



Configuring Ethernet Modules on Avaya Virtual Services Platform 9000

Release 4.1
NN46250-508
Issue 07.01
October 2015

© 2010-2015, Avaya Inc.
All Rights Reserved.

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

Avaya grants You a license within the scope of the license types described below, with the exception of Heritage Nortel Software, for which the scope of the license is detailed below. Where the order documentation does not expressly identify a license type, the applicable license will be a Designated System License. The

applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the documentation or other materials available to You. "Software" means computer programs in object code, provided by Avaya or an Avaya Channel Partner, whether as stand-alone products, pre-installed on hardware products, and any upgrades, updates, patches, bug fixes, or modified versions thereto. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users. "Instance" means a single copy of the Software executing at a particular time: (i) on one physical machine; or (ii) on one deployed software virtual machine ("VM") or similar deployment.

Licence types

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

Heritage Nortel Software

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

Copyright

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

Third Party Components

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

Preventing Toll Fraud

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

result in substantial additional charges for your telecommunications services.

Avaya Toll Fraud intervention

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

Downloading Documentation

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

Contact Avaya Support

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

Trademarks

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

Avaya is a registered trademark of Avaya Inc.

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

Contents

Chapter 1: Introduction	6
Purpose.....	6
Related resources.....	6
Documentation.....	6
Training.....	6
Viewing Avaya Mentor videos.....	6
Support.....	7
Searching a documentation collection.....	7
Locating end of sale notices.....	8
Chapter 2: New in this release	9
Features.....	9
Other changes.....	10
Chapter 3: Ethernet module fundamentals	11
Ethernet module concepts.....	11
Port speed and duplex mode.....	11
Autonegotiation.....	11
Custom Autonegotiation Advertisement.....	13
802.3x flow control for Gigabit Ethernet.....	14
Maximum transmission unit and jumbo frames.....	14
MLT and Link Aggregation Control Protocol groups port speed.....	15
Virtual Router Forwarding for Ethernet modules.....	15
Ethernet module features.....	15
First and second generation modules and modes.....	16
9048XS-2 I/O module.....	18
9024XL I/O module.....	21
9048GT interface module.....	24
9048GB interface module.....	24
9012QQ-2 I/O module.....	26
Ethernet SFPs and Autonegotiation.....	28
Chapter 4: Ethernet module configuration using CLI	30
Enabling a module.....	30
Disabling a module.....	31
Configuring the minimum module generation support.....	31
Configuring 10/100/1000 Mbps Ethernet ports with Autonegotiation.....	33
Configuring 10/100/1000 Mbps Ethernet ports without Autonegotiation.....	35
Configuring 40 Gbps Ethernet ports with Autonegotiation.....	37
Variable definitions.....	37
Configuring 100/1000 Mbps Ethernet ports.....	38
Configuring 10 Gbps and 40 Gbps Ethernet ports.....	41

Associating a port to a VRF instance.....	42
Chapter 5: Ethernet module configuration using EDM.....	44
Enabling a module.....	44
Disabling a module.....	45
Configuring the minimum module generation support.....	45
Configuring an Ethernet port.....	46
Associating a port to a VRF instance.....	51
Chapter 6: Lossless Ethernet	52
Lossless Ethernet fundamentals.....	52
Lossless Ethernet configuration using ACLI.....	55
Configuring Lossless Ethernet.....	55
Displaying Lossless Ethernet information.....	58
Lossless Ethernet configuration using EDM.....	59
Configuring Lossless Ethernet.....	59
Specifying the lossless-802.1p value.....	60
Chapter 7: Port reference.....	64
Glossary.....	67

Chapter 1: Introduction

Purpose

You can use Configuration – Ethernet Modules to configure and monitor the Ethernet modules for the Virtual Services Platform 9000.

Related resources

Documentation

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

Training

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

Course code	Course title
4D00010E	Knowledge Access: ACIS - Avaya ERS 8000 and VSP 9000 Implementation
5D00040E	Knowledge Access: ACSS - Avaya VSP 9000 Support

Viewing Avaya Mentor videos

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

About this task

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

Procedure

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

 **Note:**

Videos are not available for all products.

Support

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

Searching a documentation collection

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

Before you begin

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

Procedure

1. Extract the document collection zip file into a folder.

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

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

Locating end of sale notices

Use the following procedure to locate the most up-to-date information on end of sale notices and replacement recommendations for transceivers and optical components.

Procedure

1. Go to the Avaya Support website at <http://support.avaya.com/>.
2. Select **Support by Product > Documents**.
3. Enter the product name, and a release.
4. Select Product Lifecycle Notices.
5. Click **Enter**.
6. Select the End of Sale notice to view the information.

Chapter 2: New in this release

There are no feature-related changes in *Configuring Ethernet Modules on Avaya Virtual Services Platform 9000*, NN46250-508 for Release 4.1.

Features

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

Release 4.1 adds the following SFP+ and QSFP+ modules and cables.

Bidirectional SFP+ optical transceiver

Release 4.1 adds support for the 10GBASE-BX SFP+ bidirectional transceivers.

Model	Part number
10GBASE-BX SFP+ bidirectional transceivers, 10 kilometer	AA1403169, AA1403170

For more information, see:

- [9024XL I/O module](#) on page 21.

End of sale notices

Release 4.1 adds information on how to locate the most up-to-date information end of sale notices and replacement recommendations for transceivers and optical components. For more information, see [Locating end of sale notices](#) on page 8.

QSFP+ optical transceiver

Release 4.1 adds the following two QSFP+ optical transceivers.

