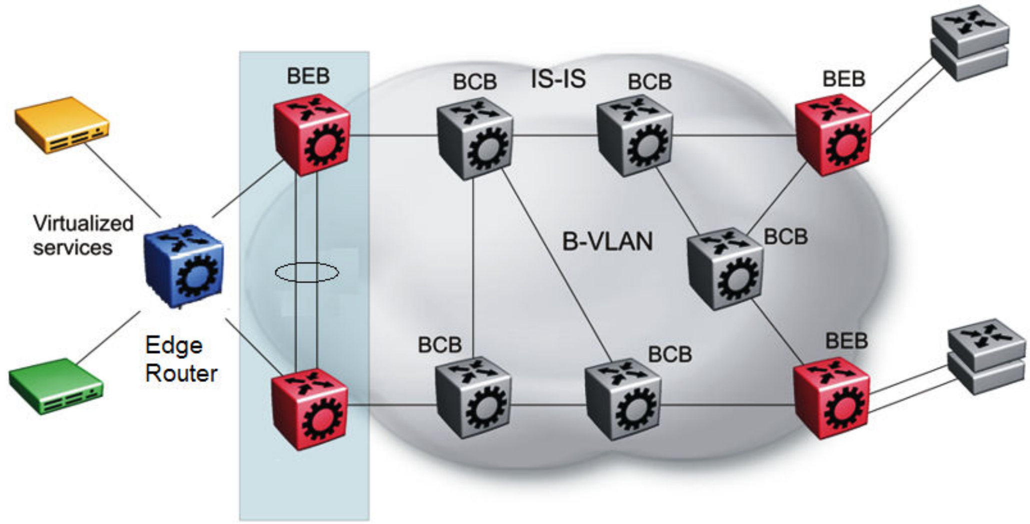

Figure 36: SPBM basic architecture

Provisioning an SPBM core is as simple as enabling SPBM and IS-IS globally on all the nodes and on the core facing links. To migrate an existing edge configuration into an SPBM network is just as simple.

The boundary between the MAC-in-MAC SPBM domain and the 802.1Q domain is handled by the BEBs. At the BEBs, VLANs or VRFs are mapped into I-SIDs based on the local service provisioning. Services (whether Layer 2 or Layer 3 VSNs) only need to be configured at the edge of the SPBM backbone (on the BEBs). There is no provisioning needed on the core SPBM nodes.

The following figure illustrates an existing edge that connects to an SPBM core.

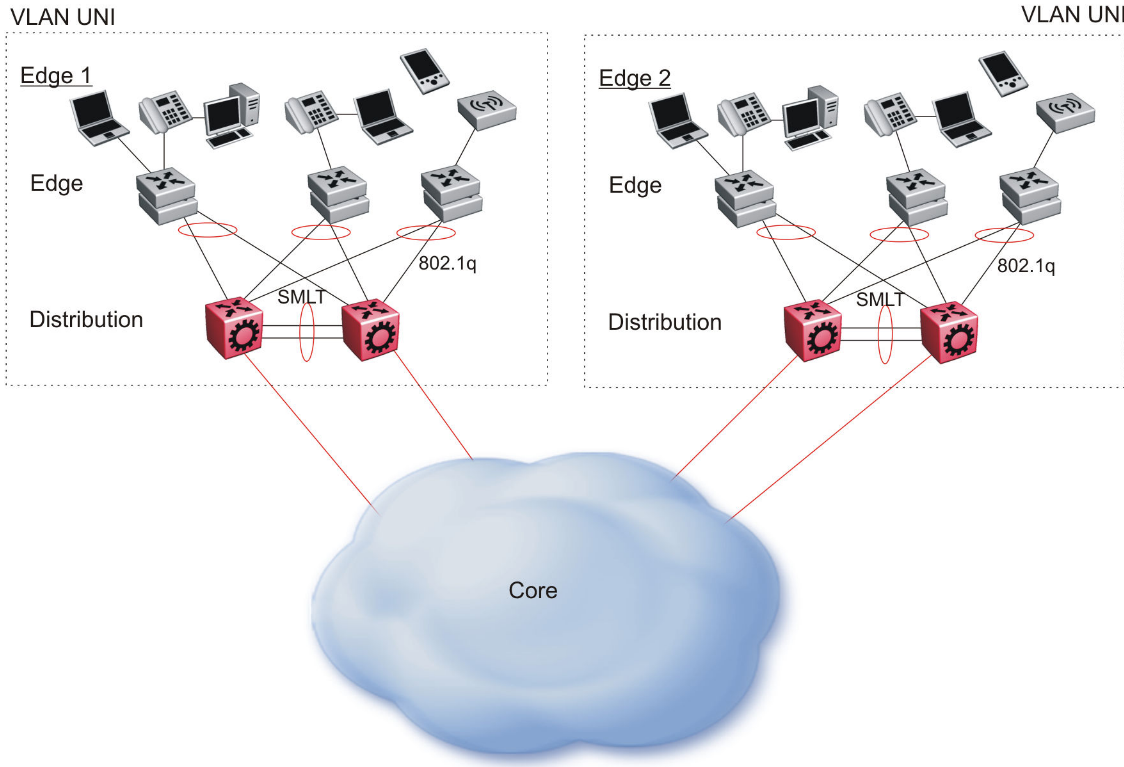

Figure 37: Access to the SPBM Core

For Layer 2 virtualized bridging (Layer 2 VSN), identify all the VLANs that you want to migrate into SPBM and assign them to an I-SID on the BEB.

For Layer 3 virtualized routing (Layer 3 VSN), map IPv4-enabled VLANs to VRFs, create an IP VPN instance on the VRF, assign an I-SID to the VRF, and then configure the desired IP redistribution of IP routes into IS-IS.

All BEBs that have the same I-SID configured can participate in the same VSN. That completes the configuration part of the migration and all the traffic flows will be back to normal.

SPBM on Virtual Services Platform 4000 supports the following traffic:

- Layer-2 bridged traffic (Layer 2 VSN)
- IPv4 unicast routed traffic on the Global Router (IP shortcuts)
- IPv4 unicast routed traffic using a VRF (Layer 3 VSN)
- IPv4 unicast routed traffic using different VSNs, which have different I-SIDs (Inter-VSN)
- Layer-2 IP multicast traffic in a bridged network (Layer 2 VSN with IP multicast over SPBM)

- IPv4 multicast routed traffic on the Global Router (IP shortcuts with IP multicast over SPBM)
- IPv4 multicast routed traffic using a VRF (Layer 3 VSN with IP multicast over SPBM)

Campus architecture

For migration purposes, you can add SPBM to an existing network that has SMLT configured. In fact, if there are other protocols already running in the network such as Open Shortest Path First (OSPF), you can leave them in place too. SPBM uses IS-IS, and operates independently from other protocols. However, Avaya recommends that you eventually eliminate SMLT in the core and other unnecessary protocols. This reduces the complexity of the network and makes it much simpler to maintain and troubleshoot.

Whether you configure SMLT in the core or not, the main point to remember is that SPBM separates services from the infrastructure. For example, in a large campus, a user may need access to other sites or data centers. SPBM enables you to grant that access by associating the user to a specific I-SID. This mechanism enables the user to work without getting access to confidential information of another department.

The following figure depicts a topology where the BEBs in the edge and data center distribution nodes are configured in SMLT clusters. Prior to implementing SPBM, the core nodes would also have been configured as SMLT clusters. When migrating SPBM onto this network design, it is important to note that you can deploy SPBM over the existing SMLT topology without network interruption. After the SPBM infrastructure is in place, you can create VSN services over SPBM or migrate them from the previous end-to-end SMLT-based design.

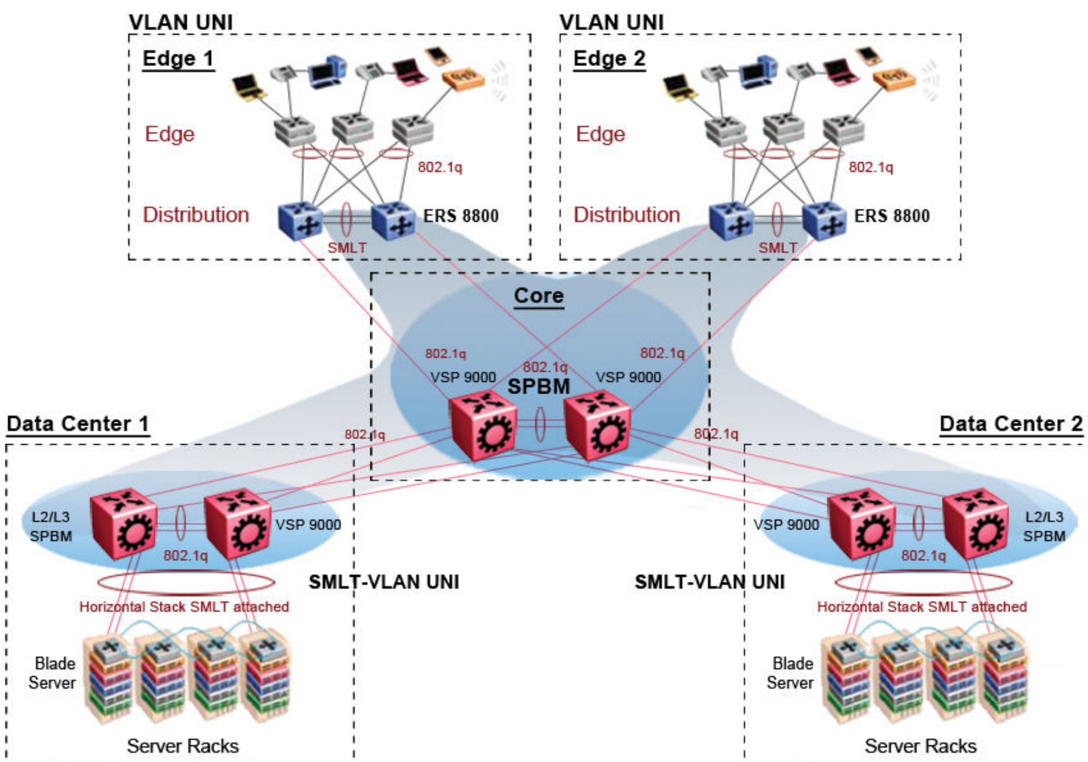

Figure 38: SPBM campus without SMLT

After migrating all services to SPBM, the customer VLANs (C-VLANs) will exist only on the BEB SMLT clusters at the edge of the SPBM network. The C-VLANs will be assigned to an I-

SID instance and then associated with either a VLAN in an Layer 2 VSN or terminated into a VRF in an Layer 3 VSN. You can also terminate the C-VLAN into the default router, which uses IP shortcuts to IP route over the SPBM core.

In an SPBM network design, the only nodes where it makes sense to have an SMLT cluster configuration is on the BEB nodes where VSN services terminate. These are the SPBM nodes where C-VLANs exist and these C-VLANs need to be redundantly extended to non-SPBM devices such as Layer 2 edge stackable switches. On the BCB core nodes where no VSNs are terminated and no Layer 2 edge stackables are connected, there is no longer any use for the SMLT clustering functionality. Therefore, in the depicted SPBM design, the SMLT/IST configuration can be removed from the core nodes because they now act as pure BCBs that have no knowledge of the VSN they transport and the only control plane protocol they need to run is IS-IS.

Because SMLT BEB nodes exist in this design (the edge BEBs) and it is desirable to use equal cost paths to load balance VSN traffic across the SPBM core, all SPBM nodes in the network are configured with the same two B-VIDs.

Where Figure 38: SPBM campus without SMLT on page 86 shows the physical topology, the following two figures illustrate a logical rendition of the same topology. In both of the following figures, you can see that the core is almost identical. Because the SPBM core just serves as a transport mechanism that transmits traffic to the destination BEB, all the provisioning is done at the edge.

In the data center, VLANs are attached to Inter-VSNs that transmit the traffic across the SPBM core between the data center on the left and the data center on the right. A common application of this service is VMotion moving VMs from one data center to another.

The following figure uses IP shortcuts that route VLANs. There is no I-SID configuration and no Layer 3 virtualization between the edge distribution and the core. This is normal IP forwarding to the BEB.

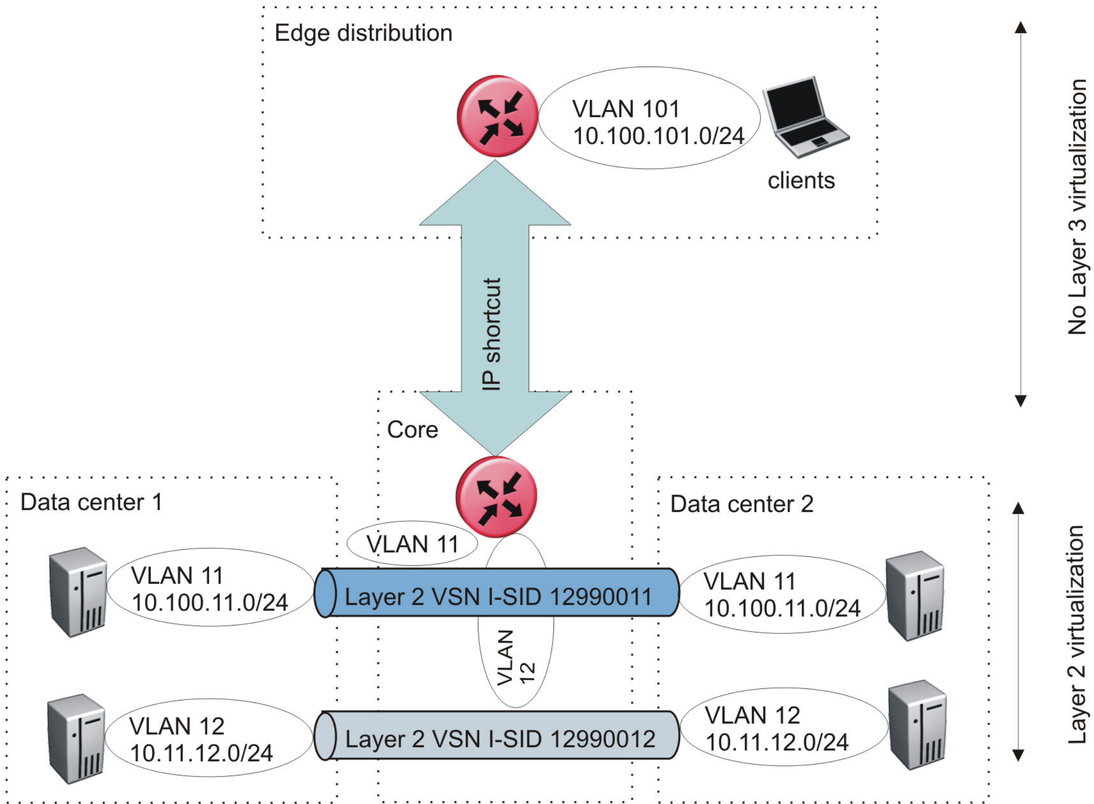

Figure 39: IP shortcut scenario to move traffic between data centers

The following figure uses Layer 3 VSNs to route VRFs between the edge distribution and the core. The VRFs are attached to I-SIDs and use Layer 3 virtualization.

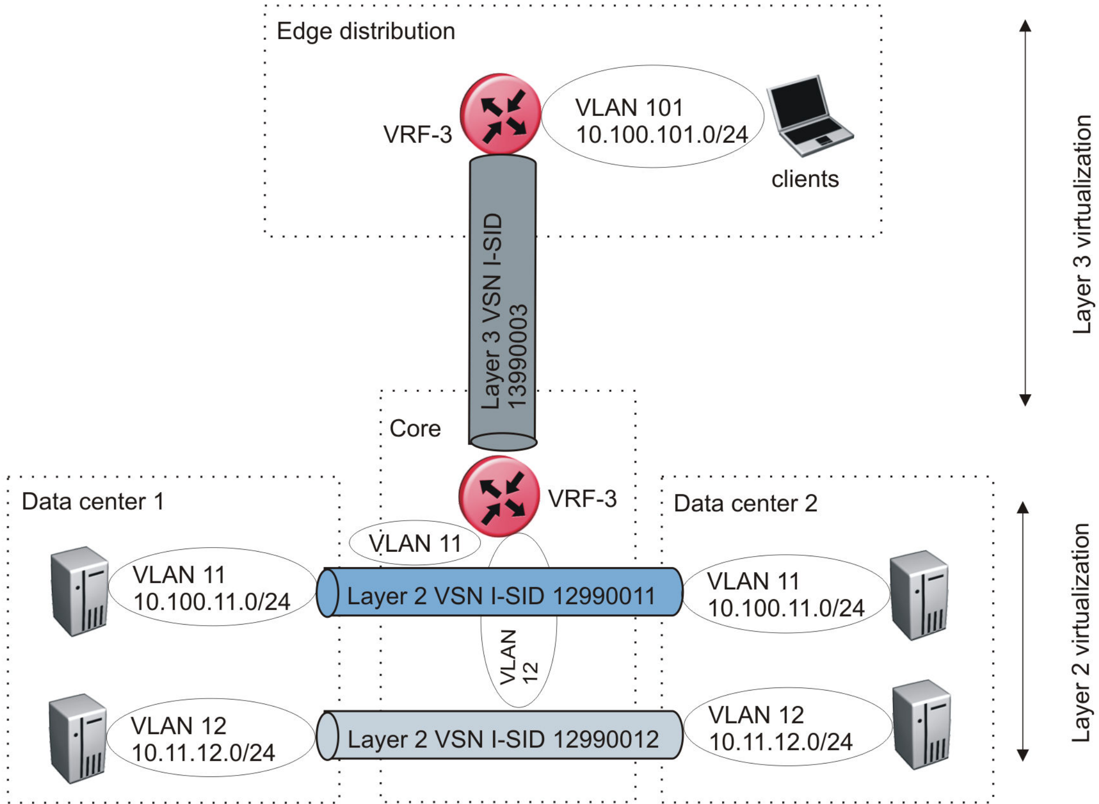

Figure 40: VRF scenario to move traffic between data centers

Multicast architecture

Networks today either have inefficient bridged IP multicast networks (IGMP) or IP multicast networks that require multiple protocols that are complex to configure and operate. IP multicast over SPBM builds on the simplicity introduced with Avaya VENA using SPBM for the control plane with support for bridged and routed IP multicast traffic, without the inefficiencies or complexities that exist in other topologies today.

This functionality extends the SPBM IS-IS control plane to additionally exchange IP multicast stream advertisement and membership information, which means that you can use SPBM for Layer 2 (unicast, broadcast, multicast) virtualization as well as Layer 3 (unicast, multicast) routing and forwarding virtualization.

IP multicast over SPBM supports three operational models:

1. Layer 2 Virtual Services Network with IGMP support on the access networks for optimized forwarding of IP multicast traffic in a bridged network (L2 VSN with multicast)
2. IP multicast routing support for Global Routing Table using SPB in the core (IP Shortcuts with multicast)
3. Layer 3 Virtual Services Network with VRF based IP multicast routing support over SPB in the core and IGMP on the access (L3 VSN with multicast)

All multicast streams are constrained within the level in which they originate, which is called the scope level. In other words, if a sender transmits a multicast stream to a BEB on a C-VLAN with IP multicast over SPBM enabled, only receivers that are part of the same Layer 2 VSN can receive that stream. Similarly, if a sender transmits a multicast stream to a BEB on a VLAN that is part of the Layer 3 VSN with IP multicast over SPBM enabled, only receivers that are part of the same Layer 3 instance can receive that stream.

IP multicast over SPBM uses BEBs to act as senders and receivers of data. After a BEB receives IP multicast data from a sender, a BEB allocates a data I-SID in the range of 16,000,000 to 16,512,000 for the stream. The stream is identified by the S,G,V tuple, which is the source IP address, group IP Address, and the stream is identified by the local VLAN on which the stream is received. The BEB also sends a TLV update to its neighbors to inform them of the presence of an IP multicast stream, along with identifying the sender. The BEB propagates the information through the SPBM cloud through IS-IS TLV updates in LSPs that result in a multicast tree being created for that stream.

IGMP handles group membership registration to enable members to receive data. IGMP snooping listens to conversations between hosts and routers, and maintains a table of links that require IP multicast streams.

The BEBs also act as IGMP queriers and send out periodic IGMP queries. The IGMP querier enables the creation of the link table. After a BEB receives an IGMP join message from a receiver, a BEB queries the IS-IS database to check if a sender exists for the requested stream within the scope of the receiver. If the requested stream does not exist, the IGMP information is kept, but no further action is taken. If the requested stream exists, the BEB sends an IS-IS TLV update to its neighbors to inform them of the presence of a receiver and this information is propagated through the SPBM cloud.

IS-IS acts dynamically using the TLV information it receives from BEBs that connect to the sender and the receivers to create a multicast tree between them.

The following figure shows how multicast senders and receivers connect to the SPBM cloud using BEBs.

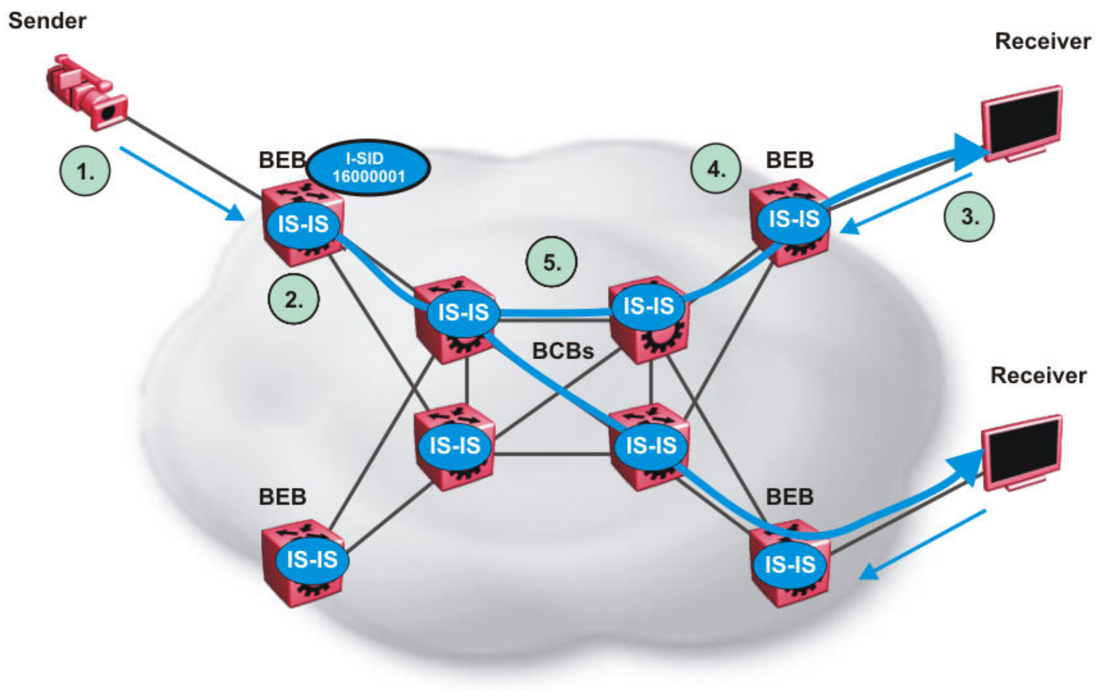

Figure 41: IP multicast over SPBM streams

The following list describes how multicast senders and receivers connect to the SPBM cloud using BEBs in the preceding diagram:

1. The sender sends multicast traffic with group IP address 233.252.0.1.
2. After the BEB receives the IP multicast stream from the sender, the BEB allocates data I-SID 16000001 for the S,G multicast stream. The BEB sends an LSP with the IPMC TLV (for Layer 2 VSN and Layer 3 VSN) or IPVPN TLV (for IP Shortcuts for multicast) with the transmit bit set. The BEB also sends an IS-IS service identifier and unicast address sub-TLV (where the unicast address has the multicast bit set and the I-SID is the data I-SID).
3. The receiver sends a join request to group 233.252.0.1.
4. The BEB (acting as the IGMP querier) queries the IS-IS database to find all senders for group 233.252.0.1. If the group exists, the BEB sends IS-IS service identifier and unicast address sub-TLV (where the unicast address has the multicast bit set and the nickname is the stream transmitter BEB and the I-SID is the data I-SID). If the requested stream does not exist, the BEB keeps the IGMP information, but no further action is taken.
5. IS-IS acts dynamically using the TLV information it receives from BEBs that connect to the sender and receivers to create a multicast tree between them and the data starts flowing from the sender.

For conceptual and configuration information on IP multicast over SPBM, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 4000*, NN46251–510.

Large data center architecture

SPBM supports data centers with IP shortcuts, Layer 2 VSNs, or Layer 3 VSNs. If you use vMotion, you must use Layer 2 between data centers (Layer 2 VSN). With Layer 2 VSNs, you can add IP addresses to the VLAN on both data centers and run Virtual Router Redundancy Protocol (VRRP) between them to allow the ESX server to route to the rest of the network.

The following figure shows an SPBM topology of a large data center. This figure represents a full-mesh Virtual Enterprise Network Architecture (VENA) data center fabric using SPBM for storage over Ethernet. This topology is optimized for storage transport because traffic never travels more than two hops.

Note:

Avaya recommends that you use a two-tier, full-mesh topology for large data centers.

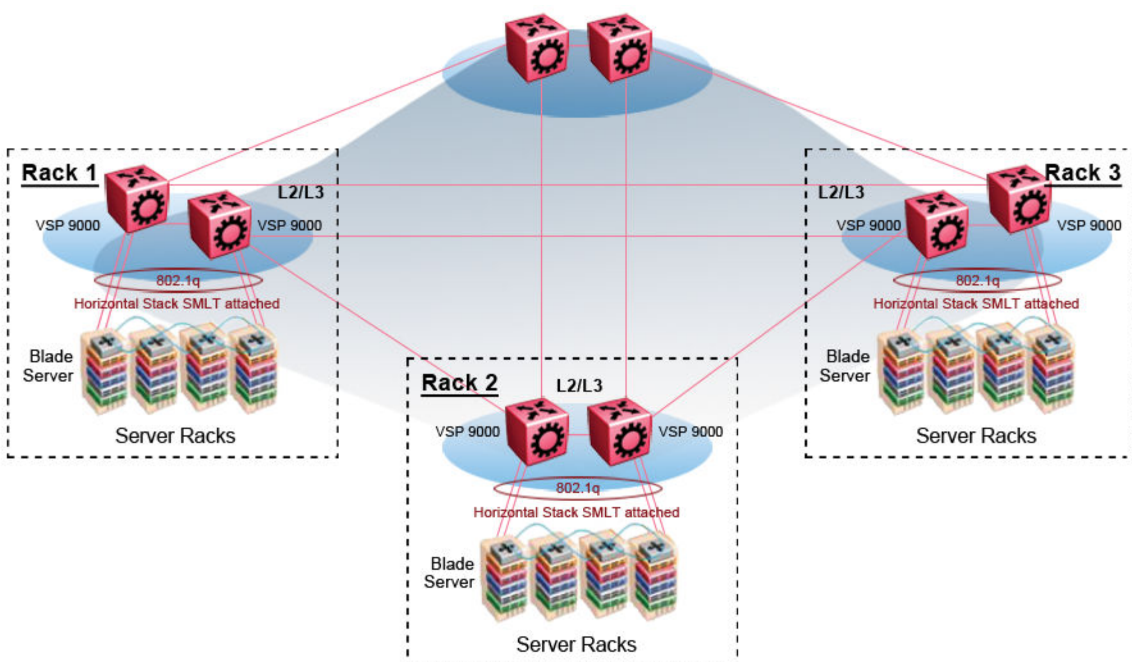

Figure 42: SPBM data center—full mesh

Traditional data center routing of VMs:

In a traditional data center configuration, the traffic flows into the network to a VM and out of the network in almost a direct path.

The following figure shows an example of a traditional data center with VRRP configured. Because end stations are often configured with a static default gateway IP address, a loss of the default gateway router causes a loss of connectivity to the remote networks. VRRP eliminates the single point of failure that can occur when the single static default gateway router for an end station is lost.

Data Center Routing — Traditional (pre VM move):

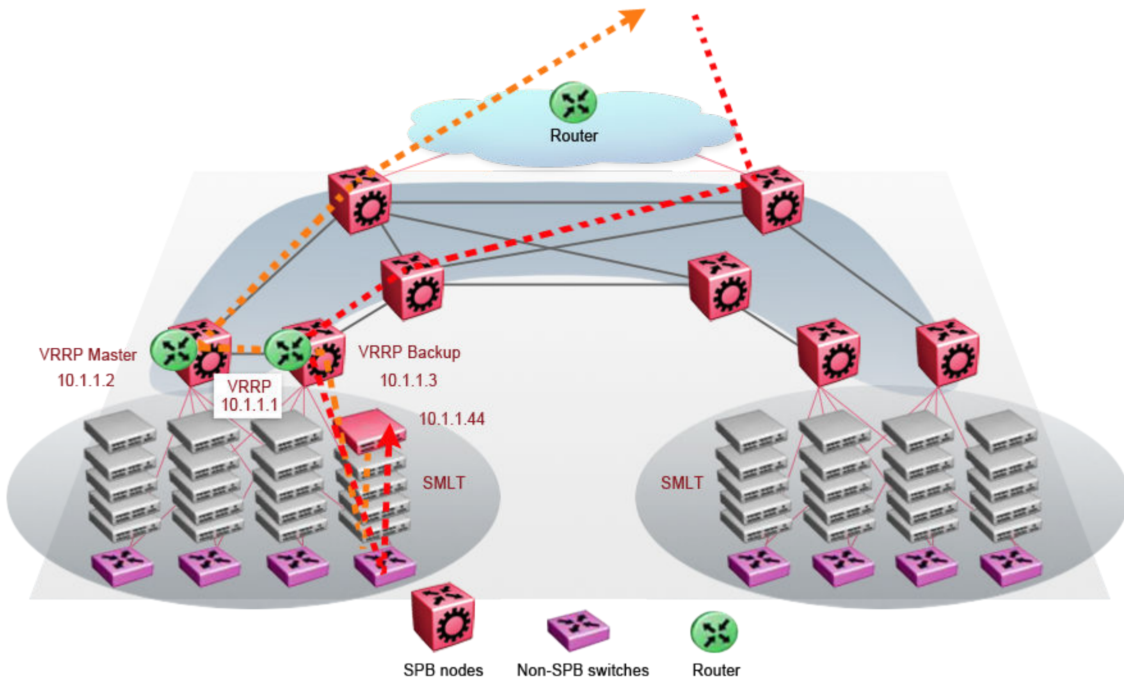

Figure 43: Traditional routing before moving VMs

A VM is a virtual server. When you move a VM, the virtual server is moved as is. This action means that the IP addresses of that server remain the same after the server is moved from one data center to the other. This in turn dictates that the same IP subnet (and hence VLAN) exist in both data centers.

In the following figure, the VM moved from the data center on the left to the data center on the right. To ensure a seamless transition that is transparent to the user, the VM retains its network connections through the default gateway. This method works, but it adds more hops to all traffic. As you can see in the figure, one VM move results in a complicated traffic path. Multiply this with many moves and soon the network look like a tangled mess that is very inefficient, difficult to maintain, and almost impossible to troubleshoot.

Data Center Routing — Traditional (post VM move):

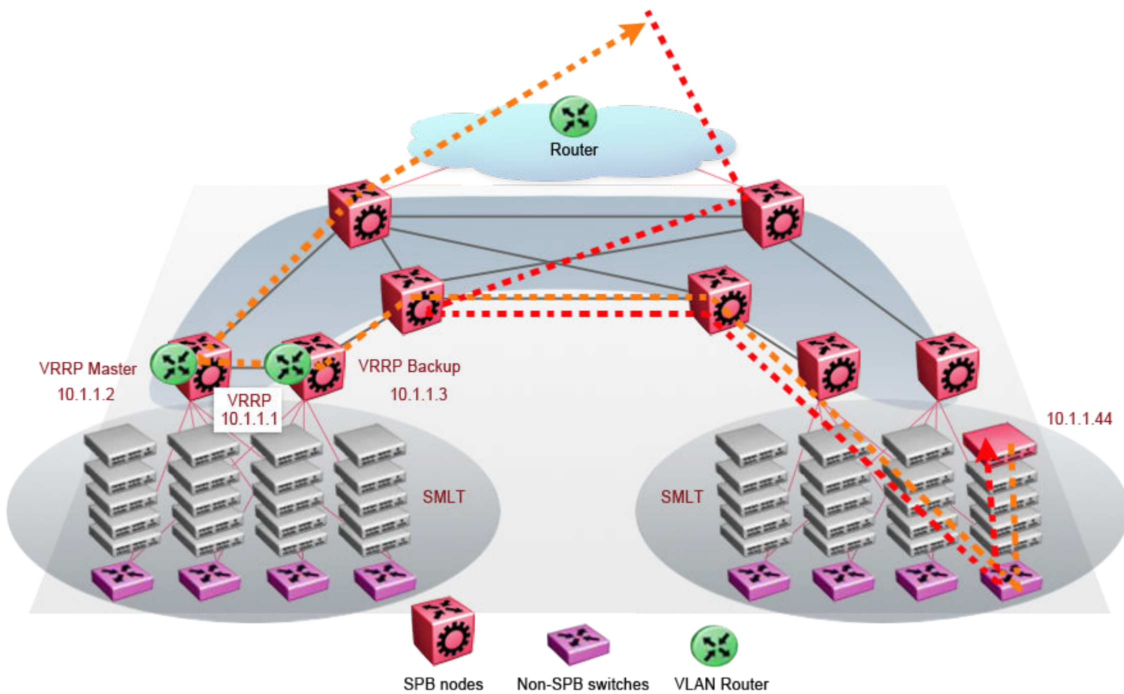

Figure 44: Traditional routing after moving VMs

Optimized data center routing of VMs:

Two features make a data center optimized:

- VLAN routers in the Layer 2 domain (green icons)
- VRRP BackupMaster

The VLAN routers use lookup tables to determine the best path to route incoming traffic (red dots) to the destination VM.

VRRP BackupMaster solves the problem of traffic congestion on the IST. Because there can be only one VRRP Master, all other interfaces are in backup mode. In this case, all traffic is forwarded over the IST link towards the primary VRRP switch. All traffic that arrives at the VRRP backup interface is forwarded, so there is not enough bandwidth on the IST link to carry all the aggregated riser traffic. VRRP BackupMaster overcomes this issue by ensuring that the IST trunk is not used in such a case for primary data forwarding. The VRRP BackupMaster acts as an IP router for packets destined for the logical VRRP IP address. All traffic is directly routed to the destined subnetwork and not through Layer 2 switches to the VRRP Master. This avoids potential limitation in the available IST bandwidth.

The following figure shows a solution that optimizes your network for bidirectional traffic flows. However, this solution turns two SPBM BCB nodes into BEBs where MAC and ARP learning will be enabled on the Inter-VSN routing interfaces. If you do not care about top-down traffic flows, you can omit the Inter-VSN routing interfaces on the SPBM BCB nodes. This makes the IP routed paths top-down less optimal, but the BCBs remain pure BCBs, thus simplifying core switch configurations.

Distributed vRouting — Optimized (pre VM move):

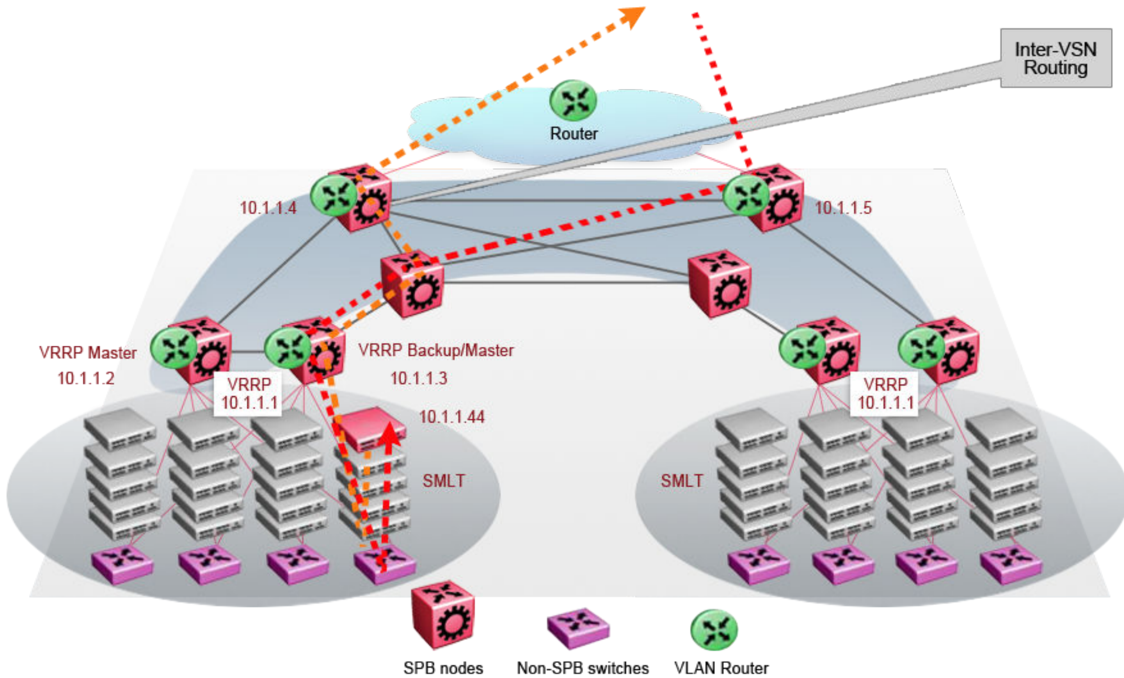

Figure 45: Optimized routing before moving VMs

In the traditional data center, chaos resulted after many VMs were moved. In an optimized data center as shown in the following figure, the incoming traffic enters the Layer 2 domain where an edge switch uses Inter-VSN routing to attach an I-SID to a VLAN. The I-SID bridges traffic directly to the destination. With VRRP BackupMaster, the traffic no longer goes through the default gateway; it takes the most direct route in and out of the network.

Distributed vRouting (post VM move):

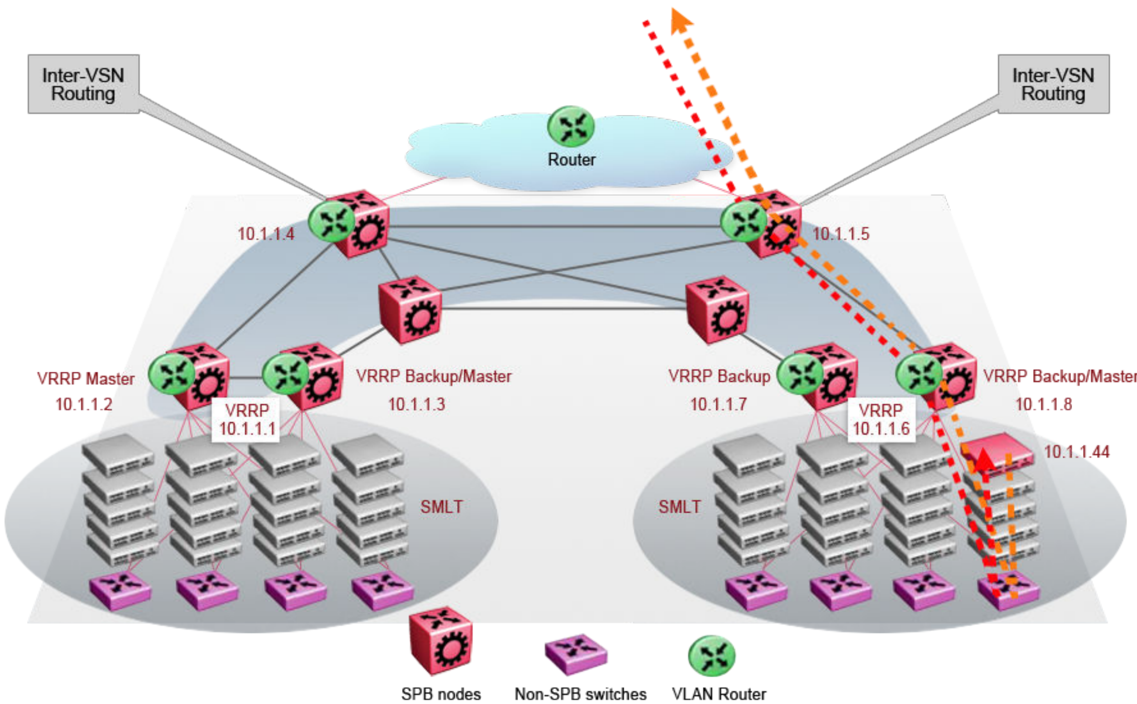

Figure 46: Optimized routing after moving VMs

Solution-specific reference architectures

The following sections describe solution-specific reference architectures, like for example for Video Surveillance or Data Center implementation, using the VSP 4000.