Model	Part number
40GBase-LM4 QSFP+	AA1404002-E6

For more information, see:

- [9012QQ-2 I/O module](#) on page 26.

QSFP+ Direct Attach Cable (DAC)

The QSFP+ to QSFP+ 40 gigabit Direct Attach Cable (DAC) assembly directly connects two QSFP+ ports. The new 10 meter cable is added in this release.

Model	Part number
QSFP+ to QSFP+ 40G, 10 meter active DAC	AA1404028-E6

For more information, see:

- [9012QQ-2 I/O module](#) on page 26.

QSFP+ breakout cable specifications

The QSFP+ to four SFP+ 10 Gigabit Ethernet (GbE) breakout cable (BOC) assembly directly connects one QSFP+ port to four SFP+ ports. The new 1, 3, 5, 10 meter cables are added in this release.

*** Note:**

Avaya Virtual Services Platform 9000 does not support the 40 Gigabit Ethernet ends of the QSFP+ breakout cables because the platform does not support channelization on the VSP 9012QQ-2 module. Avaya Virtual Services Platform 9000 supports only the four SFP+ 10 Gigabit Ethernet ends of the following QSFP+ breakout cables: AA1404033-E6, AA1404035-E6, AA1404036-E6, and AA1404041-E6.

Model	Part number
QSFP+ to 4 SFP+ breakout cable, 1 meter (Passive)	AA1404033-E6
QSFP+ to 4 SFP+ breakout cable, 3 meter (Passive)	AA1404035-E6
QSFP+ to 4 SFP+ breakout cable, 5 meter (Passive)	AA1404036-E6
Fiber QSFP+ to 4 SFP+ breakout cable, 10 meter (Active)	AA1404041-E6

For more information, see:

- [9012QQ-2 I/O module](#) on page 26.

Other changes

There are no other changes.

Chapter 3: Ethernet module fundamentals

Use the information in this section to understand Ethernet configuration concepts.

Ethernet module concepts

Extending Ethernet over local, metropolitan, and wide area networks provides cost-effective solutions for data transmission and ensures end-to-end Ethernet connectivity. To optimize Ethernet communications, you can configure a number of parameters. The following sections provide information you can use to configure Ethernet modules for optimal performance.

Port speed and duplex mode

Avaya Virtual Services Platform 9000 Ethernet modules support various data rates; these include 10 Mbps, 100 Mbps, 1000 Mbps, 10 Gbps, and 40 Gbps . Some modules support several data rates and use Autonegotiation to determine the appropriate data rate.

The duplex mode determines the direction of traffic flow during data communication. Half-duplex indicates that communication occurs between two Ethernet devices in both directions, but only in one direction at a time. Full-duplex indicates that communication occurs in both directions simultaneously. If you do not use Autonegotiation, ensure that you set both interconnected Ethernet devices to the same mode.

Some modules must use full-duplex; on those modules, you cannot change this parameter.

Autonegotiation

Use Autonegotiation to allow the device to automatically negotiate the best common data rate and duplex mode to use between two Autonegotiation-capable Ethernet devices.

Autonegotiation automatically configures devices that share a link segment to take maximum advantage of their abilities. You can use Autonegotiation to advertise abilities of devices at both ends of a link segment, acknowledge receipt and understanding of their common modes of operation, and to reject the use of operational modes that both devices do not share. Where more than one common mode exists between the two devices, a mechanism selects a single mode of operation using a predetermined priority resolution function.

Autonegotiation uses autosensing to allow devices compatible with 10BASE-T, 100BASE-TX, or 100BASE-T4 to be recognized even if they do not support Autonegotiation. In this case, only the speed can be sensed, not the duplex mode.

Autonegotiation uses a modified 10BASE-T link integrity test pulse sequence to perform Autonegotiation; no packet or upper layer protocol overhead is added to the network devices.

Disabling Autonegotiation on 1000BASE-T ports can result in inconsistent behavior. Avaya recommends that you enable Autonegotiation on all 1000BASE-T ports.

Autonegotiation is not supported on 10 Gbps small form factor pluggable plus (SFP+) ports or on 100 Mbps small form factor pluggable (SFP) ports. While 1000 Mbps SFP ports do support Autonegotiation, they do not support a change in data rate unless you change the SFP which supports a different speed; a 1 Gbps SFP can only operate at 1 Gbps. On fixed speed links, Autonegotiation does not negotiate the data rate; Autonegotiation is only used for Remote Fault Indication (RFI).

*** Note:**

The 10 GigabitEthernet fiber-based I/O module ports can operate at either 1 Gigabit per second (Gbps) or 10 Gbps, dependent upon the capabilities optical transceiver that you install.

This presents an ambiguity with respect to the autonegotiation settings of the port, while 1 Gigabit Ethernet (GbE) ports require autonegotiation; autonegotiation is not defined and is non-existent for 10 GbE ports.

For a 10GbE fiber-based I/O module, you have the capability to swap back-and-forth between 1 GbE and 10 GbE operation by simply swapping transceivers. To help with this transition between 1 GbE and 10 GbE port operation, Avaya allows you to configure autonegotiation when you install a 10 GbE transceiver, even though autonegotiation is not defined for 10GbE.

You can do this in anticipation of a port changeover from 10 GbE to 1 GbE. In this manner, you could essentially preconfigure a port in 1 GbE mode while the 10 GbE transceiver is still installed. The port is ready to go upon the changeover to the 1 GbE transceiver.