Multi tenant — fabric connect

This fabric connect based solution leverages the fabric capabilities of the VSP platforms; a VSP 7000 core and a VSP 4000 edge. This solution provides the ability to run up to 24 VRFs for each wiring closet and is well suited for multi-tenant applications. The zero-touch core is enabled by the fabric connect end-point provisioning capabilities.

If this solution must support IPv6, then a central router-pair routes all IPv6 traffic. The IPv6 traffic is tunneled from each wiring closet to the IPv6 routers by extending L2 VSNs to the q-tagged router interfaces

Note:

IPv6 is not supported in the current release of the VSP 4000.

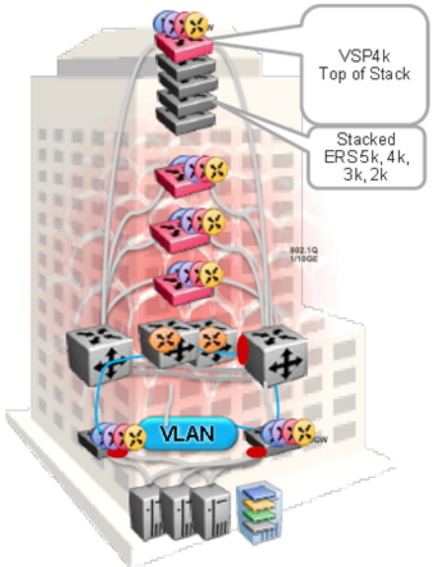

	Recommended Device Types	Connectivity and Resiliency Model	Protocol configuration
Edge	VSP4000 + ERS 5k/4k/3k/2k	1/10GE MLT/LAG SPB NNI	L3 VSNs for IPv4 routing L2 VSNs for IPv6 tunnelled to central ERS56xx(s) IP Multicast support with Release 3.1
Core – Distribution	VSP7024	SPB & MLT	SPB NNI & SPB UNI towards central router(s)
Core – IPv6 Routing	ERS56xx	VRRP	One or two routing switches with IPv6 interfaces & VRRP connected to L2 VSNs (no bridging)
Server Access	VSP7024 VSP4000	SPB & MLT	L2 VSNs L3 VSNs & L2 VSN routing
	Server	VMWare active-active Other servers active/passive	

Figure 47: Small core — Multi-tenant

The following list outlines the benefits of the fabric connect based solution:

- end-point provisioning
- fast failover
- simple to configure
- L2 and L3 virtualized

Hosted data center management solution — ETREE

In some hosted data center solutions, the hosting center operating company takes responsibility for managing customer servers. For this shared management, shown in the following figure, servers that control the operating system level of the production servers, such as the patch level, are deployed. Because customer production servers do not communicate with each other, a distributed private VLAN solution based on fabric connect is deployed to manage all production servers. This solution builds a distributed set of ETREEs for each management domain.

VSP 9000, ERS 8000, VSP 7000 as core and VSP 4000 as access, provide an elegant network-wide ETREE solution. Spokes, or managed servers, cannot communicate to each other over this network, but the shared management servers on the hub ports can access all spokes. Because of the L2 – ETREE nature of this setup, the managed servers do not require any route entries, and only require one IP interface in this management private VLAN. This solution supports tagged, untagged physical and virtual (VM) servers.

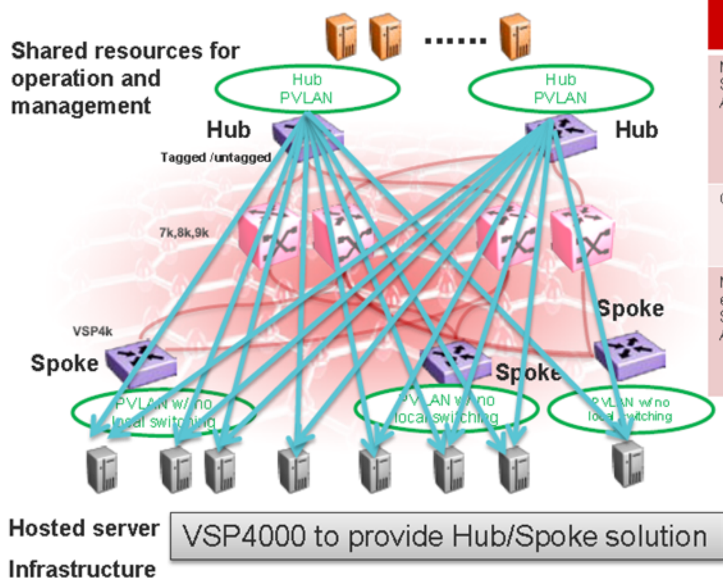

	Recom-mended Device Types	Connectivity and Resiliency Model	Protocol configura-tion
Mgmt Server Access	VSP7024/ VSP9000	SPB UNI/NNI MLT/LAG	SPB
Core	VSP7024/ VSP9000/ ERS8800	SPB NNI MLT	SPB
Managed Server Access	VSP4000	SPB-BEB MLT/LAG	SPB L2VSN ETREE

Figure 48: Data center hosting private VLAN

The following list outlines the benefits of the hosted data center management solution:

- Easy end-point provisioning
- optimal resiliency
- secure tenant separation

Video surveillance — bridged

In a video surveillance solution, optimal traffic forwarding is a key requirement to ensure proper operation of the camera and recorder solutions. However, signaling is also important to ensure quick channel switching. This is achieved by deploying a fabric connect based IP multicast infrastructure that is optimized for multicast transport, so that the cameras can be selected quickly, and so that there is no unnecessary traffic sent across the backbone.

Fabric connect enables this solution with support for ERS 8000, VSP 7000, VSP 9000 and the VSP 4000 products.

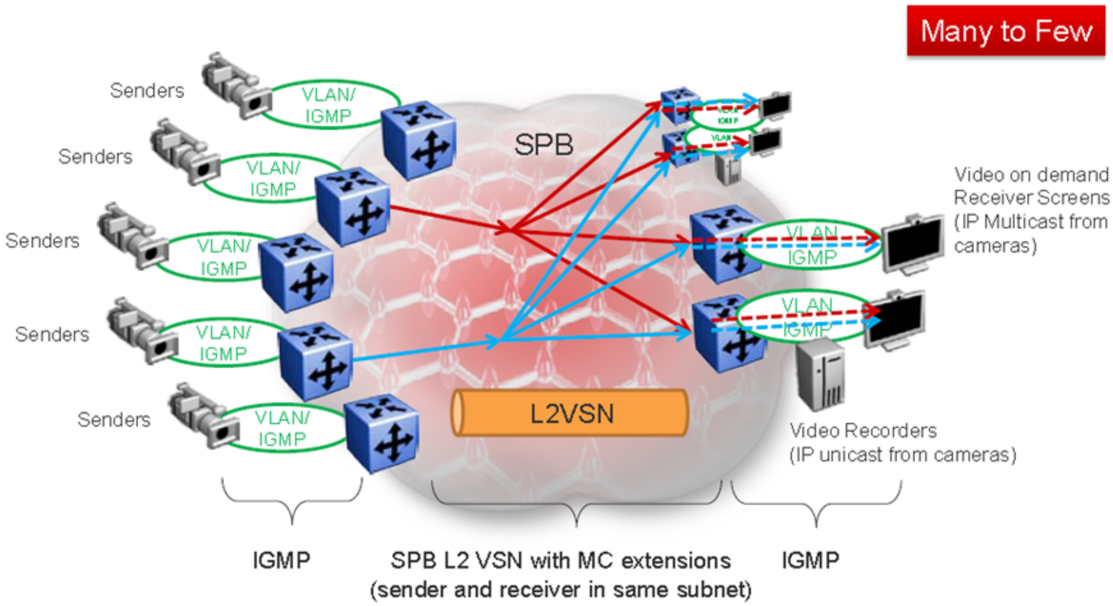

Figure 49: Deployment scenario — Bridged video surveillance and IP camera deployment for transportation, airports, and government

The following list outlines the benefits of the bridged video surveillance solution:

- Easy end-point provisioning
- sub second resiliency and mc forwarding
- secure tenant separation
- quick camera switching

Video surveillance — routed

In a video surveillance solution, optimal traffic forwarding is a key requirement to ensure proper operation of the camera and recorder solutions. However, signaling is also important to ensure quick channel switching. This is achieved by deploying an IP multicast infrastructure that is optimized for multicast transport, so that the cameras can be selected quickly, and so that there is no unnecessary traffic sent across the backbone. In the topology shown in the following figure, each camera is attached to its own IP subnet. In a larger topology, this can reduce network overhead. To increase network scalability, you can attach a set of cameras to an L2 switch that has IGMP, and then connect the cameras to the fabric edge (BEB) which has a routing instance.

In many customer scenarios, surveillance must be separated from the rest of the infrastructure. This can be achieved by deploying an L3 VSN for the surveillance traffic to keep the surveillance traffic isolated from any other tenant.

Fabric connect enables this solution with support for ERS 8000, VSP 7000, VSP 9000 and VSP 4000 products.

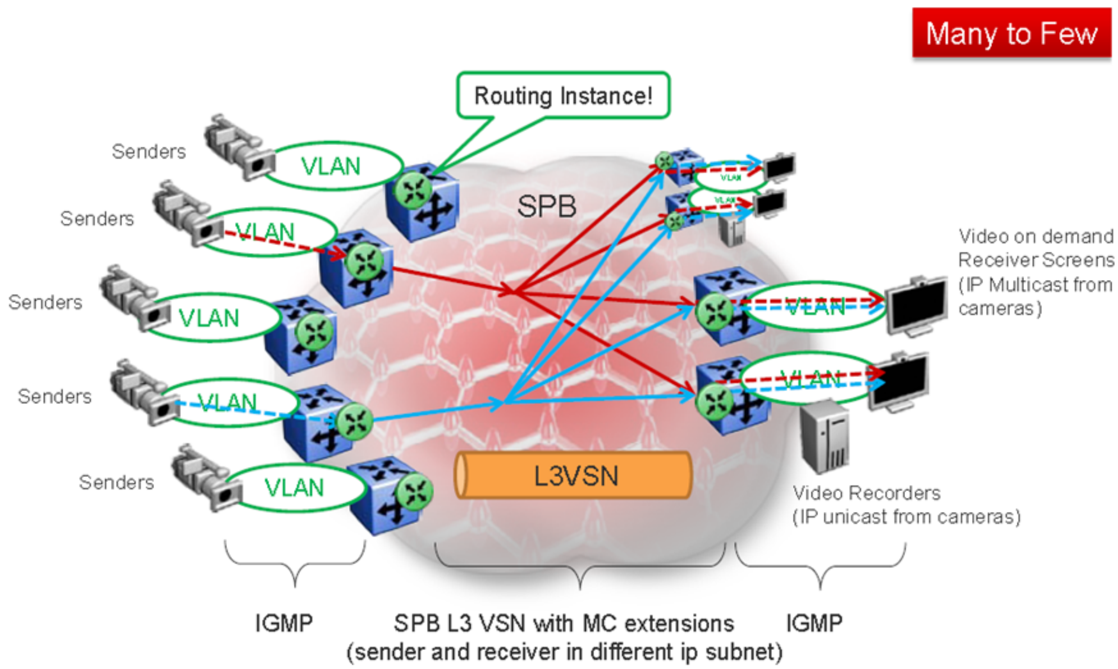

Figure 50: Deployment scenario — Routed video surveillance and IP camera deployment for transportation, airports, and government

The following list outlines the benefits of the routed video surveillance solution:

- Easy end-point provisioning
- optimal resiliency and mc forwarding
- secure tenant separation
- rapid channel/camera switching

Metro-Ethernet Provider solution

VSP 9000, ERS 8000, VSP 7000 and VSP 4000 provide an end-to-end Metro-Ethernet Provider solution. Leveraging fabric connect throughout the infrastructure enables a scalable and flexible wholesale provider infrastructure.

VSP 4000 switches are used as the access product, VSP 7000 switches build the distribution layer, and in large scale solutions VSP 9000 can be leveraged to build the core of the network.

This use case extends the Transparent UNI functionality to transparently forward any customer VLAN across the services.

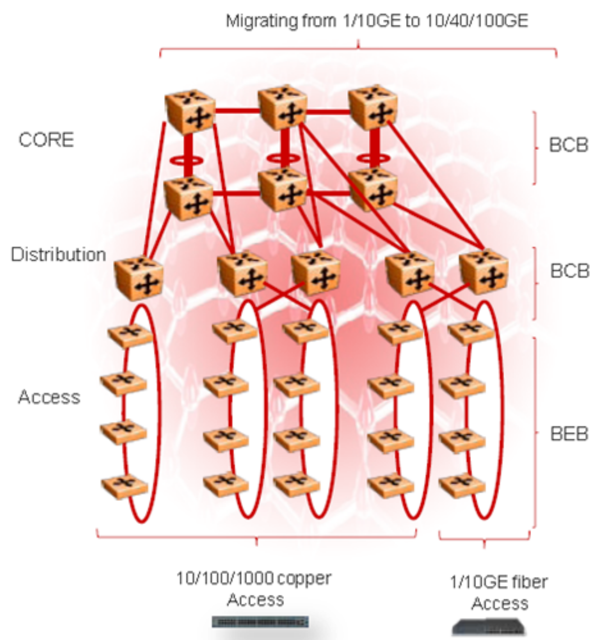

	Recommended Device Types	Connectivity and Resiliency Model	Protocol configuration
Core	VSP7024/VSP9000	SPB NNI MLT	SPB
Distribution	VSP7024/VSP9000/ERS8800	SPB NNI MLT	SPB
Access	VSP4000 VSP7000	MLT/LAG	SPB – UNI L2 VSNs L3 VSNs ETREE

Figure 51: Metro ring access solution

The following list outlines the benefits of the Metro-Ethernet Provider solution:

- Easy end-point provisioning
- optimal resiliency
- secure tenant separation

Best practices

This section provides best practices to configure an SPBM network.

IS-IS

The following list identifies best practices for IS-IS:

- Avaya recommends that you change the IS-IS system ID from the default B-MAC value to a recognizable address to easily identify a switch. This change helps to recognize source and destination addresses for troubleshooting purposes.
    - If you leave the system ID as the default value (safe practice as it ensures no duplication in the network), it can be difficult to recognize the source and destination B-MAC for troubleshooting purposes.

  - If you do manually change the system ID, take the necessary steps to ensure no duplication exists in the network.
- Create two B-VLANs to allow load distribution over both B-VLANs. This configuration is required if you use SMLT. Even if you do not use SMLT in the network, this is still good practice as adding a second B-VLAN to an existing configuration allows SPBM to load balance traffic across two equal cost multipaths if the physical topology grants it.
- In a ring topology with OSPF and IS-IS configured in the core, a core link break causes slow convergence that can lead to SPBM Layer 2 traffic loss. If the last member link of an OSPF VLAN fails, it takes down the IP interface and OSPF reconverges. While OSPF reconverges, SPBM does not have access to the CPU, which causes traffic loss.

SPBM

The following list identifies best practices for SPBM:

- Use a different, easily recognizable IS-IS nickname on each switch.
- If you enable IP shortcuts, you must configure an IS-IS IP source address on the switch.

Physical or MLT links between IS-IS switches

Virtual Services Platform supports only a single port or single MLT between a pair of IS-IS switches. For example, if two individual ports exist between a pair of IS-IS switches, you can configure IS-IS on both the ports but an IS-IS adjacency only forms on one of the links.

If you configure a single MLT between a pair of IS-IS switches, all ports in the MLT are used. You must configure the MLT before you enable IS-IS on the MLT.

CFM using manual mode configuration

Important:

Avaya recommends auto-configuration of CFM, which is simpler than the manual mode of configuration.

For more information on these commands, see *Avaya Command Line Reference, NN46251–104*.

The following list identifies best practices for manual configuration of Connectivity Fault Management (CFM):

- The CFM domain name must be the same on all switches in an IS-IS area.
- The maintenance association must be the same on all switches in an IS-IS area.
- To use CFM testing over both B-VLANs, create two maintenance associations — one for each B-VLAN.
- You can configure the same maintenance domain intermediate points (MIP) on all switches in an IS-IS area or uniquely defined per switch.

Important:

You must configure the MIP at the same level as the Maintenance Association Endpoints (MEP) on all switches in the SPBM network.

Example of a configuration using best practices

```
spbm-id : 1
BVID #1 & BVID #2 : 4040, 4041 (ignore warning message when configuring)
nick-name : b:b0:<node-id>
MEP-id : md.ma.<node-id>
BMAC : 00:bb:00:00:<node-id>:00
VirtBMAC : 00:bb:00:00:<node-id>:ff
MD : spbm (level 4)
MA : 4040 & 4041
mep : <node-id>
mip : (level 4)
isis manual area : 49.0001
```

SPBM restrictions and limitations

This section describes the restrictions and limitations associated with SPBM on Avaya Virtual Services Platform 4000.

RSTP and MSTP

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

- RSTP mode does not support SPBM.
- Since we support non-SPBM C-VLANs to also span the SPBM network, MSTP can be provisioned in the network to provide loop-free connectivity for these non-SPBM C-VLANs. Since all ports on the VSP 4000 system including IS-IS enabled NNI ports belong to MSTP instance 0, Avaya recommends provisioning the non-SPBM C-VLANs in an MSTP instance other than 0.
- SPBM NNI ports are not part of the Layer 2 VSN C-VLAN, and BPDUs are not transmitted over the SPBM tunnel. SPBM can only guarantee loop-free topologies consisting of the NNI ports. Avaya recommends that you always use Simple Loop Prevention Protocol (SLPP) for loop prevention.

  Note:

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

- SPB internally uses spanning tree group (STG) 63 or Multiple Spanning Tree Instance (MSTI) 62. STG 63 or MSTI 62 cannot be used by another VLAN or MSTI. For non-SPB customer networks, if you use STG 63 or MSTI 62 in the configuration, you must delete STG 63 or MSTI 62 before you can configure SPBM.

- You must configure SPBM B-VLANs on all devices in the same MSTP region. MSTP requires this configuration to generate the correct digest.

SPBM IS-IS

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

- The current release does not support IP over IS-IS as defined by RFC 1195. IS-IS protocol is only to facilitate SPBM.
- The current release uses level 1 IS-IS. The current release does not support level 2 IS-IS. The ACLI command **`show isis int-l2-contl-pkts`** is not supported in the current release because the IEEE 802.1aq standard currently only defines the use of one hierarchy, Level 1.
- The IS-IS standard defines wide (32bit) metrics and narrow (8 bits) metrics. The current release supports the wide metric.

Pay special attention to the expected scaling of routes in the network when you select configuration values for the **`isis l1-hello-interval`** and **`isis l1-hello-multiplier`** commands on IS-IS interfaces. The default values for these commands work well for most networks, including those using moderately-scaled routes. In highly-scaled networks, you may need to configure higher values for these commands.

VLACP

VLACP is generally used when a repeater or switch exists between connected Virtual Services Platform 4000 switches to detect when a connection is down even when the link LED is lit.

SNMP traps

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

Other

The following identifies other restrictions and limitations:

- The current release does not support I-SID filters.

IP multicast over SPBM restrictions

Review the following restrictions for the IP multicast over SPBM feature.

IGMP

The BEB must be the only IGMP querier in the network. If the BEB receives an IGMP query from any other device, it causes unpredictable behavior, including traffic loss.

SPBM supports IGMP Snooping on a C-VLAN, but it does not support PIM on a C-VLAN. If you enable IGMP Snooping on a C-VLAN, then its operating mode is Layer 2 VSN with IP multicast over SPBM.

You must enable SSM snoop before you configure IGMP version 3, and you must enable both ssm-snoop and snooping for IGMPv3.

For IGMP Snooping, ensure that the IGMP version used by multicast hosts and other devices in the network is either the same as the IGMP version configured on the IGMP Snooping VLAN, or that compatibility mode is enabled.

SSM

If you delete any **`ssm-map`** in a static range group, the switch deletes the entire static range group. For example, create an ssm-map for 232.122.122.122 to 232.122.122.122.128 and after that configure this same range in a static group. If you delete any ssm-map between 232.122.122.122. to 232.122.122.128, the switch deletes the entire static range group.

Data I-SID

The BEB matches a single multicast stream to a particular data I-SID. As a result there is a one-to-one mapping between the S,G to data I-SID for each BEB.

IP address

IP multicast over SPBM only supports IPv4 multicast traffic in this release.

Supported services

Virtual Services Platform 4000 does not support IP multicast over SPBM routing on inter-VSN routing interfaces.

Virtual Services Platform 4000 supports the following modes of IP multicast over SPBM:

- Layer 2 VSN multicast service — Multicast traffic remains within the same Layer 2 VSN across the SPBM cloud.
- Layer 3 VSN multicast service — Multicast traffic remains within the same Layer 3 VSN across the SPBM cloud.
- IP Shortcuts multicast service — Multicast traffic can cross VLAN boundaries but remains confined to the subset of VLANs with the Global Routing Table that have IP multicast over SPBM enabled.

Chapter 12: IP multicast network design

Use multicast routing protocols to efficiently distribute a single data source among multiple users in the network. This section provides information about how to design networks that support IP multicast routing.

For more information about multicast routing, see *Avaya Virtual Services Platform 4000 Configuration — IP Multicast Routing Protocols*, NN46251-504.

For design guidelines on IP Multicast over SPBM, see SPBM design guidelines on page 73.

For more conceptual and configuration information on IP Multicast over SPBM, see *Configuring Avaya VENA Fabric Connect on Avaya Virtual Services Platform 4000*, NN46251-510.

Multicast and VRF-lite

You can configure multicast routing support with the Virtual Routing and Forwarding (VRF) Lite feature. This feature is known as multicast virtualization.

Multicast virtualization enables multiple VRF routing instances on devices and supports various unicast routing protocols so that you can provide the services of many virtual routers from one physical device.

VRF-lite configurations support IGMP.

Multicast virtualization provides support for:

- Virtualization of control and data plane
- Multicast routing table managers (MRTM)
- Virtualized IGMPv1, v2, and v3
- Support for overlapping multicast address spaces
- Support for the Global Router (VRF0)

Multicast and MultiLink Trunking considerations

Multicast traffic distribution is important because the bandwidth requirements can be substantial when a large number of streams are employed. Virtual Services Platform 4000 can distribute IP multicast streams over links of a multilink trunk using one of the methods in this section.

Multicast flow distribution over MLT

MultiLink Trunking (MLT) distributes multicast streams over a multilink trunk based on the source MAC address and the destination MAC address. As a result, the load is distributed on different ports of the MLT more evenly. This functionality is enabled by default on the VSP 4000 and cannot be manually configured.

Multicast scalability design rules

The following section lists the design rules to increase multicast route scaling.

Important:

The current release of the VSP 4000 does not support the following:

- PIM
- SMLT and RSMLT
- High Availability (HA)

Multicast scalability design rules

1. Whenever possible, use simple network designs that do not use VLANs that span several switches. Instead, use routed links to connect switches.
2. Whenever possible, group sources sending to the same group in the same subnet. Virtual Services Platform 4000 uses a single egress forwarding pointer for all sources in the same subnet sending to the same group. Be aware that these streams have separate hardware forwarding records on the ingress side.
3. Do not configure multicast routing on edge switch interfaces that do not contain multicast senders or receivers. By following this rule, you:
    - Provide secure control over multicast traffic that enters or exits the interface.
    - Reduce the load on the switch, as well as the number of routes. This improves overall performance and scalability.
4. Avoid initializing many (several hundred) multicast streams simultaneously. Initial stream setup is a resource-intensive task, and initializing a large number can increase the setup time. In some cases, this delay can result in stream loss.
5. Whenever possible, do not connect IP multicast sources and receivers by using VLANs that interconnect switches (see the following figure). In some cases, this can result in excessive hardware record use. By placing the source on the interconnected VLAN, traffic takes two paths to the destination, depending on the RPF checks and the shortest path to the source.

    For example, if a receiver is on VLAN 1 on switch S1 and another receiver is on VLAN 2 on switch S1, traffic can be received from two different paths to the two receivers, which results in the use of two forwarding records. If the source on switch

S2 is on a different VLAN than VLAN 3, traffic takes a single path to switch S1 where the receivers are located.

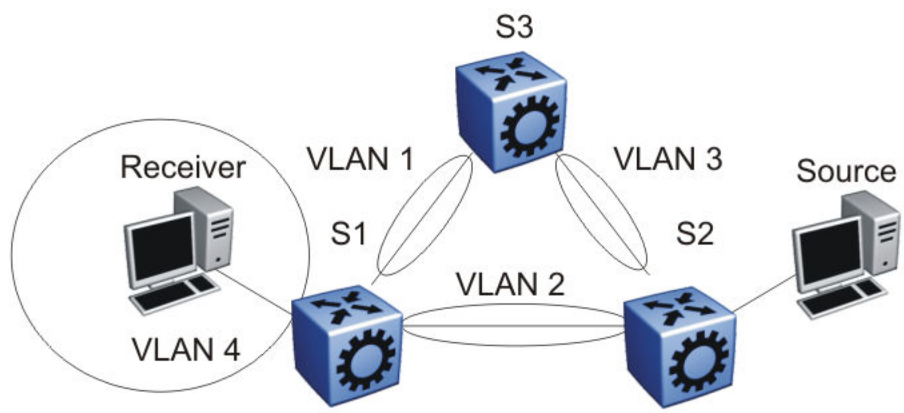

Figure 52: IP multicast sources and receivers on interconnected VLANs

6. Avaya recommends the use of Static group-range-to-RP mappings in an SMLT topology as opposed to RP set learning via the Bootstrap Router (BSR) mechanism. Static RP allows for faster convergence in box failure, reset, and HA failover scenarios; whereas there are inherent delays in the BSR mechanism as follows:

    - When a router comes back up after a failover or reset, to accept and propagate (*,g) Join requests from surrounding routers (either PIM Join messages or local IGMP membership reports) to the RP, a PIM router must determine the address of the RP for each group for which they desire (*,g) state. The PIM router needs to know the unicast route to the RP address. The route to the RP address is learned by using a unicast routing protocol such as OSPF, and the RP address is either statically configured or dynamically learned using the BSR mechanism.
    - When a system comes up after a reset or the standby CP becomes master after an HA failover, if the RP is not statically configured, it must wait for the BSR to select the RP from candidate RP routers, and then propagate the RP set hop-by-hop to all PIM routers. This must be done before a Join message can be processed. If the PIM router receives a Join message before it learns the RP set, the router drops the Join message and waits for another Join or Prune message to arrive before it creates the multicast router, and propagates the Join messages to the RP. The default Join/Prune timer is 60 seconds, and because of this and the delays inherent in BSR RP-set learning, significant multicast traffic interruptions can occur. If the RP is statically configured, the only delay is in the unicast routing table convergence and the arrival of the Join/Prune messages from surrounding boxes.

IP multicast address range restrictions

IP multicast routers use D class addresses, which range from 224.0.0.0 to 239.255.255.255. Although you can use subnet masks to configure IP multicast address ranges, the concept of subnets does not exist for multicast group addresses. Consequently, the usual unicast conventions—where you reserve the all 0s subnets, all 1s subnets, all 0s host addresses, and all 1s host addresses—do not apply.

Internet Assigned Numbers Authority (IANA) reserves addresses from 224.0.0.0 through 224.0.0.255 for link-local network applications. Multicast-capable routers do not forward packets with an address in this range. For example, Open Shortest Path First (OSPF) uses 224.0.0.5 and 224.0.0.6, and Virtual Router Redundancy Protocol (VRRP) uses 224.0.0.18 to communicate across local broadcast network segments.

IANA also reserves the range of 224.0.1.0 through 224.0.1.255 for well-known applications. IANA assigns these addresses to specific network applications. For example, the Network Time Protocol (NTP) uses 224.0.1.1, and Mtrace uses 224.0.1.32. RFC1700 contains a complete list of these reserved addresses.

Multicast addresses in the 232.0.0.0/8 (232.0.0.0 to 232.255.255.255) range are reserved only for source-specific multicast (SSM) applications, such as one-to-many applications. While this range is the publicly reserved range for SSM applications, private networks can use other address ranges for SSM.

Finally, addresses in the range 239.0.0.0/8 (239.0.0.0 to 239.255.255.255) are administratively scoped addresses; they are reserved for use in private domains. Do not advertise these addresses outside the private domain. This multicast range is analogous to the 10.0.0.0/8, 172.16.0.0/20, and 192.168.0.0/16 private address ranges in the unicast IP space.

In a private network, only assign multicast addresses from 224.0.2.0 through 238.255.255.255 to applications that are publicly accessible on the Internet. Assign addresses in the 239.0.0.0/8 range to multicast applications that are not publicly accessible.

Although you can use a multicast address you choose on your own private network, it is generally not good design practice to allocate public addresses to private network entities. Do not use public addresses for unicast host or multicast group addresses on private networks.

Multicast MAC address mapping considerations

Like IP, Ethernet has a range of multicast MAC addresses that natively support Layer 2 multicast capabilities. While IP has a total of 28 addressing bits available for multicast addresses, Ethernet has only 23 addressing bits assigned to IP multicast. The Ethernet multicast MAC address space is much larger than 23 bits, but only a subrange of that larger space is allocated to IP multicast. Because of this difference, 32 IP multicast addresses map to one Ethernet multicast MAC address.

IP multicast addresses map to Ethernet multicast MAC addresses by placing the low-order 23 bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01:00:5E:00:00:00. Thus, more than one multicast address maps to the same Ethernet address (see the following figure). For example, all 32 addresses 224.1.1.1, 224.129.1.1, 225.1.1.1, 225.129.1.1, 239.1.1.1, 239.129.1.1 map to the same 01:00:5E:01:01:01 multicast MAC address.

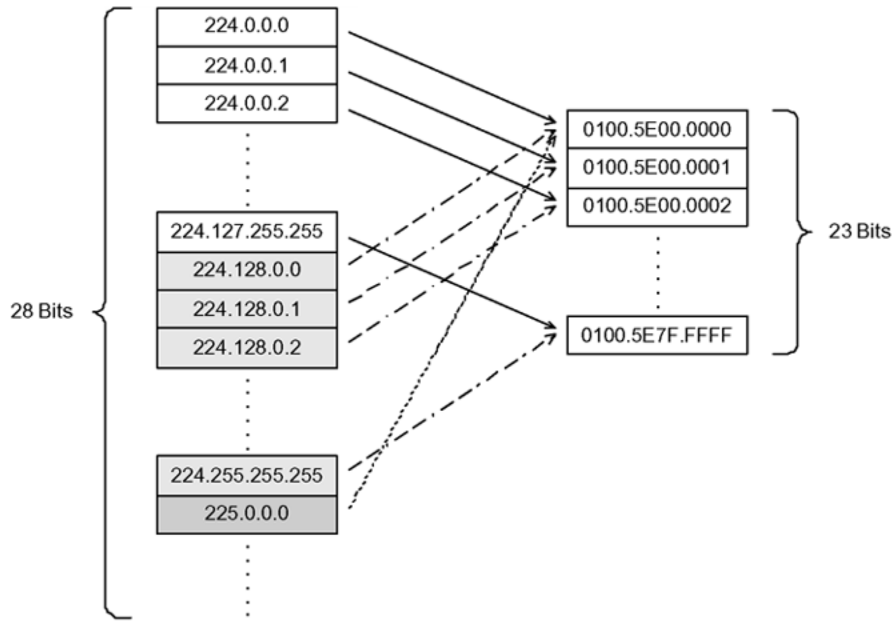

Figure 53: Multicast IP address to MAC address mapping

Most Ethernet switches handle Ethernet multicast by mapping a multicast MAC address to multiple switch ports in the MAC address table. Therefore, when you design the group addresses for multicast applications, take care to efficiently distribute streams only to hosts that are receivers. Virtual Services Platform 4000 switches IP multicast data based on the IP multicast address, not the MAC address, and thus, does not have this issue.

As an example, consider two active multicast streams using addresses 239.1.1.1 and 239.129.1.1. Suppose that two Ethernet hosts, receiver A and receiver B, connect to ports on the same switch and only want the stream addressed to 239.1.1.1. Suppose also that two other Ethernet hosts, receiver C and receiver D, also connect to the ports on the same switch as receiver A and B, and want to receive the stream addressed to 239.129.1.1. If the switch uses the Ethernet multicast MAC address to make forwarding decisions, then all four receivers receive both streams—even though each host only wants one stream. This transmission increases the load on both the hosts and the switch. To avoid this extra load, Avaya recommends that you manage the IP multicast group addresses used on the network.

Virtual Services Platform 4000 does not forward IP multicast packets based on multicast MAC addresses—even when bridging VLANs at Layer 2. Thus, the platform does not encounter this problem. Instead, the platform internally maps IP multicast group addresses to the ports that contain group members.

When an IP multicast packet is received, the lookup is based on the IP group address, regardless of whether the VLAN is bridged or routed. While Virtual Services Platform 4000 does not suffer from the problem described in the previous example, other switches in the network can. This problem is particularly true of pure Layer 2 switches.

In a network that includes non Virtual Services Platform 4000 equipment, the easiest way to ensure that this issue does not arise is to use only a consecutive range of IP multicast addresses that correspond to the lower order 23 bits of that range. For example, use an address range from 239.0.0.0 through 239.127.255.255. A group address range of this size can still easily accommodate the needs of even the largest private enterprise.

Dynamic multicast configuration changes

Avaya recommends that you do not perform dynamic multicast configuration changes when multicast streams flow in a network. For example, do not change the routing protocol that runs on an interface, or the IP address, or the subnet mask for an interface until multicast traffic ceases.

For such changes, Avaya recommends that you temporarily stop all multicast traffic. If the changes are necessary and you have no control over the applications that send multicast data, you can disable the multicast routing protocols before you perform the change. For example, consider disabling multicast routing before making interface address changes. In all cases, these changes result in traffic interruptions because they impact neighbor state machines and stream state machines.

In addition, Avaya recommends that when removing port members of an MLT group that you first disable the ports before removing them from that MLT group. Changing the group set without first shutting the ports down can result in high CPU utilization and processing in a scaled multicast environment due to the necessary hardware reprogramming on the multicast records.

IGMPv3 backward compatibility

IGMPv3 for PIM is backward compatible with IGMPv1/v2. According to RFC3376, the multicast router with IGMPv3 can use one of two methods to handle older query messages:

- If an older version of IGMP is present on the router, the querier must use the lowest version of IGMP present on the network.
- If a router that is not explicitly configured to use IGMPv1 or IGMPv2, hears an IGMPv1 query or IGMPv2 general query, it logs a rate-limited warning.

You can configure the IGMP version of an interface to version 3 regardless of the PIM or snooping mode.

You can configure whether the switch downgrades the version of IGMP to handle older query messages. If the switch downgrades, the host with IGMPv3 only capability does not work. If you do not configure the switch to downgrade the version of IGMP, the switch logs a warning.

Note:

If you enable the explicit host tracking option on an IGMPv3 interface, you cannot downgrade to IGMPv1 or IGMPv2. You must disable explicit host tracking to downgrade the IGMP version.

IGMP Layer 2 Querier

In a multicast network, if you only need to use Layer 2 switching for the multicast traffic, you do not need multicast routing. However, you must have an IGMP querier on the network for multicast traffic to flow from sources to receivers. A multicast router normally provides the IGMP querier function. You can use the IGMP Layer 2 Querier to provide a querier on a Layer 2 network without a multicast router.

The Layer 2 querier function originates queries for multicast receivers, and processes the responses accordingly. On the connected Layer 2 VLANs, IGMP snoop continues to provide services as normal. IGMP snoop responds to queries and identifies receivers for the multicast traffic.

You must enable Layer 2 querier and configure an IP address for the querier before it can originate IGMP query messages. If a multicast router exists on the network, Virtual Services Platform 4000 automatically disables the Layer 2 querier.

In a Layer 2 multicast network, enable Layer 2 querier on only one of the switches in the VLAN. A Layer 2 multicast domain supports only one Layer 2 querier. No querier election exists.

For more information about how to configure IGMP Layer 2 Querier, see *Avaya Virtual Services Platform 4000 Configuration — IP Multicast Routing Protocols*, NN46251–504.

TTL in IP multicast packets

Virtual Services Platform 4000 treats multicast data packets with a time-to-live (TTL) of 1 as expired packets and sends them to the CPU before dropping them. To avoid this, ensure that the originating application uses a hop count large enough to enable the multicast stream to traverse the network and reach all destinations without reaching a TTL of 1. Avaya recommends that you use a TTL value of 33 or 34 to minimize the effect of looping in an unstable network.

Multicast MAC filtering

Certain network applications, such as the Microsoft Network Load Balancing solution, require multiple hosts to share a multicast MAC address. Instead of flooding all ports in the VLAN with this multicast traffic, you can use Multicast MAC filtering to forward traffic to a configured subset of the ports in the VLAN. This multicast MAC address is not an IP multicast MAC address.

At a minimum, map the multicast MAC address to a set of ports within the VLAN. In addition, if traffic is routed on the local Virtual Services Platform 4000, you must configure an Address Resolution Protocol (ARP) entry to map the shared unicast IP address to the shared multicast MAC address. You must configure an ARP entry because the hosts can also share a virtual IP address, and packets addressed to the virtual IP address need to reach each host.

Avaya recommends that you limit the number of such configured multicast MAC addresses to a maximum of 100. This number is related to the maximum number of possible VLANs you can configure because for every multicast MAC filter that you configure the maximum number of configurable VLANs reduces by one. Similarly, configuring large numbers of VLANs reduces the maximum number of configurable multicast MAC filters downwards from 100.

Although you can configure addresses starting with 01.00.5E, which are reserved for IP multicast address mapping, do not enable IP multicast with streams that match the configured addresses. This can result in incorrect IP multicast forwarding and incorrect multicast MAC filtering.

Guidelines for multicast access policies

Use the following guidelines when you configure multicast access policies:

- Use masks to specify a range of hosts. For example, 10.177.10.8 with a mask of 255.255.255.248 matches hosts addresses 10.177.10.8 through 10.177.10.15. The host subnet address and the host mask must be equal to the host subnet address. An easy way to determine this is to ensure that the mask has an equal or fewer number of trailing zeros than the host subnet address. For example, 3.3.0.0/255.255.0.0 and 3.3.0.0/255.255.255.0 are valid. However, 3.3.0.0/255.0.0.0 is not.
- Apply receive access policies to all eligible receivers on a segment. Otherwise, one host joining a group makes that multicast stream available to all.
- Receive access policies are initiated after the switch receives reports with addresses that match the filter criteria.
- Transmit access policies apply after the switch receives the first packet of a multicast stream.

Multicast access policies can apply to a routed PIM interface if Internet Group Management Protocol (IGMP) reports the reception of multicast traffic.

The following rules and limitations apply to IGMP access policy parameters when you use them with IGMP instead of PIM:

- The static member parameter applies to IGMP snooping and PIM on both interconnected links and edge ports.
- The Static Not Allowed to Join parameter applies to IGMP snooping and PIM on both interconnected links and edge ports.
- For multicast access control, the denyRx parameter applies to IGMP snooping and PIM. The DenyTx and DenyBoth parameters apply only to IGMP snooping.

Multicast for multimedia

Virtual Services Platform 4000 provides a flexible and scalable multicast implementation for multimedia applications. Several features are dedicated to multimedia applications and in particular, to television distribution.

Join and leave performance

For TV applications, you can attach several TVs directly, or through an IGMP-capable ethernet switch, to the Virtual Services Platform 4000. Base this implementation on IGMP; the set-top boxes use IGMP reports to join a TV channel and IGMP leaves to exit the channel. After a viewer changes channels, an IGMPv2 leave for the old channel (multicast group) is issued, and a membership report for the new channel is sent. If viewers change channels continuously, the number of joins and leaves can become large, particularly if many viewers attach to the switch.