In addition, you can use a saved configuration file with autonegotiation enabled to boot a system with either 10 GbE or 1 GbE transceivers installed. If you install a 1 GbE transceiver, the system applies autonegotiation. If you install a 10 GbE transceiver, the system does not remove the autonegotiation settings from the configuration, but the system simply ignores the configuration because autonegotiation settings are irrelevant to a 10 GbE transceiver. The system preserves the saved configuration for autonegotiation when resaved no matter which speed of transceiver you install.

Autonegotiation recommendations

Avaya recommends that you set the following Autonegotiation parameters between two 1000BASE-T ports on devices A and B.

Table 1: Recommended Autonegotiation configuration on 1000BASE-T ports

Port on device A	Port on device B	Remarks	Recommendation
Autonegotiation enabled	Autonegotiation enabled	Ports negotiate on highest supported mode on both sides.	This configuration is the recommended configuration if both ports support Autonegotiation mode.
Autonegotiation disabled; set to full-duplex	Autonegotiation disabled; set to full-duplex	Both sides require the same mode.	This configuration is the recommended configuration if you use full-duplex, but Autonegotiation is not supported.
Autonegotiation disabled; set to half-duplex	Autonegotiation enabled	This is an invalid configuration.	Use fixed speed/duplex modes on both sides.

*** Note:**

Although all Virtual Services Platform 9000 1000BASE-T ports support Autonegotiation, there can be situations where Autonegotiation does not function properly and a link remains down. If this situation occurs, Avaya recommends that you disable Autonegotiation and configure matching speed and duplex on both sides of the link (usually 10 Mbps and full-duplex). This operation does not comply with the latest IEEE 802.3ab standard, which recommends use of Autonegotiation on 1000BASE-T ports for 1000 Mbps links.

Remote Fault Indication

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

The 802.3z Gigabit Ethernet (GbE) standard defines RFI as part of the Autonegotiation function.

Because RFI is part of the Autonegotiation function, if you disable Autonegotiation, you automatically disable RFI.

+ Tip:

Avaya recommends that you enable Autonegotiation on GbE links in all cases where the devices on both ends of a fiber link support Autonegotiation.

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 and the link remains active.

Custom Autonegotiation Advertisement

The 10/100/1000 Mbps ports support Custom Autonegotiation Advertisement (CANA). Use CANA to control the speed and duplex settings that these modules advertise during Autonegotiation sessions between Ethernet devices. You 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 Mbps to 1000 Mbps on host and server connections. Using Autonegotiation only, the switch always uses the fastest possible data rates. In scenarios where uplink bandwidth is limited, CANA provides control over negotiated access speeds, and thus improves control over traffic load patterns.

Virtual Services Platform 9000 supports CANA on only 10/100/1000 Mbps RJ-45 ports. To use CANA, you must enable Autonegotiation.

802.3x flow control for Gigabit Ethernet

Use flow control to protect against packet loss. Flow control prevents Ethernet ports from receiving more data than they can process.

The 802.3x flow control mechanism is effective only at the Data Link Layer (Layer 2). When Virtual Services Platform 9000 routes packets, flow control packets are ignored.

The Ethernet modules support IEEE 802.3x flow control in the following way:

- temporarily halt packet transmission when an 802.3x Pause frame is received from the remote end

*** Note:**

802.3x flow control pause frames are not transmitted by default. Transmit flow control can be enabled if required.

Flow control can be enabled on ports of any speed.

Maximum transmission unit and jumbo frames

Jumbo frames are larger than the maximum Ethernet frame size, or maximum transmission unit (MTU) specified in the IEEE 802.3 standard. For untagged frames, the maximum standard size is 1518 bytes. For tagged frames, the maximum standard size increases by 4 bytes to 1522 bytes.

The default maximum Ethernet frame size supported by Virtual Services Platform 9000 is 1950 bytes. The 10 Gigabits per second (Gbps) module also supports this frame size. The 40 Gigabits per second (Gbps) module also supports all frame sizes including jumbo. To support a jumbo frame system, configure the CP MTU to 9600 bytes. Interface ports cannot pass jumbo frames by default.

*** Note:**

All Virtual Services Platform 9000 modules supports jumbo frames.

If a port on a Gigabit Ethernet module receives a packet bigger than the configured system MTU size, the MAC chip increments the byte counter with the MTU size instead of the actual packet size, which reflects in the statistics of that port.

*** Note:**

Because Virtual Services Platform 9000 does not negotiate the maximum transmission unit (MTU) for IPv4 interfaces, the interface MTU is the maximum-sized packet that the CP transmits. Virtual Services Platform 9000 receives and processes any packet less than the system MTU. In the fastpath, Virtual Services Platform 9000 receives and sends packets less than, or equal to, the system MTU.

For more information about the system MTU, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

MLT and Link Aggregation Control Protocol groups port speed

Avaya recommends that all ports that belong to the same MLT group or Link Aggregation Control Protocol (LACP) group use the same port speed, for example, 1 Gbps, even if you use Autonegotiation. In the case of MLTs, the software does not enforce this requirement.

Virtual Router Forwarding for Ethernet modules

You can use Virtual Router Forwarding (VRF) to have Virtual Services Platform 9000 act as multiple virtual routers.

You can configure each VRF instance as a separate router; this means that you can configure different routing protocols and associated parameters for each instance. You can associate non—zero VRF instances with module ports.

The Ethernet parameters (such as: Autonegotiate; AdminDuplex; AdminSpeed) that a user can edit for a VRF instance depend on whether the port belongs to only one, or more than one, VRF instance.

The following examples show what is possible:

- If a port belongs to only one VRF, you can edit the Ethernet parameters of the VRF.
- If a port belongs to more than one VRF instance, you cannot edit the Ethernet parameters of the VRF instance; you can only edit the GlobalRouter port parameters.
- If a port belongs to a single non—zero VRF, this VRF can change the port Ethernet parameters.
- If a port belongs to multiple VRF instances, then only the Global Router can change this port configuration.

Ethernet module features

This section details the features of Virtual Services Platform 9000 Ethernet modules.

+ Tip:

To determine if SFPs are present on your switch, use the CLI command `show pluggable-optical-modules <basic|detail>` to display SFP and SFP+ information.

For the `show pluggable-optical-modules basic` command, the device reports qualified optics as Avaya in the type field. The device reports non-qualified best-effort optics as a different manufacturer in the type field. Unsupported optics display as unsupported in the type field, and do not operate in the system.

The following table shows supported Ethernet modules you can use to provide Ethernet services using Virtual Services Platform 9000.

Table 2: Ethernet modules

Module	Module description
9048XS-2	48 port 10 gigabits per second (Gbps) small form-factor (SFP+) Ethernet interface module. The 9048XS-2 module also supports 1 Gbps SFP connections. See <i>Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000</i> , NN46250-305, for a list of supported SFPs.
9024XL	24 port 10 Gbps SFP+ Ethernet interface module. The 9024XL module also supports 1 Gbps SFP connections. See <i>Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000</i> , NN46250-305, for a list of supported SFPs.
9048GT	48 port 10/100/1000 megabits per second (Mbps) Ethernet Copper interface module with RJ45 connectors.
9048GB	48 port 1 Gbps SFP Ethernet interface module that supports multimode fiber (MMF), single-mode fiber (SMF), and copper connections. The 9048GB module also supports 100 MB FX connections. See <i>Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000</i> , NN46250-305, for a list of supported SFPs.
9012QQ-2	The 9012QQ-2 module is a 12-port 40 gigabits per second (Gbps) module, that supports the 40GBASE-R QSFP+ module configuration. See <i>Installing Transceivers and Optical Components on Avaya Virtual Services Platform 9000</i> , NN46250-305, for a list of supported SFPs.

First and second generation modules and modes

First generation and second generation modules operate in any I/O module slot in both Virtual Services Platform 9010 and Virtual Services Platform 9012.

*** Note:**

Avaya recommends you update your device fully to Release 4.0.1.0 or higher, and ensure that the upgrade is fully complete, before you install new 9048XS-2 or 9012QQ-2 I/O modules. Once the upgrade is fully complete, insert the new 9048XS-2 or 9012QQ-2 I/O module into the chassis, one module at a time.

The 9048XS-2 or 9012QQ-2 go through a series of steps as part of the upgrade process, including burning of images into the FPGAs on the module and can go through multiple module resets to activate those firmware images. Up to 35 minutes may be required for the upgrades on each module to be complete. Allow the upgrade process to complete successfully. Failure to do so could result in a failed or an incorrect upgrade or incorrect commissioning of your device. Refer to *Release Notes for Avaya Virtual Services Platform 9000*, NN46250-401 for more information.

The following table identifies the first generation and second generation module types and shows the various support configurations for each mode.

Chassis mode	Modules	Supported on both VSP 9010 and VSP 9012	High speed cooling module support
First generation mode	First generation modules: <ul style="list-style-type: none"> • 9024XL • 9048GB • 9048GT 	<ul style="list-style-type: none"> • Supports all first generation modules and second generation modules. • Switch Fabric (SF) modules: <ul style="list-style-type: none"> - First generation modules: require a minimum of three SF modules. - Second generation modules: require a minimum of five SF modules. 	<ul style="list-style-type: none"> • First generation modules on VSP 9010: use the 9012FC cooling module. The 9012FCHS is not required on the VSP 9010. • First generation modules on VSP 9012: use the 9012FC cooling modules or the 9012FCHS cooling modules. <p>* Note: Second generation modules running in first generation mode on the VSP 9012 still require the 9012FCHS cooling module.</p>
	Second generation modules: <ul style="list-style-type: none"> • 9048XS-2 • 9012QQ-2 		
Second generation mode	Second generation modules: <ul style="list-style-type: none"> • 9048XS-2 • 9012QQ-2 	<ul style="list-style-type: none"> • Supports only second generation modules. • Requires a minimum of five SF modules. • Achieves enhanced scaling capabilities for FIB routes. 	<ul style="list-style-type: none"> • Second generation modules on VSP 9010: use the 9012FC cooling module. The 9012FCHS is not required on the VSP 9010. • Second generation modules on VSP 9012: require the 9012FCHS cooling modules.

*** Note:**

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

Virtual Services Platform 9012, the system checks which cooling modules are in use, and if the system does not use the 9012FCHS cooling modules, the second generation modules remain offline.

First generation mode configuration

Virtual Services Platform 9000 with first generation modules are based on 3 x 48Gbps Network Processing Unit (NPU). If you want to use both first generation and second generation modules, you must use the default parameter `gen1` with the `boot config linecard-compatibility-mode {gen1|gen2}` command, which configures the system to first generation (gen1). After you configure the device for first generation modules, the system reboots in first generation (gen1) mode.

First generation mode is the default mode of operation.

Second generation mode configuration

Virtual Services Platform 9000 with second generation modules is a fully-featured high performance high-end platform based on 3 x 160Gbps NPUs per I/O module. If you want to use only second generation modules and achieve full scaling capability, you must use the parameter `gen2` with the `boot config linecard-compatibility-mode {gen1|gen2}` command, which configures the system to second generation (gen2). The system reboots in second generation mode and disables all first generation modules.