Virtual Services Platform 4000 supports more than a thousand joins and leaves per second, which is well adapted to TV applications.

Important:

For IGMPv3, Avaya recommends that you ensure a join rate of 1000 per second or less. This ensures the timely processing of join requests.

If you use the IGMP proxy functionality at the receiver edge, you reduce the number of IGMP reports received by Virtual Services Platform 4000. This provides better overall performance and scalability.

Fast Leave

IGMP Fast Leave supports two modes of operation: single user mode and multiple user mode.

In single user mode, if more than one member of a group is on the port and one of the group members leaves the group, everyone stops receiving traffic for this group. A group-specific-query is not sent before the effective leave takes place.

Multiple user mode allows several users on the same port or VLAN. If one user leaves the group and other receivers exist for the same stream, the stream continues. The switch tracks the number of receivers that join a given group. For multiple user mode to operate properly, do not suppress reports. This ensures that the switch properly tracks the correct number of receivers on an interface.

The Fast Leave feature is particularly useful in IGMP-based TV distribution where only one receiver of a TV channel connects to a port. In the event that a viewer changes channels quickly, you create considerable bandwidth savings if you use Fast Leave.

You can implement Fast Leave on a VLAN and port combination; a port that belongs to two different VLANs can have Fast Leave enabled on one VLAN (but not on the other). Thus, with the Fast Leave feature enabled, you can connect several devices on different VLANs to the same port. This strategy does not impact traffic after one device leaves a group to which another device subscribes. For example, you can use this feature when two TVs connect to a port through two set-top boxes, even if you use the single user mode.

To use Fast Leave, you must first enable explicit host tracking. IGMP uses explicit host tracking to track all source and group members. Explicit host tracking is disabled by default. For configuration information, see *Avaya Virtual Services Platform 4000 Configuration — IP Multicast Routing Protocols*, NN46251–504.

Last member query interval tuning

If an IGMPv2 host leaves a group, it notifies the router by using a leave message. Because of the IGMPv2 report suppression mechanism, the router is unaware of other hosts that require the stream. Thus, the router broadcasts a group-specific query message with a maximum response time equal to the last member query interval (LMQI).

Because this timer affects the latency between the time that the last member leaves and when the stream actually stops, you must properly tune this parameter. This timer can especially affect TV delivery or other large-scale, high-bandwidth multimedia applications. For instance, if you assign a value that is too low, this can lead to a storm of membership reports if a large number of hosts are subscribed. Similarly, assigning a value that is too high can cause unwanted high-bandwidth stream propagation across the network if users change channels rapidly. Leave latency also depends on the robustness value, so a value of two equates to a leave latency of twice the LMQI.

Determine the proper LMQI value for your particular network through testing. If a very large number of users connect to a port, assigning a value of three can lead to a storm of report messages after a group-specific query is sent. Conversely, if streams frequently start and stop in short intervals, as in a TV delivery network, assigning a value of ten can lead to frequent congestion in the core network.

Another performance-affecting factor that you need to be aware of is the error rate of the physical medium. For links that have high packet loss, you can find it necessary to adjust the robustness variable to a higher value to compensate for the possible loss of IGMP queries and reports.

In such cases, leave latency is adversely impacted as numerous group-specific queries are unanswered before the stream is pruned. The number of unanswered queries is equal to the robustness variable (default two). The assignment of a lower LMQI can counterbalance this effect. However, if you configure the LMQI too low, it can actually exacerbate the problem by inducing storms of reports on the network. LMQI values of three and ten, with a robustness value of two, translate to leave latencies of six tenths of a second and two seconds, respectively.

When you choose an LMQI, consider all of these factors to determine the best configuration for the given application and network. Test that value to ensure that it provides the best performance.

Important:

In networks that have only one user connected to each port, Avaya recommends that you use the Fast Leave feature instead of LMQI, because no wait is required before the stream stops. Similarly, the robustness variable does not impact the Fast Leave feature, which is an additional benefit for links with high loss.

Chapter 13: System and network stability and security

Use the information in this section to design and implement a secure network.

You must provide security mechanisms to prevent your network from attack. If links become congested due to attacks, you can immediately halt end-user services. During the design phase, study availability issues for each layer.

To provide additional network security, you can use the Avaya VSP 9000 or your own high-performance stateful firewalls.

DoS protection mechanisms

Several internal mechanisms and features protect Virtual Services Platform 4000 against Denial-of-Service (DoS) attacks.

Broadcast and multicast rate limiting

To protect the switch and other devices from excessive broadcast traffic, you can use broadcast and multicast rate limiting on an individual port basis.

For more information about how to configure the rate limits for broadcast or multicast packets on a port, see *Avaya Virtual Services Platform 4000 Configuration — QoS and IP Filtering*, NN46251-502.

Directed broadcast suppression

You can enable or disable forwarding for directed broadcast traffic on an IP-interface basis. A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. By disabling or suppressing directed broadcasts on an interface, you cause all frames sent to the subnet broadcast address for a local router interface to be dropped. Directed broadcast suppression protects hosts from possible DoS attacks.

To prevent the flooding of other networks with DoS attacks, such as the Smurf attack, Virtual Services Platform 4000 is protected by directed broadcast suppression. This feature is enabled by default. Avaya recommends that you not disable it.

For more information about directed broadcast suppression, see *Avaya Virtual Services Platform 4000 Security*, NN46251-601.

Prioritization of control traffic

Virtual Services Platform 4000 uses a sophisticated prioritization scheme to schedule control packets on physical ports. This scheme involves two levels with both hardware and software queues to guarantee proper handling of control packets regardless of the switch load. In turn, this guarantees the stability of the network. Prioritization also guarantees that applications that use many broadcasts are handled with lower priority.

You cannot view, configure, or modify control traffic queues.

ARP request threshold recommendations

The Address Resoluion Protocol (ARP) request-threshold defines the maximum number of outstanding, unresolved ARP requests. The default value for this function is 500 ARP requests. To avoid excessive amounts of subnet scanning that a virus can cause, Avaya recommends that you change the ARP request threshold to a value between 100 to 50. This configuration protects the CPU from causing excessive ARP requests, protects the network, and lessens the spread of the virus to other PCs. The following list provides further recommended ARP threshold values:

- default: 500
- severe conditions: 50
- continuous scanning conditions: 100
- moderate: 200
- relaxed: 500

For more information about how to configure the ARP threshold, see *Avaya Virtual Services Platform 4000 Configuration — IP Routing* , NN46251-505.

Multicast Learning Limitation

The Multicast Learning Limitation feature protects the CPU from multicast data packet bursts generated by malicious applications. If more than a certain number of multicast streams enter the CPU through a port during a sampling interval, the port is shut down until the user or administrator takes the appropriate action.

For more information, see *Avaya Virtual Services Platform 4000 Configuration — IP Multicast Routing Protocols*, NN46251-504.

Damage prevention

To further reduce the chance that unauthorized users can use your network to damage other existing networks, take the following actions:

1. Prevent IP spoofing.

   You can use the spoof-detect feature.
2. Prevent the use of the network as a broadcast amplification site.
3. To block illegal IP addresses, enable the **`hsecure`** flag (High Secure mode).

For more information, see *Avaya Virtual Services Platform 4000 Security*, NN46251-601.

4. Prevent unknown devices from influencing the spanning tree topology.

Packet spoofing

You can stop spoofed IP packets by configuring the switch to only forward IP packets that contain the correct source IP address of your network. By denying all invalid source IP addresses, you minimize the chance that your network is the source of a spoofed DoS attack.

A spoofed packet is one that comes from the Internet into your network with a source address equal to one of the subnet addresses on your network. The source address belongs to one of the address blocks or subnets on your network. To provide spoofing protection, you can use a filter that examines the source address of all outside packets. If that address belongs to an internal network or a firewall, the packet is dropped.

To prevent DoS attack packets that come from your network with valid source addresses, you need to know the IP network blocks in use. You can create a generic filter that:

- permits valid source addresses
- denies all other source addresses

To do so, configure an ingress filter that drops all traffic based on the source address that belongs to your network.

If you do not know the address space completely, it is important that you at least deny private (see RFC1918) and reserved source IP addresses. The following table lists the source addresses to filter.

Table 12: Source addresses to filter

Address	Description
0.0.0.0/8	Historical broadcast. High Secure mode blocks addresses 0.0.0.0/8 and 255.255.255.255/16. If you enable this mode, you do not need to filter these addresses.
10.0.0.0/8	RFC1918 private network
127.0.0.0/8	Loopback
169.254.0.0/16	Link local networks
172.16.0.0/12	RFC1918 private network
192.0.2.0/24	TEST-NET
192.168.0.0/16	RFC1918 private network
224.0.0.0/4	Class D multicast
240.0.0.0/5	Class E reserved
248.0.0.0/5	Unallocated
255.255.255.255/32	Broadcast1

You can also enable the spoof-detect feature on a port.

For more information about the spoof-detect feature, see *Avaya Virtual Services Platform 4000 Configuration — VLANs and Spanning Tree*, NN46251-500.

High Secure mode

To ensure that Virtual Services Platform 4000 does not route packets with an illegal source address of 255.255.255.255 (RFC1812 Section 4.2.2.11 and RFC971 Section 3.2), you can enable High Secure mode.

By default, this feature is disabled. After you enable this flag, the feature applies to all ports.

For more information about hsecure, see *Avaya Virtual Services Platform 4000 Security*, NN46251-601.

Data plane security

Data plane security mechanisms include VLANs, filters, routing policies, and routing protocol protection.

VLANs and traffic isolation

You can use Virtual Services Platform 4000 to build secure VLANs. If you configure port-based VLANs, each VLAN is completely separate from the others. Virtual Services Platform 4000 supports the IEEE 802.1Q specification for tagging frames and coordinating VLANs across multiple switches.

Virtual Services Platform 4000 analyzes each packet independently of preceding packets. This mode, as opposed to the cache mode that other vendors use, allows complete traffic isolation.

For more information about VLANs, see *Avaya Virtual Services Platform 4000 Configuration — VLANs and Spanning Tree*, NN46251-500.

Management of access policies

At Layer 2, Virtual Services Platform 4000 provides the following security mechanisms:

- access policies

  If you enable access policies globally, the system creates a default policy (1) that allows File Transfer Protocol (FTP), Hypertext Transfer Protocol (HTTP), Telnet, and Secure Shell (SSH). If you enable access policies globally but disable the default policy, the system denies FTP, HTTP, rlogin, SSH, Simple Network Management Protocol (SNMP), Telnet, and Trivial FTP (TFTP).

  The **access-strict** parameter ties to the **accesslevel** parameter. If you enable **access-strict**, the access policy looks at the **accesslevel** parameter, and only applies to that access level. Use the following configuration as an example:

  ```
  VSP-9012:1(config)#show access-policy

    AccessPolicyEnable: off
  ```

```
                Id: 1
              Name: default
      PolicyEnable: false
              Mode: allow
           Service: ftp|http|telnet|ssh
        Precedence: 128
       NetAddrType: any
           NetAddr: N/A
           NetMask: N/A
   TrustedHostAddr: N/A
TrustedHostUserName: none
       AccessLevel: readOnly
      AccessStrict: false
             Usage: 0
```

If you disable **access-strict** (false), the policy looks at the value for **accesslevel**, and then the system applies the policy to anyone with equivalent rights or higher. In this example, all levels include readonly so the default policy applies to l1, l2, l3, rw, ro, and rwa. If you enable **access-strict**, the system applies the policy only to ro.

For SNMP and access policies, you must apply the service to the access policy - the only choice is snmpv3 but this parameter applies to all versions of SNMP. The additional command **access-policy <1-65535> snmp-group WORD<1-32> <snmpv1|snmpv2|usm>** applies the policy to the SNMP community or the SNMP group.

- filters

  ACL filters are used by individual VLANs to filter out packets based on source MAC, destination MAC and other criteria.

  For more information about these filters, see *Avaya Virtual Services Platform 4000 Configuration — QoS and IP Filtering*, NN46251-502.

- limited MAC learning

  This feature limits the number of FDB-entries learned on a particular port to a user-specified value. After the number of learned FDB-entries reaches the maximum limit, the switch drops packets with unknown source MAC addresses.

  Note:

  The current release of the VSP 4000 allows you to enable limit-learning on a port and configure the maximum number of MAC entries on this port.

```
VSP-4850GTS(config-if)#mac-security limit-learning ?
  enable     Enable limit-learning on this port
  max-addrs  Set the maximum number of entries on this port
```

Security at Layer 3: filtering

At Layer 3 and higher, Virtual Services Platform 4000 provides enhanced filtering capabilities as part of its security strategy to protect the network from different attacks.

Virtual Services Platform 4000 supports advanced filters based on Access Control Lists (ACL).

Customer Support Bulletins (CSBs) are available on the Avaya Technical Support Web site to provide information and configuration examples about how to block some attacks.

Routing protocol security

You can protect OSPF and BGP updates with a Message Digest 5 (MD5) key on each interface. At most, you can configure two MD5 keys for each interface. You can also use multiple MD5 key configurations for MD5 transitions without bringing down an interface.

For more information, see *Avaya Virtual Services Platform 4000 Configuration — OSPF and RIP*, NN46251-506 and *Avaya Virtual Services Platform 4000 Configuration — BGP Services*, NN46251-507.

Control plane security

The control plane physically separates management traffic using the in-band interface. The control plane facilitates High Secure mode, access policies, authentication, SSH and Secure Copy, and SNMP.

High Secure mode

Use High Secure to disable all unsecured applications and daemons, for example, FTP, TFTP, and rlogin. Avaya strongly recommends that you do not use unsecured protocols. See also High Secure mode on page 122.

Use Secure Copy (SCP) rather than FTP or TFTP.

Security and access policies

Access policies permit secure switch access by specifying a list of IP addresses or subnets that can manage the switch for a specific daemon, such as Telnet, SNMP, HTTP, SSH, TFTP, FTP, RSH, and rlogin. Rather than using a management VLAN that is spread out among all of the switches in the network, you can build a full Layer 3 routed network and securely manage the switch with one of the in-band IP addresses attached to one of the VLANs (see the following figure).

You can use route policies to selectively accept or announce some networks and to block the propagation of some routes. Route policies enhance the security in a network by hiding the visibility of some networks (subnets) from other parts of the network.

You can apply one policy for one purpose. For example, you can apply a RIP announce policy on a given RIP interface. In such cases, all sequence numbers under the given policy apply to that filter. A sequence number also acts as an implicit preference (that is, a lower sequence number is preferred).

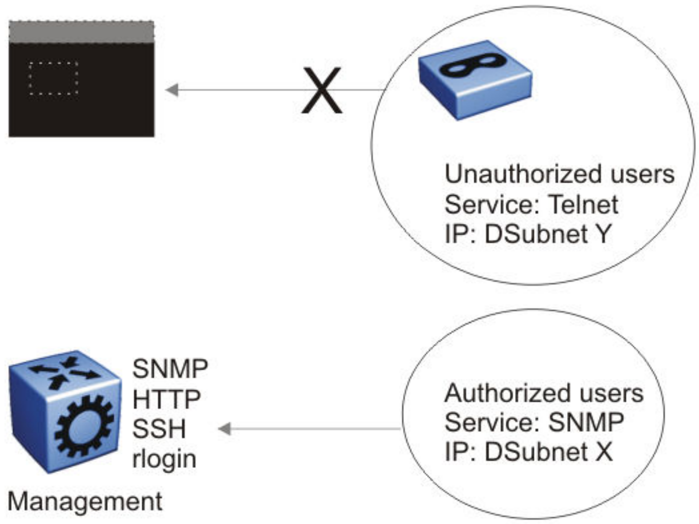

Figure 54: Access levels

Avaya recommends that you use access policies for in-band management to secure access to the switch. By default, all services are denied. You must enable the default policy or enable a custom policy to provide access. A lower precedence takes higher priority if you use multiple policies. Preference 120 has priority over preference 128.

RADIUS authentication

You can enforce access control by using RADIUS. RADIUS provides a high degree of security against unauthorized access and centralizes the knowledge of security access based on a client and server architecture. The database within the RADIUS server stores a list of pertinent information about client information, user information, password, and access privileges including the use of the shared secret.

When the switch acts as a Network Access Server, it operates as a RADIUS client. The switch is responsible for passing user information to the designated RADIUS servers. Because the switch operates in a LAN environment, it allows user access through Telnet, rlogin, and console logon.

You can configure a list of up to 10 RADIUS servers on the switch. If the first server is unavailable, Virtual Services Platform 4000 tries the second, and so on, until it establishes a successful connection.

RADIUS authentication supports: WEB, CLI, SNMP, or Extensible Authentication Protocol over LAN (EAPoL). You can configure a list of up to 10 RADIUS servers for all four methods combined. If you configure six servers for EAPoL, you can configure four servers for the other methods.

You can use the RADIUS server as a proxy for stronger authentication (see the following figure), such as:

- SecurID cards
- Kerberos

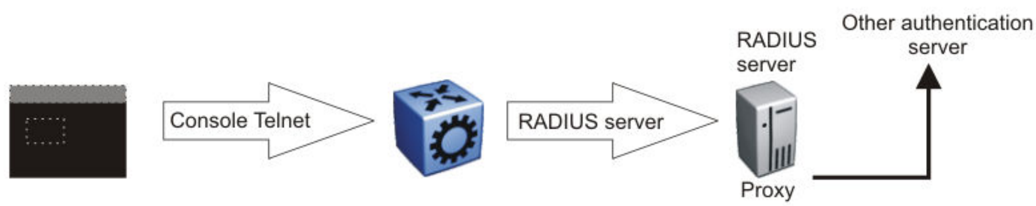

Figure 55: RADIUS server as proxy for stronger authentication

You must configure each RADIUS client to contact the RADIUS server. When you configure a client to work with a RADIUS server, complete the following configurations:

- Enable RADIUS.
- Provide the IP address of the RADIUS server.
- Ensure the shared secret matches what is defined in the RADIUS server.
- Provide the attribute value.
- Provide the use by value.

  The use by value can be CLI, SNMP, IGAP, or EAPoL.
- Indicate the order of priority in which the RADIUS server is used. (Order is essential when more than one RADIUS server exists in the network.)
- Specify the User Datagram Protocol (UDP) port that the client and server use during the authentication process. The UDP port between the client and the server must have the same or equal value. For example, if you configure the server with UDP 1812, the client must use the same UDP port value.

Other customizable RADIUS parameters require careful planning and consideration, for example, switch timeout and retry. Use the switch timeout to define the number of seconds before the authentication request expires. Use the retry parameter to indicate the number of retries the server accepts before sending an authentication request failure.

Avaya recommends that you use the default value in the attribute-identifier field. If you change the default value, you must alter the dictionary on the RADIUS server with the new value. To configure the RADIUS feature, you require Read-Write-All access to the switch.

For more information about RADIUS, see *Avaya Virtual Services Platform 4000 Security*, NN46251-601.

Encryption of control plane traffic

Control plane traffic encryption involves SSHv1/v2, SCP, and SNMPv3.

Use SSH to conduct secure communications over a network between a server and a client. The switch supports only the server mode (supply an external client to establish communication). The server mode supports SSHv1 and SSHv2.

The SSH protocol offers

- Authentication

SSH determines identities. During the logon process, the SSH client asks for digital proof of the identity of the user.

- Encryption

  SSH uses encryption algorithms to scramble data. This data is rendered unintelligible except to the intended receiver.

- Integrity

  SSH guarantees that data is transmitted from the sender to the receiver without alteration. If a third party captures and modifies the traffic, SSH detects this alteration.

Virtual Services Platform 4000 supports

- SSH version 1, with password and Rivest, Shamir, Adleman (RSA) authentication
- SSH version 2 with password and Digital Signature Algorithm (DSA) authentication
- Digital Encryption Standard (DES)
- Triple DES (3DES)
- Advanced Encryption Standard (AES)

You must load the encryption module before you can enable it. For more information about how to load encryption modules, see *Avaya Virtual Services Platform 4000 Security*, NN46251-601.