For more information, see *Configuring Ethernet Modules on Avaya Virtual Services Platform 9000*, NN46250-508.

9048XS-2 I/O module

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

This module supports standard management information base (MIB).

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

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

! Important:

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

Model	Part number	Description
1000BASE-T SFP	AA1419043-E6	gigabit Ethernet, RJ-45 connector 850 nm, gigabit Ethernet, duplex LC connector
1000BASE-SX DDI SFP	AA1419048-E6	850 nm, gigabit Ethernet, duplex LC connector
1000BASE-XD DDI SFP	AA1419050-E6*	1310 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-XD DDI SFP	AA1419051-E6*	1550 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-ZX DDI SFP	AA1419052-E6*	1550 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-BX10 DDI SFP	AA1419069-E6 (10 km at 1310 nm) and mating pair AA1419070-E6 (10 km at 1490 nm)	1310 nm (tx) and 1490 nm (rx) gigabit Ethernet, single-fiber LC connector
1000BASE-BX40 DDI SFP	AA1419076-E6 (40 km at 1310 nm) and mating pair AA1419077-E6 (40 km at 1490 nm)	1490 nm (tx) 1310 nm (rx) gigabit Ethernet, single-fiber LC connector
1000BASE-EX DDI SFP	AA1419071-E6** This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.	1550 nm, gigabit Ethernet, duplex LC connector

Table continues...

Model	Part number	Description
1000BASE–XD DDI CWDM 40 km SFP	AA1419053-E6 to AA1419060-E6** These transceivers have reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.	gigabit Ethernet, duplex LC connector
1000BASE–ZX DDI CWDM 70 km SFP	AA1419061-E6 to AA1419068-E6** These transceivers have reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8. ! Important: AA1419065–E6 remains available for purchase.	gigabit Ethernet, duplex LC connector
100BASE-FX SFP	AA1419074-E6	1300 nm, 100-Mbps Ethernet, multimode fiber, duplex LC connector

Table 3: Supported SFP+ transceivers

Model	Part number	Description
10GBASE-LR/LW SFP+	AA1403011-E6	10 km, 1310 nm SMF
10GBASE-ER/EW SFP+	AA1403013-E6	40 km, 1550 nm SMF
10GBASE-SR/SW SFP+	AA1403015-E6	400 m, 850 nm MMF
10GBASE-ZR/ZW SFP+	AA1403016-E6	70 km, 1550 nm SMF
10GBASE-LRM SFP+	AA1403017-E6	220 m, 1260 to 1355 nm; 1310 nm nominal MMF
10GBASE-CX SFP+ 2-pair twinaxial copper cable that plugs into the SFP+ socket	AA1403018-E6	10 meter

Table continues...

Model	Part number	Description
and connects two 10-gigabit ports		
10GBASE-ER CWDM DDI SFP+	AA1403153-E6 to AA1403160-E6	40 km, 1471 to 1611 nm
10GBASE-ZR CWDM DDI SFP+	AA1403161-E6 to AA1403168-E6	70 km, 1471 to 1611 nm
10GBASE-BX10 SFP+	AA1403169-E6 and AA1403170-E6	10 km

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

The 9048XS-2 has the following characteristics:

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

9024XL I/O module

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

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

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

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

Important:

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

Model	Part number	Description
1000BASE-T SFP	AA1419043-E6	gigabit Ethernet, RJ-45 connector

Table continues...

Model	Part number	Description
		850 nm, gigabit Ethernet, duplex LC connector
1000BASE-SX DDI SFP	AA1419048-E6	850 nm, gigabit Ethernet, duplex LC connector
1000BASE-XD DDI SFP	AA1419050-E6*	1310 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-XD DDI SFP	AA1419051-E6*	1550 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-ZX DDI SFP	AA1419052-E6*	1550 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-BX10 DDI SFP	AA1419069-E6 (10 km at 1310 nm) and mating pair AA1419070-E6 (10 km at 1490 nm)	1310 nm (tx) and 1490 nm (rx) gigabit Ethernet, single-fiber LC connector
1000BASE-BX40 DDI SFP	AA1419076-E6 (40 km at 1310 nm) and mating pair AA1419077-E6 (40 km at 1490 nm)	1490 nm (tx) 1310 nm (rx) gigabit Ethernet, single-fiber LC connector
1000BASE-EX DDI SFP	AA1419071-E6** This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.	1550 nm, gigabit Ethernet, duplex LC connector
1000BASE-XD DDI CWDM 40 km SFP	AA1419053-E6 to AA1419060-E6** These transceivers have reached end-of-sale (EOS). For more information about	gigabit Ethernet, duplex LC connector

Table continues...

Model	Part number	Description
	EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.	
1000BASE-ZX DDI CWDM 70 km SFP	AA1419061-E6 to AA1419068-E6** These transceivers have reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8. ! Important: AA1419065-E6 remains available for purchase.	gigabit Ethernet, duplex LC connector
100BASE-FX SFP	AA1419074-E6	1300 nm, 100-Mbps Ethernet, multimode fiber, duplex LC connector

Table 4: Supported SFP+ transceivers

Model	Part number	Description
10GBASE-LR/LW SFP+	AA1403011-E6	10 km, 1310 nm SMF
10GBASE-ER/EW SFP+	AA1403013-E6	40 km, 1550 nm SMF
10GBASE-SR/SW SFP+	AA1403015-E6	400 m, 850 nm MMF
10GBASE-ZR/ZW SFP+	AA1403016-E6	70 km, 1550 nm SMF
10GBASE-LRM SFP+	AA1403017-E6	220 m, 1260 to 1355 nm; 1310 nm nominal MMF
10GBASE-CX SFP+ 2-pair twinaxial copper cable that plugs into the SFP+ socket and connects two 10-gigabit ports	AA1403018-E6	10 meter
10GBASE-ER CWDM DDI SFP+	AA1403153-E6 to AA1403160-E6	40 km, 1471 to 1611 nm
10GBASE-ZR CWDM DDI SFP+	AA1403161-E6 to AA1403168-E6	70 km, 1471 to 1611 nm

Table continues...

Model	Part number	Description
10GBASE-BX10 SFP+	AA1403169-E6 and AA1403170-E6	10 km

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

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

The 9024XL has the following characteristics:

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

9048GT interface module

The 9048GT interface module is a 48 port 10/100/1000 megabits per second (Mbps) Ethernet Copper interface module with RJ45 connectors.

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

The 9048GT has the following characteristics:

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

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

9048GB interface module

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

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

Important:

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

Model	Part number	Description
1000BASE-T SFP	AA1419043-E6	gigabit Ethernet, RJ-45 connector 850 nm, gigabit Ethernet, duplex LC connector
1000BASE-SX DDI SFP	AA1419048-E6	850 nm, gigabit Ethernet, duplex LC connector
1000BASE-XD DDI SFP	AA1419050-E6*	1310 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-XD DDI SFP	AA1419051-E6*	1550 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-ZX DDI SFP	AA1419052-E6*	1550 nm, gigabit Ethernet, duplex LC connector * This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.
1000BASE-BX10 DDI SFP	AA1419069-E6 (10 km at 1310 nm) and mating pair AA1419070-E6 (10 km at 1490 nm)	1310 nm (tx) and 1490 nm (rx) gigabit Ethernet, single-fiber LC connector
1000BASE-BX40 DDI SFP	AA1419076-E6 (40 km at 1310 nm) and mating pair AA1419077-E6 (40 km at 1490 nm)	1490 nm (tx) 1310 nm (rx) gigabit Ethernet, single-fiber LC connector
1000BASE-EX DDI SFP	AA1419071-E6** This transceiver has reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.	1550 nm, gigabit Ethernet, duplex LC connector
1000BASE-XD DDI CWDM 40 km SFP	AA1419053-E6 to AA1419060-E6** These transceivers have reached end-of-	gigabit Ethernet, duplex LC connector

Table continues...

Model	Part number	Description
	sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8.	
1000BASE-ZX DDI CWDM 70 km SFP	AA1419061-E6 to AA1419068-E6** These transceivers have reached end-of-sale (EOS). For more information about EOS transceivers, and recommended replacements for your product, see Locating end of sale notices on page 8. ! Important: AA1419065-E6 remains available for purchase.	gigabit Ethernet, duplex LC connector
100BASE-FX SFP	AA1419074-E6	1300 nm, 100-Mbps Ethernet, multimode fiber, duplex LC connector

The 9048GB is 100/1000 Mbps capable.

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

The 9048GB has the following characteristics:

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

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

9012QQ-2 I/O module

The second generation 9012QQ-2 Input/Output (I/O) module is a 12-port 40 Gigabits per second (Gbps) module. The 9012QQ-2 module supports the 40GBASE-R QSFP+ transceivers. You must also have a minimum of five Switch Fabric modules installed, if you install the 9012QQ-2 module on

the Virtual Services Platform 9010 or Virtual Services Platform 9012. The Virtual Services Platform 9012 requires the High-Speed Front Cooling Modules be installed before you install the 9012QQ-2 module.

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

*** Note:**

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

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

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

! Important:

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

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

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

*** Note:**

Avaya Virtual Services Platform 9000 does not support the 40 Gigabit Ethernet ends of the QSFP+ breakout cables because the platform does not support channelization on the VSP 9012QQ-2 module. Avaya Virtual Services Platform 9000 supports only the four SFP+ 10 Gigabit Ethernet ends of the following QSFP+ breakout cables: AA1404033-E6, AA1404035-E6, AA1404036-E6, and AA1404041-E6.

Table 5: Compatible 40 gigabit QSFP+ transceivers

Hardware	Description	Minimum software version	Part number
Active optical DAC	10 meter	4.1	AA1404028-E6
Passive copper DAC	1 meter	4.0.1	AA1404029-E6
Passive copper DAC	3 meter	4.0.1	AA1404031-E6
Passive copper DAC	5 meter	4.0.1	AA1404032-E6
Passive copper DAC	0.5 meter	4.0.1	AA1404037-E6GS
Passive copper DAC	1 meter	4.0.1	AA1404038-E6GS
Passive copper DAC	3 meter	4.0.1	AA1404039-E6GS
Passive copper breakout cable	1 meter	4.1	AA1404033-E6

Table continues...

Hardware	Description	Minimum software version	Part number
Passive copper breakout cable	3 meter	4.1	AA1404035-E6
Passive copper breakout cable	5 meter	4.1	AA1404036-E6
Active optical breakout cable	10 meter	4.1	AA1404041-E6
40GBASE-LR4 QSFP+	10 km 1300 nm with duplex LC connector	4.0.1	AA1404001-E6
40GBASE-LM4 QSFP+	80 meters on 50 µm multimode fiber	4.1	AA1404002-E6
40GBASE-SR4	150 m, (OM4) 850 nm with MPO/MTP connector	4.0.1	AA1404005-E6

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

The 9012QQ-2 has the following characteristics:

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

Ethernet SFPs and Autonegotiation

Two interface modules accept SFPs: the 9024XL (10 Gbps) and the 9048GB (1 Gbps). Both accept a limited number of lower-speed SFPs, which affects Autonegotiation support.