SNMP header network address

You can direct an IP header to have the same source address as the management virtual IP address for self-generated UDP packets. If you configure a management virtual IP address and enable the **`udpsrc-by-vip`** flag, the network address in the SNMP header is always the management virtual IP address. This configuration is true for all traps routed out on the I/O ports or on the out-of-band management Ethernet port.

SNMPv3 support

SNMP version 1 and version 2 are not secure because communities are not encrypted.

Avaya strongly recommends that you use SNMP version 3. SNMPv3 provides stronger authentication services and the encryption of data traffic for network management.

Other security equipment

Avaya offers other devices that increase the security of your network.

For sophisticated state-aware packet filtering (real stateful inspection), you can add an external firewall to the architecture. State-aware firewalls can recognize and track application flows that use not only static TCP and UDP ports, like Telnet or HTTP, but also applications that create and use dynamic ports, such as FTP, and audio and video streaming. For every packet, the state-aware firewall finds a matching flow and conversation.

The following figure shows a typical configuration used in firewall load balancing.

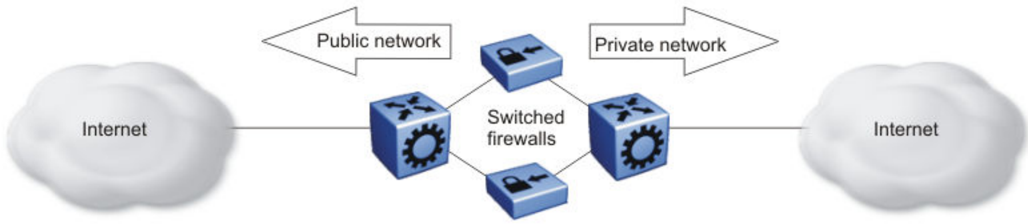

Figure 56: Firewall load balancing configuration

Use this configuration to redirect incoming and outgoing traffic to a group of firewalls and to automatically load balance across multiple firewalls. The benefits of such a configuration are

- increased firewall performance
- reduced response time
- redundant firewalls ensure Internet access

Virtual private networks (VPN) replace the physical connection between the remote client and access server with an encrypted tunnel over a public network. VPN technology employs IP security (IPsec) and Secure Sockets Layer (SSL) services.

Several Avaya products support IPSec and SSL, including Avaya VPN Gateway and Secure Router.

Additional information

The following organizations provide the most up-to-date information about network security attacks and recommendations about good practices:

- The Center of Internet Security Expertise (CERT)
- The Research and Education Organization for Network Administrators and Security Professionals (SANS)
- The Computer Security Institute (CSI)

Chapter 14: QoS design guidelines

This section provides design guidelines to provide Quality of Service (QoS) to user traffic on the network.

For more information about fundamental QoS mechanisms, and how to configure QoS, see *Avaya Virtual Services Platform 4000 Configuration — QoS and IP Filtering*, NN46251-502.

QoS mechanisms

Virtual Services Platform 4000 has a solid, well-defined architecture to handle QoS in an efficient and effective manner. The following sections briefly describe several QoS mechanisms that the platform uses.

QoS classification and mapping

Virtual Services Platform 4000 provides a hardware-based QoS platform through hardware packet classification. Packet classification is based on the examination of the QoS fields within the Ethernet packet, primarily the DiffServ Codepoint (DSCP) and the 802.1p fields.

You can configure ingress interfaces in one of two ways. In the first type of configuration, the interface does not classify traffic, but it forwards the traffic based on the packet markings. This mode of operation applies to trusted interfaces (core port mode) because the DSCP or 802.1p field is trusted to be correct, and the edge switch performs the mapping without classification.

In the second type of configuration, the interface classifies traffic as it enters the port, and marks the packet for further treatment as it traverses Virtual Services Platform 4000 network. This mode of operation applies to untrusted interfaces (access port mode) because the DSCP or 802.1p field is not trusted to be correct.

Virtual Services Platform 4000 assigns an internal QoS level to each packet that enters a port.

The Avaya QoS strategy simplifies QoS implementation by providing a mapping of various traffic types and categories to a Class of Service. These service classes are termed Avaya Service Classes (ASC). The following table provides a summary of the mappings and their typical traffic types.

Table 13: Traffic categories and ASC mappings

Traffic category	Application example	ASC
Network Control	Alarms and heartbeats	Critical
	Routing table updates	Network

Traffic category		Application example	ASC
Real-Time, Delay Intolerant		IP telephony; interhuman communication	Premium
Real-Time, Delay Tolerant		Video conferencing; interhuman communication.	Platinum
		Audio and video on demand; human-host communication	Gold
NonReal-Time Mission Critical	Interactive	eBusiness (B2B, B2C) transaction processing	Silver
	NonInteractive	Email; store and forward	Bronze
NonReal Time, NonMission Critical		FTP; best effort	Standard
		PointCast; Background/standby	Custom/ best effort

QoS and filters

Filters help you provide QoS by permitting or dropping traffic based on the parameters you configure. You can use filters to mark packets for specific treatment.

Typically, filters act as firewalls or are used for Layer 3 redirection. In more advanced cases, traffic filters can identify Layer 3 and Layer 4 traffic streams. The filters cause the streams to be re-marked and classified to attain a specific QoS level at both Layer 2 (802.1p) and Layer 3 (DSCP).

Traffic filtering is a key QoS feature. Virtual Services Platform 4000, by default, determines incoming packet 802.1p or DiffServ markings, and forwards traffic based on their assigned QoS levels. However, situations exist where the markings are incorrect, or the originating user application does not have 802.1p or DiffServ marking capabilities. Also, you can give a higher priority to select users (executive class). In these situations, use filters to prioritize specific traffic streams.

You can use filters to assign QoS levels to devices and applications. To help you decide whether or not to use a filter, key questions include:

1. Does the user or application have the ability to mark QoS information on data packets?
2. Is the traffic source trusted? Are the QoS levels configured appropriately for each data source?

   Users can maliciously configure QoS levels on their devices to take advantage of higher priority levels.
3. Do you want to prioritize traffic streams?

This decision-making process is outlined in the following figure.

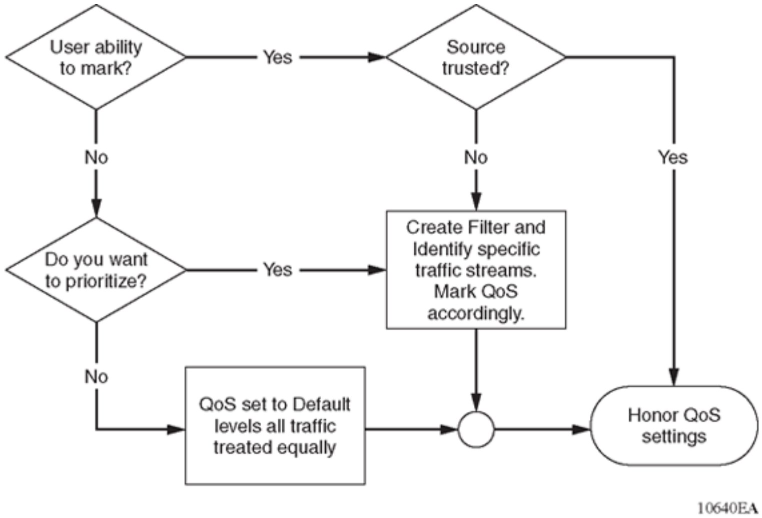

Figure 57: Filter decision-making process

Configure filters through the use of Access Control Lists (ACL) and Access Control Entries (ACE), which are implemented in hardware. An ACL can include both security and QoS type ACEs. The platform supports 2048 ACLs and 1000 ACEs for each ACL to a maximum of 16 000 ACEs for each plaform.

The following steps summarize the filter configuration process:

1. Determine your desired match fields.
2. Create an ACL.
3. Create an ACE within the ACL.
4. Configure the desired precedence, traffic type, and action.

   You determine the traffic type by creating an ingress or egress ACL.
5. Modify the parameters for the ACE.

Policing and shaping

As part of the filtering process, you can police ingress traffic. Policing is performed according to the traffic filter profile assigned to the traffic flow. For enterprise networks, policing ensures that traffic flows conform to the criteria assigned by network managers.

Traffic policers identify traffic using a traffic policy. Traffic that conforms to this policy is guaranteed for transmission, whereas nonconforming traffic is considered to be in violation. Traffic policers drop packets if traffic is excessive, or remark the DSCP or 802.1p markings by using filter actions. With Virtual Services Platform 4000, you can define multiple actions in case of traffic violation.

For service providers, policing at the network edge provides different bandwidth options as part of a Service Level Agreement (SLA). For example, in an enterprise network, you can police the traffic rate from one department to give critical traffic unlimited access to the network. In a service provider network, you can control the amount of traffic customers send to ensure that they comply with their SLA. Policing ensures that users do not exceed their traffic contract for a QoS level.

VSP 4000 supports two-rate, three-color marking for policers as described in RFC 2698. Policers mark packets as Green, Yellow, or Red. Red packets are dropped automatically. Out of profile packets cannot be re-marked to a lower QoS level.

The system can perform rate metering only on a Layer 3 basis.

Traffic shapers buffer and delay violating traffic. These operations occur at the egress level. Virtual Services Platform 4000 supports traffic shaping at the port level.

QoS interface considerations

Four QoS interface types are explained in detail in the following sections. You can configure an interface as trusted or untrusted, and for bridging or routing operations. Use these parameters to properly apply QoS to network traffic.

Trusted and untrusted interfaces

You can configure an interface as trusted (core) or untrusted (access). The default is trusted (core).

Use trusted interfaces (core) to mark traffic in a specific way, and to ensure that packets are treated according to the service level of those markings. Use a core interface if you need control over network traffic prioritization. For example, use 802.1p-bits to apply desired class of service (CoS) attributes to the packets before they are forwarded to the access node. You can also classify other protocol types ahead of IP packets.

A core port preserves the DSCP and 802.1p-bits markings. The device uses these values to assign a corresponding QoS level to the packets.

Use an access port to control the classification and mapping of traffic for delivery through the network. Untrusted interfaces require you to configure filter sets to classify and re-mark ingress traffic. For untrusted interfaces in the packet forwarding path, the DSCP is mapped to an IEEE 802.1p user priority field in the IEEE 802.1Q frame, and both of these fields are mapped to an IP Layer 2 drop precedence value that determines the forwarding treatment at each network node along the path. Traffic that enters an access port is re-marked with the appropriate DSCP and 802.1p markings, and given an internal QoS level. The switch performs this re-marking based on the filters and traffic policies that you configure.

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

Table 14: Data packet ingress mapping

Enable DiffServ	Access DiffServ	802.1p Override	Routed Packet	Tagged Ingress Packet	Internal QoS Derived From	Egress Packet DSCP Derived from	Egress Packet 802.1p Derived from
1	0, L3T=1	0, L2T=1	1	1	DSCP	Stays untouched	iQoS

Enable DiffServ	Access DiffServ	802.1p Override	Routed Packet	Tagged Ingress Packet	Internal QoS Derived From	Egress Packet DSCP Derived from	Egress Packet 802.1p Derived from
1	0, L3T=1	0, L2T=1	0	1	.1p	Stays untouched	iQoS
1	0, L3T=1	0, L2T=1	X	0	DCSP	Stays untouched	iQoS
1	1, L3T=0	0, L2T=1	X	1	.1p	iQoS	iQoS
1	1, L3T=0	0, L2T=1	X	0	Port QoS	iQoS	iQoS
0	X, L3T=0	0, L2T=1	X	1	.1p	Stays untouched	iQoS
0	X, L3T=0	0, L2T=1	X	0	Port QoS	Stays untouched	iQoS
1	0, L3T=1	1, L2T=0	X	X	DSCP	Stays untouched	iQoS
1	1, L3T=0	1, L2T=0	X	X	Port QoS	iQoS	iQoS

Bridged and routed traffic

In a service provider network, access nodes use Virtual Services Platform 4000 for bridging. In this case, Virtual Services Platform 4000 uses DiffServ to manage network traffic and resources, but some QoS features are unavailable in the bridging mode of operation.

In an enterprise network, access nodes use Virtual Services Platform 4000 for bridging, and core nodes use it for routing. For bridging, ingress traffic is mapped from the 802.1p-bit marking to a QoS level. For routing, ingress traffic is mapped from the DSCP marking to the appropriate QoS level.

802.1p and 802.1Q recommendations

In a network, to map the 802.1p user priority bits, use 802.1Q-tagged encapsulation on customer premises equipment (CPE). You require encapsulation because Virtual Services Platform 4000 does not provide classification when it operates in bridging mode.

To ensure consistent Layer 2 QoS boundaries within the service provider network, you must use 802.1Q encapsulation to connect a CPE directly to Virtual Services Platform 4000 access node. If you do not require packet classification, use Ethernet Routing Switch 5600 to connect to the access node. In this case, configure the traffic classification functions in the Ethernet Routing Switch 5600.

At the egress access node, packets are examined to determine if their IEEE 802.1p or DSCP values must be re-marked before leaving the network. Upon examination, if the packet is a tagged packet, the IEEE 802.1p tag is configured based on the QoS level-to-IEEE 802.1p-bit mapping. For bridged packets, the DSCP is re-marked based on the QoS level.

Network congestion and QoS design

When you provide QoS in a network, one of the major elements you must consider is congestion, and the traffic management behavior during congestion. Congestion in a network is caused by many different conditions and events, including node failures, link outages, broadcast storms, and user traffic bursts.

At a high level, three main types or stages of congestion exist:

1. no congestion
2. bursty congestion
3. severe congestion

In a noncongested network, QoS actions ensure that delay-sensitive applications, such as real-time voice and video traffic, are sent before lower-priority traffic. The prioritization of delay-sensitive traffic is essential to minimize delay and reduce or eliminate jitter, which has a detrimental impact on these applications.

A network can experience momentary bursts of congestion for various reasons, such as network failures, rerouting, and broadcast storms. Virtual Services Platform 4000 has sufficient capacity to handle bursts of congestion in a seamless and transparent manner. If the burst is not sustained, the traffic management and buffering process on the switch allows all the traffic to pass without loss.

Severe congestion is defined as a condition where the network or certain elements of the network experience a prolonged period of sustained congestion. Under such congestion conditions, congestion thresholds are reached, buffers overflow, and a substantial amount of traffic is lost.

After the switch detects severe congestion, Virtual Services Platform 4000 discards traffic based on drop precedence values. This mode of operation ensures that high-priority traffic is not discarded before lower-priority traffic.

When you perform traffic engineering and link capacity analysis for a network, the standard design rule is to design the network links and trunks for a maximum average-peak utilization of no more than 80%. This value means that the network peaks to up to 100% capacity, but the average-peak utilization does not exceed 80%. The network is expected to handle momentary peaks above 100% capacity.

QoS examples and recommendations

The sections that follow present QoS network scenarios for bridged and routed traffic over the core network.

Bridged traffic

If you bridge traffic over the core network, you keep customer VLANs separate (similar to a Virtual Private Network). Normally, a service provider implements VLAN bridging (Layer 2) and no routing. In this case, the 802.1p-bit marking determines the QoS level assigned to each packet. If DiffServ is active on core ports, the level of service received is based on the highest of the DiffServ or 802.1p settings.

The following cases provide sample QoS design guidelines you can use to provide and maintain high service quality in a network.

If you configure a core port, you assume that, for all incoming traffic, the QoS value is properly marked. All core switch ports simply read and forward packets; they are not re-marked or reclassifiied. All initial QoS markings are performed at the customer device or on the edge devices.

The following figure illustrates the actions performed on three different bridged traffic flows (that is VoIP, video conference, and e-mail) at access and core ports throughout the network.

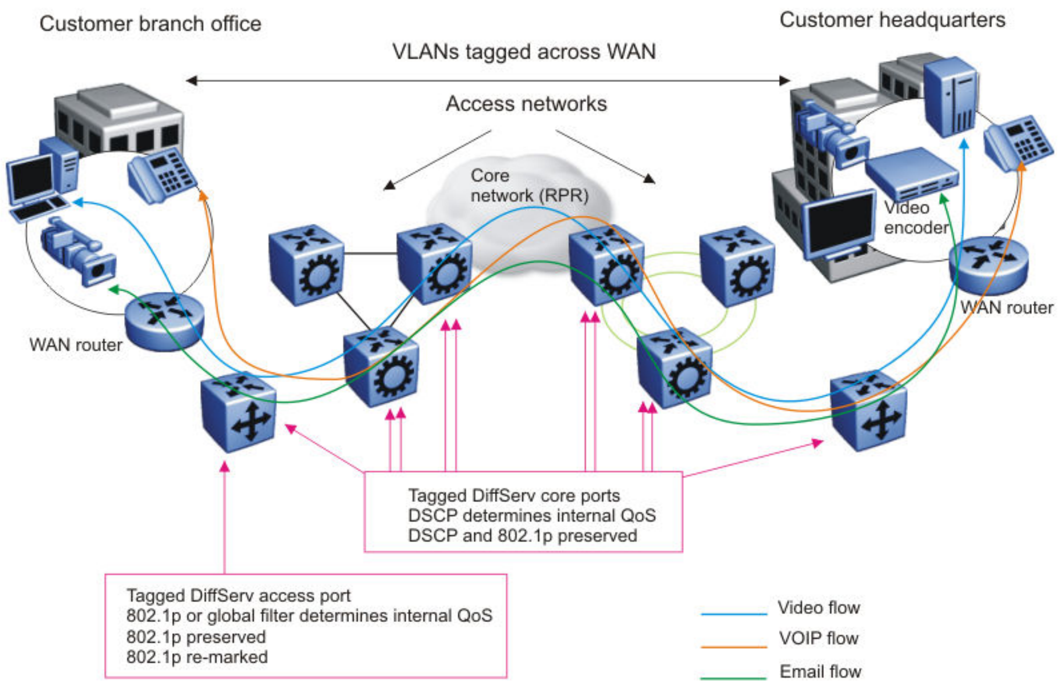

Figure 58: Trusted bridged traffic

For bridged, untrusted traffic, if you configure the port to access, mark and prioritize traffic on the access node using global filters. Reclassify the traffic to ensure it complies with the class of service specified in the SLA.

For Resilient Packet Ring (RPR) interworking, you can assume that, for all incoming traffic, the QoS configuration is properly marked by the access nodes. The core switch ports, configured as core or trunk ports, perform the RPR interworking. These ports preserve the DSCP marking and re-mark the 802.1p bit to match the 802.1p bit of the RPR. The following figure shows the actions performed on three different traffic flows (VoIP, video conference, and e-mail) over an RPR core network.

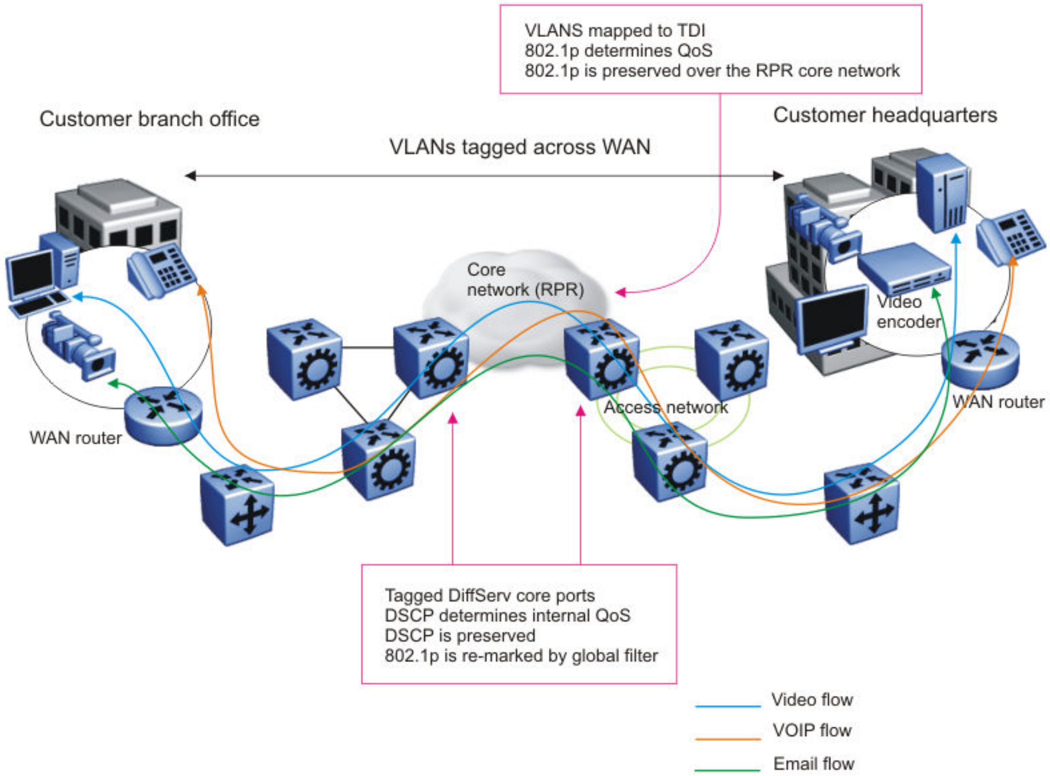

Figure 59: RPR QoS internetworking

Routed traffic

If you route traffic over the core network, VLANs are not kept separate.