The following restrictions apply to the 9024XL:

- You can use 1 Gbps SFPs but not 100 Mbps SFPs.
- Autonegotiation is undefined for 10 Gbps interfaces. You cannot use the CLI to modify Autonegotiation on a port with a 10 Gbps SFP inserted.
- You can use CLI to change the Autonegotiation configuration for an empty port before you insert a 1 Gbps SFP.
- After you insert a 1 Gbps SFP into a port, you can enable Autonegotiation on that port only for the detection of remote faults. You cannot select the speed or duplex mode, or configure CANA.
- The default configuration for Autonegotiation is disabled, even if you insert a 1 Gbps SFP.

The following restrictions apply to the 9048GB:

- Autonegotiation is enabled by default for all ports.
- Autonegotiation is disabled if you insert a 100 Mbps SFP. This module does not support Autonegotiation on reduced speed SFPs.

- You can change the Autonegotiation configuration for a port, as long as you do not insert a 100 Mbps SFP.

Because you can change SFPs, inconsistencies can arise in the configuration file if you swap SFPs of different speeds. The port operates correctly after you change the SFP, but the configuration file can contain inappropriate Autonegotiation settings because of prior CLI actions. The system ignores these inconsistencies and does not generate an error. You can save the running configuration to update the file with correct Autonegotiation settings for the port.

Chapter 4: Ethernet module configuration using ACLI

Configure an Ethernet port to ensure proper operation and optimum performance of the port.

Enabling a module

About this task

Enable an Ethernet module to allow traffic to flow through it.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable a module:

```
no slot shutdown {slot[-slot][,...]}
```

Example

```
VSP-9012:>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#no slot shutdown 3/1-3/48, 4/1-4/24
```

Variable definitions

Use the data in the following table to use the `no slot shutdown` command.

Variable	Value
{slot[-slot][,...]}	Specifies one or more slot numbers. Valid slots are 3 to 12.

Disabling a module

About this task

Disable an Ethernet module before you remove it from the chassis to minimize traffic loss. Traffic does not flow on a disabled module.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Disable a module:

```
slot shutdown {slot[-slot][, ...]}
```

Example

```
VSP-9012:>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#slot shutdown 3/1
```

Variable definitions

Use the data in the following table to use the `slot shutdown` command.

Variable	Value
{slot[-slot][, ...]}	Specifies one or more slot numbers. Valid slot numbers are 3 to 12.

Configuring the minimum module generation support

Configure the minimum generation of modules that the Virtual Services Platform 9000 supports. Both Virtual Services Platform 9010 and Virtual Services Platform 9012 support first and second generation modules.

The default is gen1.

About this task

In gen1 mode, VSP 9000 supports both first and second generation modules.

In gen2 mode, VSP 9000 supports only second generation modules. After reboot in second generation-mode, the system disables all first generation modules, and provides full fastpath routing for up to one million IPv4 FIB routes.

*** Note:**

When you change the `linecard-compatibility-mode` the system saves the configuration and reboots.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the minimum module generation support:

```
boot config linecard-compatibility-mode {gen1|gen2}
```

3. Configure the minimum module generation support to the default:

```
default boot config linecard-compatibility-mode
```

4. Display the configuration:

```
show boot config master
```

Example

Configure the minimum module generation support to the default of gen1:

```
VSP-9010:>enable
VSP-9010:#configure terminal
VSP-9010:(config)#default boot config linecard-compatibility-mode
VSP-9010:1(config)#show boot config master
delay 10
master 1
logfile 100 1024 90
loadconfigtime 60
multicast 0
linecard-compatibility-mode gen1
```

Variable definitions

Use the data in the following table to use the `boot config linecard-compatibility-mode` command.

Variable	Value
{gen1 gen2}	<p>Specifies the minimum module generation support for VSP 9000.</p> <ul style="list-style-type: none"> • gen1 specifies VSP 9000 supports both first and second generation modules. When you specify gen1 the device configures the maximum IPv4 FIB route entries to 400,000. • gen2 specifies VSP 9000 supports only second generation modules. In second generation mode, the device configures the maximum IPv4 FIB route entries to one million. After reboot in second

Variable	Value
	<p>generation-mode, the system disables all first generation modules, and provides full fastpath routing for up to one million IPv4 FIB routes.</p> <p>The default is gen1.</p>

Configuring 10/100/1000 Mbps Ethernet ports with Autonegotiation

Configure 10/100/1000 Mbps Ethernet ports with Autonegotiation so that they operate optimally for your network conditions. To use Custom Autonegotiation Advertisement (CANA), you must enable Autonegotiation.