If you configure the port to core, you assume that, for all incoming traffic, the QoS configuration is properly marked. All core switch ports simply read and forward packets. The switch does not re-mark or classify the packets. The customer device or the edge devices perform all initial QoS markings.

The following figure shows the actions performed on three different routed traffic flows (that is VoIP, video conference, and e-mail) at access and core ports throughout the network.

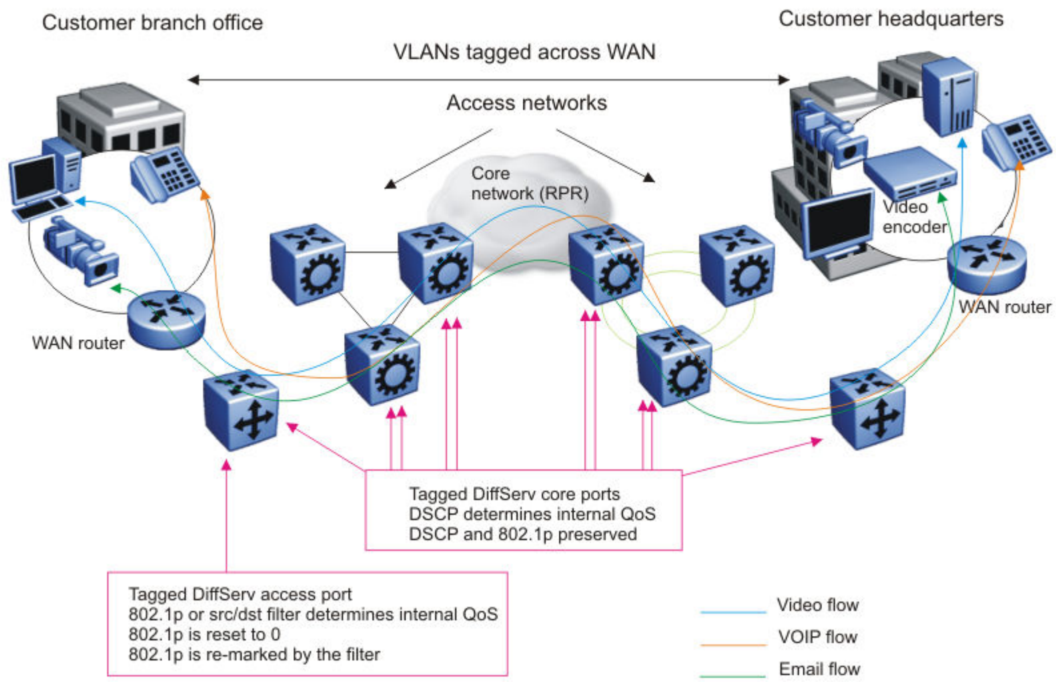

Figure 60: Trusted routed traffic

For routed, untrusted traffic, in an access node, packets that enter through a tagged or untagged access port exit through a tagged or untagged core port.

Chapter 15: Layer 1, 2, and 3 design examples

This section provides examples to help design your network. Layer 1 examples deal with the physical network layouts. Layer 2 examples map Virtual Local Area Networks (VLAN) on top of the physical layouts. Layer 3 examples show the routing instances that Avaya recommends to optimize IP for network redundancy.

Layer 1 example

This section describes a Layer 1 network design example that focuses primarily on the physical network layout. In this example, a VSP 4000 switch can function as an access switch.

Layer 1: Design example

This example uses double physical links and distributed MultiLink Trunking (DMLT) to provide a redundant network.

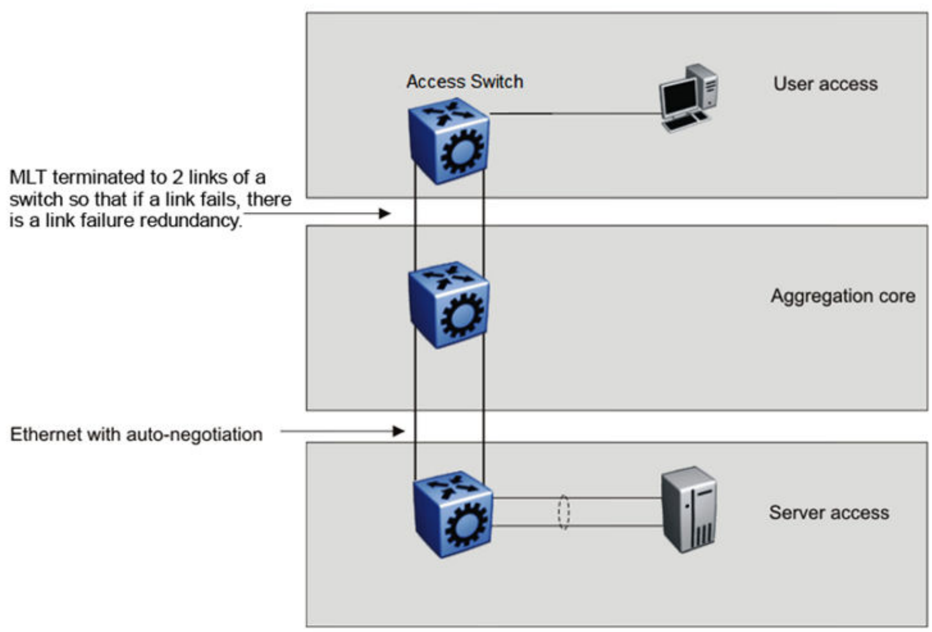

Figure 61: Layer 1 design example

Layer 2 example

This section describes a Layer 2 network design example that maps VLANs over the physical network layout.

Layer 2: Design example

The following example shows a redundant device network that uses one VLAN for all switches. To support multiple VLANs, you need 802.1Q tagging on the links with trunks.

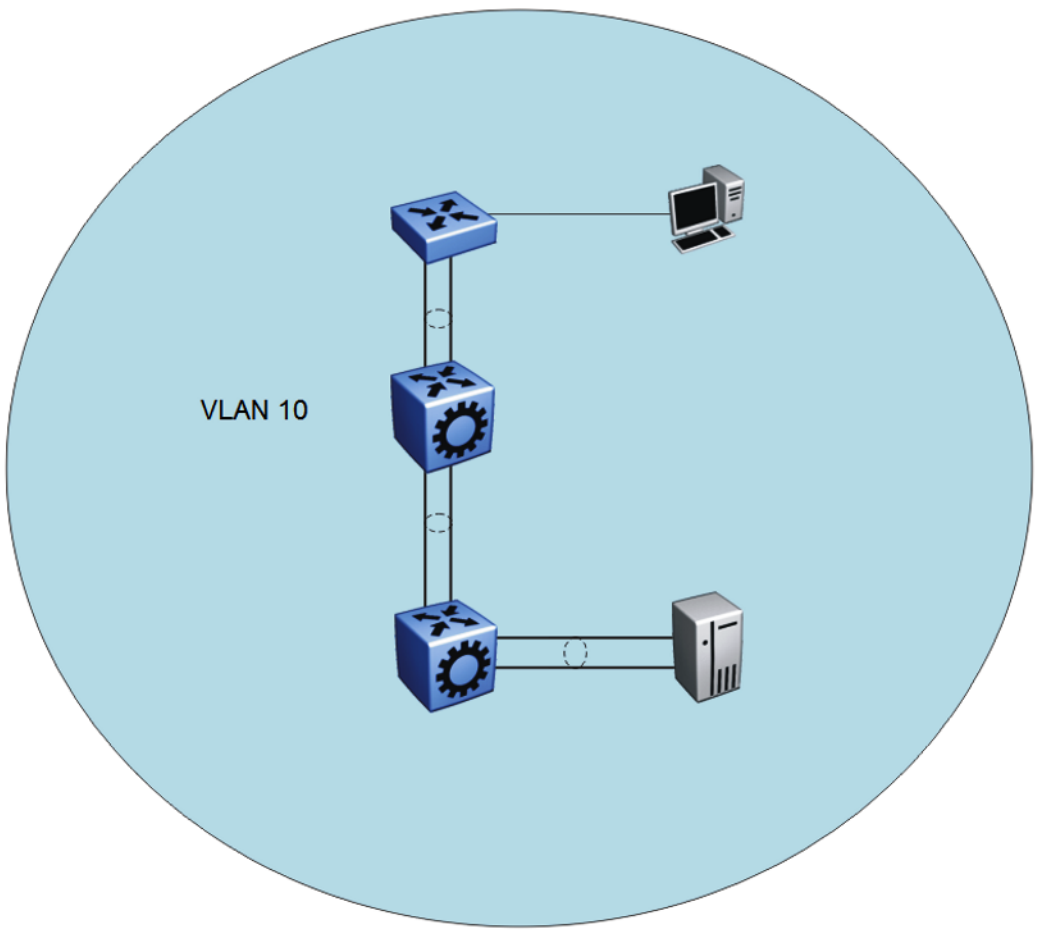

Figure 62: Layer 2 design example

Layer 3 example

This section describes a Layer 3 network design example that shows the routing instances that Avaya recommends you use to optimize IP for network redundancy.

Layer 3: Design example

Figure 63: Layer 3 design example on page 141 uses redundant links.

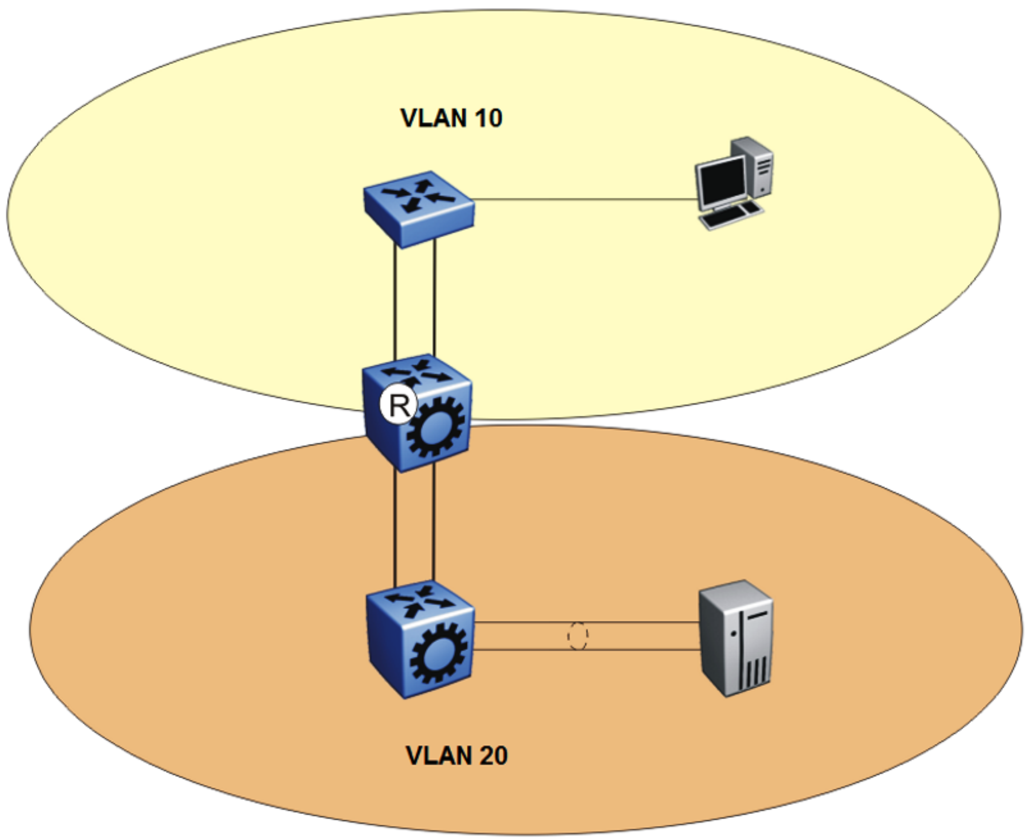

Figure 63: Layer 3 design example

Chapter 16: Software scaling capabilities

This section lists software scaling capabilities of Avaya Virtual Services Platform 4000.

Table 15: Software scaling capabilities

	Maximum number supported
Layer 2	
IEEE/Port-based VLANs	4060
LACP	24 aggregators
LACP ports per aggregator	8 active and 8 standby
MACs in forwarding database (FDB)	32,000
Multi-Link Trunking (MLT)	24 groups
Multiple Spanning Tree Protocol (MSTP)	12 instances
Protocol-based VLANs	1
Rapid Spanning Tree Protocol (RSTP)	1 instance
SLPP	128 VLANs
VLACP Interfaces	50
Layer 3	
RIP interfaces	24
RIP routes	500
OSPF interfaces	48 (24 of these can be passive)
OSPF adjacencies	24
OSPF areas (per system)	64
OSPF routes per VRF	16,000 **Note:** The maximum routes supported per VRF is 16000. The 16000 routes can be distributed across the 24 VRFs (+ GRT) in any manner. If all 24 VRFs are operational, 640 routes per VRF are supported.
OSPF routes	16,000

	Maximum number supported
OSPF VRF support	24
e-BGP peers	12
e-BGP routes	16,000
Address Resolution Protocol (ARP) for each port, VRF, or VLAN (IPv4)	6,000 entries total
Circuitless IP interfaces	64
Maximum B-MACs	1000
ECMP routes	1000
ECMP groups	512 groups with a maximum of 4 ECMP paths per group **Note:** The maximum number of ECMP routes per VSP 4000 system is 1000. So, for example, if 500 ECMP groups are configured, the maximum number of ECMP paths per group is 2 and if 250 ECMP groups are configured, the maximum number of ECMP paths per group is 4.
ECMP paths per route	4
FIB IPv4 routes	16,000
RIB IPv4 routes	16,000
IPv4 interfaces	256
Maximum VRFs	24
IPv4 CLIP interfaces	64
IP routing policies	500 for each VRF 5,000 for each system
IPv4 FTP sessions	4
IPv4 Rlogin sessions	8
IPv4 SSH sessions	8
IPv4 Telnet sessions	8
IPv4 VRF instances	24
Static ARP entries (IPv4)	200 for each VRF 1,000 for each system
Static routes (IPv4)	1,000 per VRF/per system
UDP/DHCP forwarding entries	128 for each system

	Maximum number supported
VRRP interfaces (IPv4)	64
VRRP interfaces fast timers (200 ms)	24
Diagnostics	
Mirrored ports	49
Remote Mirroring Termination (RMT) ports	4
Filters and QoS	
Port shapers (IPv4)	50
ACEs per ACL (a combination of Security and QoS ACEs)	1,000
Unique redirect next hop values for ACE Actions (IPv4)	Ingress: 1,536, Egress: 256
SPBM	
C-VLANs per VSP 4000 node	1000
Maximum number of nodes per region	1000
MAC entries	16,000 (combination of ARP entries and Layer 2 MACs)
Backbone MAC	1,000
IP routes in the Global Router	16,000
Maximum IS-IS IP routes	16,000
IS-IS interfaces	24
IS-IS adjacencies per VSP 4000 node	24
Layer 2 VSN ISIDs per VSP 4000 node	1,000
Layer 3 VSN ISIDs per VSP 4000 node	24
IP Multicast over SPB	
Maximum unique IGMP group records per node	1000
Maximum unique Multicast Streams (S,G,V) sourced per node	1000
Maximum number of Multicast ISIDs (VSP 4000 acting as a BEB and/or BCB)	32,000
Maximum number of Layer 2 VSNs with Multicast enabled	1000
Maximum number of Layer 3 VSNs with Multicast enabled	24

	Maximum number supported
Maximum number of IP interfaces with Multicast enabled	256
Number of remote senders that can be received on each VSP 4000 node, for the Universal Plug and Play Group (239.255.255.250)	3500
Maximum unique multicast streams sourced per VSP 4000 node	1000
T-UNI	
T-UNI ISIDs per VSP 4000 node	48
Maximum MAC limit on a T-Uni I-SID **Note:** This is also the device limit.	32,000

Chapter 17: Supported standards, RFCs, and MIBs

This chapter details the standards, request for comments (RFC), and Management Information Bases (MIB) that Avaya Virtual Services Platform 4000 supports.

Supported IEEE standards

The following table details the IEEE standards that Avaya Virtual Services Platform 4000 supports.

Table 16: Supported IEEE standards

IEEE standard	Description
802.1aq	Shortest Path Bridging (SPB)
802.1D	MAC bridges (Spanning Tree)
802.1AX	Link Aggregation Control Protocol (LACP)
802.1p	VLAN prioritization
802.1Q	Virtual Local Area Network (VLAN) tagging
802.1s	Multiple Spanning Tree Protocol
802.1t	802.1D maintenance
802.1w-2001	Rapid Spanning Tree protocol (RSTP)
802.1X	Extended Authentication Protocol (EAP), and EAP over LAN (EAPoL)
802.1X-2004	Port Based Network Access Control
802.3 CSMA/CD Ethernet ISO/IEC 8802	International Organization for Standardization (ISO) /International Eletrotechnical Commission (IEC) 8802-3
802.3ab	Gigabit Ethernet 1000BaseT 4 pair Category 5 (Cat5) Unshieled Twisted Pair (UTP)
802.3ae	10 Gigabit Ethernet
802.3af and 802.3at	PoE – Power Over Ethernet

IEEE standard	Description
802.3i	10BaseT
802.3u	100BaseT
802.3x	flow control
802.3z	Gigabit Ethernet

Supported RFCs

The following table and sections list the RFCs that Avaya Virtual Services Platform 4000 supports.

Table 17: Supported request for comments

Request for comment	Description
RFC768	UDP Protocol
RFC783	Trivial File Transfer Protocol (TFTP)
RFC791	Internet Protocol (IP)
RFC792	Internet Control Message Protocol (ICMP)
RFC793	Transmission Control Protocol (TCP)
RFC826	Address Resolution Protocol (ARP)
RFC854	Telnet protocol
RFC894	A standard for the Transmission of IP Datagrams over Ethernet Networks
RFC896	Congestion control in IP/TCP internetworks
RFC906	Bootstrap loading using TFTP
RFC950	Internet Standard Subnetting Procedure
RFC951	BootP
RFC959, RFC1350, and RFC2428	FTP and TFTP client and server
RFC1027	Using ARP to implement transparent subnet gateways/Nortel Subnet based VLAN
RFC1122	Requirements for Internet Hosts
RFC1256	ICMP Router Discovery
RFC1305	Network Time Protocol v3 Specification, Implementation and Analysis

Request for comment	Description
RFC1340	Assigned Numbers
RFC1519	Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy
RFC1541	Dynamic Host Configuration Protocol1
RFC1542	Clarifications and Extensions for the Bootstrap Protocol
RFC1591	DNS Client
RFC1812	Router requirements
RFC1866	HyperText Markup Language version 2 (HTMLv2) protocol
RFC2068	Hypertext Transfer Protocol
RFC2131	Dynamic Host Control Protocol (DHCP)
RFC2138	RADIUS Authentication
RFC2139	RADIUS Accounting
RFC2338	VRRP: Virtual Redundancy Router Protocol
RFC2616	Hypertext Transfer Protocol 1.1
RFC2819	RMON
RFC2992	Analysis of an Equal-Cost Multi-Path Algorithm
RFC3046	DHCP Option 82
RFC3621	PoE – Power Over Ethernet
RFC4250–RFC4256	SSH server and client support
RFC6329	IS-IS Extensions supporting Shortest Path Bridging

Quality of service

Table 18: Supported request for comments

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

Request for comment	Description
RFC2598	An Expedited Forwarding PHB

Network management

Table 19: Supported request for comments

Request for comment	Description
RFC1155	SMI
RFC1157	SNMP
RFC1215	Convention for defining traps for use with the SNMP
RFC1271	Remote Network Monitoring Management Information Base
RFC1305	Network Time Protocol v3 Specification, Implementation and Analysis3
RFC1350	The TFTP Protocol (Revision 2)
RFC1354	IP Forwarding Table MIB
RFC1757	Remote Network Monitoring Management Information Base
RFC1907	Management Information Base for Version 2 of the Simple Network Management Protocol (SNMPv2)
RFC1908	Coexistence between v1 & v2 of the Internet-standard Network Management Framework
RFC1930	Guidelines for creation, selection, and registration of an Autonomous System (AS)
RFC2541	Secure Shell Protocol Architecture
RFC2571	An Architecture for Describing SNMP Management Frameworks
RFC2572	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
RFC2573	SNMP Applications

Request for comment	Description
RFC2574	User-based Security Model (USM) for v3 of the Simple Network Management Protocol (SNMPv3)
RFC2575	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
RFC2576	Coexistence between v1, v2, & v3 of the Internet standard Network Management Framework
RFC2819	Remote Network Monitoring Management Information Base

MIBs

Table 20: Supported request for comments

Request for comment	Description
RFC1156	MIB for network management of TCP/IP
RFC1212	Concise MIB definitions
RFC1213	TCP/IP Management Information Base
RFC1354	IP Forwarding Table MIB
RFC1389	RIPv2 MIB Extensions
RFC1398	Ethernet MIB
RFC1442	Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)
RFC1450	Management Information Base for v2 of the Simple Network Management Protocol (SNMPv2)
RFC1573	Interface MIB
RFC1650	Definitions of Managed Objects for the Ethernet-like Interface Types
RFC1657	BGP-4 MIB using SMIv2
RFC1850	OSPF MIB
RFC2096	IP Forwarding Table MIB

Request for comment	Description
RFC2578	Structure of Management Information v2 (SMIv2)
RFC2674	Bridges with Traffic MIB
RFC2787	Definitions of Managed Objects for the Virtual Router Redundancy Protocol
RFC2863	Interface Group MIB
RFC2925	Remote Ping, Traceroute & Lookup Operations MIB
RFC3416	v2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)
RFC4022	Management Information Base for the Transmission Control Protocol (TCP)
RFC4113	Management Information Base for the User Datagram Protocol (UDP)

Standard MIBs

The following table details the standard MIBs that Avaya Virtual Services Platform 4000 supports.

Table 21: Supported MIBs

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

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STDMIB6—Simple Network Management Protocol (SNMP)	RFC1157	rfc1157.mib
STDMIB7—MIB for network management of Transfer Control Protocol/Internet Protocol (TCP/IP) based Internet MIB2	RFC1213	rfc1213.mib
STDMIB8—A convention for defining traps for use with SNMP	RFC1215	rfc1215.mib
STDMIB10—Definitions of Managed Objects for Bridges	RFC1493	rfc1493.mib
STDMIB11—Evolution of the Interface Groups for MIB2	RFC2863	rfc2863.mib
STDMIB12—Definitions of Managed Objects for the Ethernet-like Interface Types	RFC1643	rfc1643.mib
STDMIB15—Remote Network Monitoring (RMON)	RFC2819	rfc2819.mib
STDMIB17—Management Information Base of the Simple Network Management Protocol version 2 (SNMPv2)	RFC1907	rfc1907.mib
STDMIB21—Interfaces Group MIB using SMIv2	RFC2233	rfc2233.mib
STDMIB26a—An Architecture for Describing SNMP Management Frameworks	RFC2571	rfc2571.mib
STDMIB26b—Message Processing and Dispatching for the SNMP	RFC2572	rfc2572.mib
STDMIB26c—SNMP Applications	RFC2573	rfc2573.mib
STDMIB26d—User-based Security Model (USM) for version 3 of the SNMP	RFC2574	rfc2574.mib

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STDMIB26e—View-based Access Control Model (VACM) for the SNMP	RFC2575	rfc2575.mib
STDMIB26f —Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework	RFC2576	rfc2576.mib
STDMIB29—Definitions of Managed Objects for the Virtual Router Redundancy Protocol	RFC2787	rfc2787.mib
STDMIB31—Textual Conventions for Internet Network Addresses	RFC2851	rfc2851.mib
STDMIB32—The Interface Group MIB	RFC2863	rfc2863.mib
STDMIB33—Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations	RFC2925	rfc2925.mib
STDMIB38—SNMPv3 These Request For Comments (RFC) make some previously named RFCs obsolete	RFC3411, RFC3412, RFC3413, RFC3414, RFC3415	rfc2571.mib, rfc2572.mib, rfc2573.mib, rfc2574.mib, rfc2575.mib
STDMIB39—Entity Sensor Management Information Base	RFC3433	
STDMIB40—The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model	RFC3826	rfc3826.mib
STDMIB41—Management Information Base for the Transmission Control protocol (TCP)	RFC4022	rfc4022.mib

Standard MIB name	Institute of Electrical and Electronics Engineers/ Request for Comments (IEEE/RFC)	File name
STDMIB43—Management Information Base for the User Datagram Protocol (UDP)	RFC4113	rfc4113.mib
STDMIB44—Entity MIB	RFC4133	rfc4133.mib
STDMIB45 – Definitions of Managed Power Over Ethernet	RFC3621	rfc3621.mib

Proprietary MIBs

The following table details the proprietary MIBs that Avaya Virtual Services Platform 4000 supports.

Table 22: Proprietary MIBs

Proprietary MIB name	File name
PROMIB1 – Rapid City MIB	rapid_city.mib
PROMIB 2 – SynOptics Root MIB	synro.mib
PROMIB3 – Other SynOptics definitions	s5114roo.mib
PROMIB4 – Other SynOptics definitions	s5tcs112.mib
PROMIB5 – Other SynOptics definitions	s5emt103.mib
PROMIB6 – Avaya RSTP/MSTP proprietary MIBs	nnrst000.mib, nnmst000.mib
PROMIB11 – Avaya MIB definitions	wf_com.mib
PROMIB12 – Other SynOptic definition for Combo Ports	s5ifx.mib
PROMIB31 – Other SynOptic definition for PoE	bayStackPethExt.mib

Glossary

Backbone Core Bridge (BCB)	Backbone Core Bridges (BCBs) form the core of the SPBM network. The BCBs are SPBM nodes that do not terminate the VSN services. BCBs forward encapsulated traffic based on the Backbone MAC Destination Address (BMAC-DA). A BCB is unaware of the VSN traffic it transports. A BCB simply knows how to reach any other Backbone Edge Bridges (BEBs) in the SPBM backbone.
Backbone Edge Bridge (BEB)	Backbone Edge Bridges (BEBs) are SPBM nodes where Virtual Services Networks (VSNs) terminate. BEBs handle the boundary between the core MAC-in-MAC Shortest Bath Bridging MAC (SPBM) domain and the edge customer 802.1Q domain. A BEB node performs 802.1ah MAC-in-MAC encapsulation and decapsulation for the Virtual Services Network (VSN).
Backbone MAC (BMAC)	Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation encapsulates customer MAC addresses in Backbone MAC (BMAC) addresses. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that SPBM uses for delivery from end to end. As the MAC header stays the same across the network, no need exists to swap a label or do a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end. In Shortest Path Bridging MAC (SPBM), each node has a System ID, which is used in the topology announcement. This same System ID also serves as the switch Backbone MAC address (B-MAC), which is used as the source and destination MAC address in the SPBM network.
Backbone VLAN identifier (B-VID)	The Backbone VLAN identifier (B-VID) indicates the Shortest Path Bridging MAC (SPBM) B-VLAN associated with the SPBM instance.
bit error rate (BER)	The ratio of the number of bit errors to the total number of bits transmitted in a specific time interval.
Connectivity Fault Management (CFM)	Connectivity Fault Management is a mechanism to debug connectivity issues and to isolate faults within the Shortest Path Bridging-MAC (SPBM) network. CFM operates at Layer 2 and provides the equivalent of ping and traceroute. IEEE 802.1ag Connectivity Fault Management (CFM) divides or separates a network into administrative domains called Maintenance Domains (MD).
coarse wavelength division	A technology that uses multiple optical signals with different wavelengths to simultaneously transmit in the same direction over one fiber, and then separates by wavelength at the distant end.

multiplexing (CWDM)	
Customer MAC (C-MAC)	For customer MAC (C-MAC) addresses, which is customer traffic, to forward across the service provider back, SPBM uses IEEE 802.1ah Provider Backbone Bridging MAC-in-MAC encapsulation. The system encapsulates C-MAC addresses within a backbone MAC (B-MAC) address pair made up of a BMAC destination address (BMAC-DA) and a BMAC source address (BMAC-SA).
dense wavelength division multiplexing (DWDM)	A technology that uses many optical signals (16 or more) with different wavelengths to simultaneously transmit in the same direction across one fiber, and then separates by wavelength at the distant end.
Designated Intermediate System (DIS)	A Designated Intermediate System (DIS) is the designated router in Intermediate System to Intermediate System (IS-IS) terminology. You can modify the priority to affect the likelihood of a router being elected the designated router. The higher the priority, the more likely the router is to be elected as the DIS. If two routers have the same priority, the router with the highest MAC address (Sequence Number Packet [SNP] address) is elected as the DIS.
Global routing engine (GRE)	The base router or routing instance 0 in the Virtual Routing and Forwarding (VRF).
Intermediate System to Intermediate System (IS-IS)	Intermediate System to Intermediate System(IS-IS) is a link-state, interior gateway protocol. ISO terminology refers to routers as Intermediate Systems (IS), hence the name Intermediate System-to-Intermediate System (IS-IS). IS-IS operation is similar to Open Shortest Path First (OSPF). In Shortest Path Bridging-MAC (SPBM) networks, IS-IS discovers network topology and builds shortest path trees between network nodes that IS-IS uses for forwarding unicast traffic and determining the forwarding table for multicast traffic. SPBM employs IS-IS as the interior gateway protocol and implements additional Type-Length-Values (TLVs) to support additional functionality.
Internet Protocol security (IPsec)	A secure version of the Internet Protocol (IP) that provides optional authentication and encryption at the packet level.
jitter	The delay variance between received packets. Packets may not arrive at the destination address in consecutive order, or on a timely basis, and the signal can vary from its original reference timing. This distortion damages multimedia traffic.

last member query interval (LMQI)	The time between when the last IGMP member leaves the group and the stream stops.
latency	The time between when a node sends a message and receipt of the message by another node; also referred to as propagation delay.
Layer 1	Layer 1 is the Physical Layer of the Open System Interconnection (OSI) model. Layer 1 interfaces with the MAC sublayer of Layer 2, and performs character encoding, transmission, reception, and character decoding.
Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are: Ethernet and Frame Relay.
Layer 2 Virtual Services Network	Layer 2 Virtual Services Network (L2 VSN) feature provides IP connectivity over SPBM for VLANs. Backbone Edge Bridges (BEBs) handle Layer 2 virtualization. At the BEBs you map the end-user VLAN to a Service Instance Identifier (I-SID). BEBs that have the same I-SID configured can participate in the same Layer 2 Virtual Services Network (VSN).
Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
Layer 3 Virtual Services Network	The Layer 3 Virtual Services Network (L3 VSN) feature provides IP connectivity over SPBM for VRFs. Backbone Edge Bridges (BEBs) handle Layer 3 virtualized. At the BEBs through local provisioning, you map the end-user IP enabled VLAN or VLANs to a Virtualized Routing and Forwarding (VRF) instance. Then you map the VRF to a Service Instance Identifier (I-SID). VRFs that have the same I-SID configured can participate in the same Layer 3 Virtual Service Network (VSN).
Layer 4	The Transport Layer of the OSI model. An example of a Layer 4 protocol is Transmission Control Protocol (TCP).
link-state advertisement (LSA)	Packets that contain state information about directly connected links (interfaces) and adjacencies. Each Open Shortest Path First (OSPF) router generates the packets.
link-state database (LSDB)	A database built by each OSPF router to store LSA information. The router uses the LSDB to calculate the shortest path to each destination in the Autonomous System (AS), with itself at the root of each path.
Link Aggregation Control Protocol (LACP)	A protocol that exists between two endpoints to bundle links into an aggregated link group for bandwidth increase and link redundancy.
Link Aggregation Control Protocol	Link aggregation control protocol data unit (LACPDU) is used for exchanging information among LACP-enabled devices.

Data Units (LACPDU)	
link aggregation group (LAG)	A group that increases the link speed beyond the limits of one single cable or port, and increases the redundancy for higher availability.
load balancing	The practice of splitting communication into two (or more) routes or servers.
MAC-in-MAC encapsulation	MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses. The originating node creates a MAC header that the device uses for delivery from end to end. As the MAC header stays the same across the network, there is no need to swap a label or do a route lookup at each node, allowing the frame to follow the most efficient forwarding path end to end.
management information base (MIB)	The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
multicast router discovery (MRDISC)	Provides the automatic discovery of multicast capable routers. By listening to multicast router discovery messages, Layer 2 devices can determine where to send multicast source data and Internet Group Management Protocol (IGMP) host membership reports.
multihomed AS	An AS that has multiple connections to one or more ASs and does not carry transit traffic.
multiplexing	Carriage of multiple channels over a single transmission medium; a process where a dedicated circuit is shared by multiple users. Typically, data streams intersperse on a bit or byte basis (time division), or separate by different carrier frequencies (frequency division).
next hop	The next hop to which a packet can be sent to advance the packet to the destination.
not so stubby area (NSSA)	Prevents the flooding of external link-state advertisements (LSA) into the area by providing them with a default route. An NSSA is a configuration of the Open Shortest Path First (OSPF) protocol.
out of band (OOB)	Network dedicated for management access to chassis.
packet loss	Expressed as a percentage of packets dropped over a specified interval. Keep packet loss to a minimum to deliver effective IP telephony and IP video services.
policing	Ensures that a traffic stream follows the domain service provisioning policy or service level agreement (SLA).

Protocol Data Units (PDUs)	A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
Provider Backbone Bridge (PBB)	To forward customer traffic across the service provider backbone, SPBM uses IEEE 802.1ah Provider Backbone Bridging (PBB) MAC-in-MAC encapsulation, which hides the customer MAC (C-MAC) addresses in a backbone MAC (B-MAC) address pair. MAC-in-MAC encapsulation defines a BMAC-DA and BMAC-SA to identify the backbone source and destination addresses.
quality of service (QoS)	QoS features reserve resources in a congested network, allowing you to configure a higher priority to certain devices. For example, you can configure a higher priority to IP deskphones, which need a fixed bit rate, and, split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.
Read Write All (RWA)	An access class that lets users access all menu items and editable fields.
remote login (rlogin)	An application that provides a terminal interface between hosts (usually UNIX) that use the TCP/IP network protocol. Unlike Telnet, rlogin assumes the remote host is, or behaves like, a UNIX host.
remote monitoring (RMON)	A remote monitoring standard for Simple Network Management Protocol (SNMP)-based management information bases (MIB). The Internetwork Engineering Task Force (IETF) proposed the RMON standard to provide guidelines for remote monitoring of individual LAN segments.
resilient packet ring (RPR)	A shared packet edge ring connection, where both paths around the ring carry traffic, that allows double bandwidth on each ring.
reverse path forwarding (RPF)	Prevents a packet from forging its source IP address. Typically, the system examines and validates the source address of each packet.
route flapping	An instability that is associated with a prefix, where the associated prefix routes can exhibit frequent changes in availability over a period of time.
routing policy	A form of routing that is influenced by factors other than the default algorithmically best route, such as the shortest or quickest path.
Secure Copy (SCP)	Secure Copy securely transfers files between the switch and a remote station.
Secure Shell (SSH)	SSH uses encryption to provide security for remote logons and data transfer over the Internet.

Secure Sockets Layer (SSL)
An Internet security encryption and authentication protocol for secure point-to-point connections over the Internet and intranets, especially between clients and servers.

Service Instance Identifier (I-SID)
The SPBM B-MAC header includes a Service Instance Identifier (I-SID) with a length of 24 bits. SPBM uses this I-SID to identify and transmit any virtualized traffic in an encapsulated SPBM frame. SPBM uses I-SIDs to virtualize VLANs (Layer 2 Virtual Services Network [VSN]) or VRFs (Layer 3 Virtual Services Network [VSN]) across the MAC-in-MAC backbone. With Layer 2 VSNs, you associate the I-SID with a customer VLAN, which is then virtualized across the backbone. With Layer 3 VSNs, you associate the I-SID with a customer VRF, which is also virtualized across the backbone.

service level agreement (SLA)
A service contract that specifies the forwarding service that traffic receives.

Shortest Path Bridging (SPB)
Shortest Path Bridging is a control Link State Protocol that provides a loop free Ethernet topology. There are two versions of Shortest Path Bridge: Shortest Path Bridging VLAN and Shortest Path Bridging MAC. Shortest Path Bridging VLAN uses the Q-in-Q frame format and encapsulates the source bridge ID into the VLAN header. Shortest Path Bridging MAC uses the 802.1 ah MAC-in-MAC frame format and encapsulates the source bridge identifier into the B-MAC header.

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

shortest path tree (SPT)
Creates a direct route between the receiver and the source for group members in a Protocol Independent Multicast—Spare Mode (PIM-SM) domain.

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

single mode fiber (SMF)
One of the various light waves transmitted in an optical fiber. Each optical signal generates many modes, but in single-mode fiber only one mode is transmitted. Transmission occurs through a small diameter core (approximately ten micrometers), with a cladding that is 10 times the core

	diameter. These fibers have a potential bandwidth of 50 to 100 GHz per kilometer.
small form factor pluggable (SFP)	A hot-swappable input and output enhancement component used with Avaya products to allow gigabit Ethernet ports to link with other gigabit Ethernet ports over various media types.
small form factor pluggable plus (SFP+)	SFP+ transceivers are similar to SFPs in physical appearance but SFP + transceivers provide Ethernet at 10 gigabit per second (Gb/s).
spanning tree	A simple, fully-connected active topology formed from the arbitrary physical topology of connected bridged Local Area Network components by relaying frames through selected bridge ports. The protocol parameters and states that are used and exchanged to facilitate the calculation of the active topology and to control the bridge relay function.
Spanning Tree Group (STG)	A collection of ports in one spanning tree instance.
SysAdmin, Audit, Network, Security (SANS) Institute	The research and education organization for network administrators and security professionals.
time-to-live (TTL)	The field in a packet used to determine the valid duration for the packet. The TTL determines the packet lifetime. The system discards a packet with a TTL of zero.
traffic engineering	A method that guarantees performance in a network.
Trivial File Transfer Protocol (TFTP)	A protocol that governs transferring files between nodes without protection against packet loss.
trunk	A logical group of ports that behaves like a single large port.
unshielded twisted pair (UTP)	A cable with one or more pairs of twisted insulated copper conductors bound in a single plastic sheath.
user-based security model (USM)	A security model that uses a defined set of user identities for authorized users on a particular Simple Network Management Protocol (SNMP) engine.
User Datagram Protocol (UDP)	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.

view-based access control model (VACM)	Provides context, group access, and group security levels based on a predefined subset of management information base (MIB) objects.
Virtual Link Aggregation Control Protocol (VLACP)	Virtual Link Aggregation Control Protocol (VLACP) is a Layer 2 handshaking protocol that can detect end-to-end failure between two physical Ethernet interfaces.
Virtual Router Redundancy Protocol (VRRP)	A protocol used in static routing configurations, typically at the edge of the network. This protocol operates on multiple routers on an IP subnet and elects a primary gateway router. When the primary router fails, a backup router is quickly available to take its place.
Voice over IP (VOIP)	The technology that delivers voice information in digital form in discrete packets using the Internet Protocol (IP) rather than the traditional circuit-committed protocols of the public switched telephone network (PSTN).
wavelength division multiplexing (WDM)	Simultaneously transmits many colors (wavelengths) of laser light down the same optical fiber to increase the amount of transferred information.
wiring closet	A central termination area for telephone or network cabling or both.