About this task

Important:

- Avaya recommends that all ports that belong to the same MLT or Link Aggregation Control Protocol (LACP) group use the same port speed. In the case of MLTs, the software does not enforce this.
- 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), you must apply the same configuration to all other ports of the MLT group (if they support CANA).
- If a 10/100/1000 Mbps port that supports CANA is in a MLT group with port type that does not support CANA, use CANA only if it does not conflict with MLT abilities.
- If you disable CANA on a port using the **auto-negotiation-advertisements none** or **no auto-negotiation-advertisements** commands, the link state goes down until a valid CANA setting is applied to the port. Resetting the device is not required for the setting to take effect.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port][,...]}
```

2. Enable Autonegotiation:

```
auto-negotiate [port {slot/port[-slot/port][,...]}] enable
```

3. Configure CANA:

```
auto-negotiation-advertisements [port {slot/port[-slot/port][,...]}]
<10-full|10-half|100-full|100-half|1000-full|1000-half|none>
```

Example

```
VSP-9012:>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 4/2
VSP-9012:1(config-if)#auto-negotiate enable
VSP-9012:1(config-if)#auto-negotiation-advertisements 100-full
```

Variable definitions

Use the data in following table to use the `auto-negotiate` command.

Variable	Value
port {slot/port[-slot/port][,...]}	Specifies the port or ports that you want to configure.
enable]	<p>Enables Autonegotiation for the port or other ports of the module.</p> <p>The default form of this command is <code>default auto-negotiate [port {slot/port[-slot/port][,...]] [enable]</code>.</p> <p>The no form of this command is <code>no auto-negotiate [port {slot/port[-slot/port][,...]] [enable]</code>.</p> <p>* Note:</p> <p>The 10 GigabitEthernet fiber-based I/O module ports can operate at either 1 Gigabit per second (Gbps) or 10 Gbps, dependent upon the capabilities optical transceiver that you install.</p> <p>This presents an ambiguity with respect to the autonegotiation settings of the port, while 1 Gigabit Ethernet (GbE) ports require autonegotiation; autonegotiation is not defined and is non-existent for 10 GbE ports.</p> <p>For a 10GbE fiber-based I/O module, you have the capability to swap back-and-forth between 1 GbE and 10 GbE operation by simply swapping transceivers. To help with this transition between 1 GbE and 10 GbE port operation, Avaya allows you to configure autonegotiation when you install a 10 GbE transceiver, even though autonegotiation is not defined for 10GbE.</p> <p>You can do this in anticipation of a port changeover from 10 GbE to 1 GbE. In this manner, you could essentially preconfigure a port in 1 GbE mode while the 10 GbE</p>

Table continues...

Variable	Value
	<p>transceiver is still installed. The port is ready to go upon the changeover to the 1 GbE transceiver.</p> <p>In addition, you can use a saved configuration file with autonegotiation enabled to boot a system with either 10 GbE or 1 GbE transceivers installed. If you install a 1 GbE transceiver, the system applies autonegotiation. If you install a 10 GbE transceiver, the system does not remove the autonegotiation settings from the configuration, but the system simply ignores the configuration because autonegotiation settings are irrelevant to a 10 GbE transceiver. The system preserves the saved configuration for autonegotiation when resaved no matter which speed of transceiver you install.</p>

Use the data in following table to use the `auto-negotiation-advertisements` command.

Variable	Value
port {slot/port[-slot/port][,...]}	<p>Specifies which port to configure the CANA advertisement settings of. You must enable Autonegotiation to use CANA.</p> <p>The default form of this command is <code>default auto-negotiation-advertisements [port {slot/port[-slot/port][,...]}]</code> .</p> <p>The no form of this command is <code>no auto-negotiation-advertisements [port {slot/port[-slot/port][,...]}]</code>.</p>
<10-full 10-half 100-full 100-half 1000-full 1000-half none>	Configures the CANA speed and duplex mode.

Configuring 10/100/1000 Mbps Ethernet ports without Autonegotiation

Configure 10/100/1000 Mbps Ethernet ports without Autonegotiation so that they operate optimally for your network conditions.

About this task

! Important:

Ensure that all ports that belong to the same MLT or Link Aggregation Control Protocol (LACP) group use the same port speed.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port][,...]}
```

2. Configure the duplex mode:

```
duplex [port {slot/port[-slot/port][,...]] <half|full>
```

3. Configure the port speed:

```
speed [port {slot/port[-slot/port][,...]] <10|100>
```

Example

Configure port 4/47 to use full-duplex mode and a speed of 100 Mbps.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 4/47
VSP-9012:1(config-if)#duplex full
VSP-9012:1(config-if)#speed 100
```

Variable definitions

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

Variable	Value
port {slot/port[-slot/port][,...]}	Specifies which port to configure the duplex settings of.
<half full>	Specifies half-duplex or full-duplex mode for 10/100/1000 Mbps ports; 1 and 10 Gbps optical ports must use full-duplex mode.

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

Variable	Value
port {slot/port[-slot/port][,...]}	Specifies the port speed. This option does not apply to 1 Gigabit or 10 Gigabit Ethernet modules. Use <code>port {slot/port[-slot/port][,...]}</code> to change other ports to the same settings.

Table continues...

Variable	Value
	The default form of this command is <code>default speed [port {slot/port[-slot/port] [,...]}</code> .
<10 100>	Specifies the speed of the port in Mbps.

Configuring 40 Gbps Ethernet ports with Autonegotiation

Configure 40 Gbps Ethernet ports with Autonegotiation so that they operate optimally for your network conditions.

* Note:

You cannot configure Custom Autonegotiation Advertisement (CANA) for 10 Gbps and 40 Gbps ports.

Autonegotiation for 40Gbps is supported only for CR4 ports.

! Important:

Avaya recommends that all ports that belong to the same MultiLink Trunking (MLT) or Link Aggregation Control Protocol (LACP) group use the same port speed. In the case of MLTs, the software does not enforce this.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port] [,...]}
```

2. Enable Autonegotiation:

```
auto-negotiate [port {slot/port[-slot/port] [,...]}] enable
```

Example

```
Switch>enable
Switch#configure terminal
Switch(config)#interface gigabitethernet 4/2
Switch(config-if)#auto-negotiate enable
```

Variable definitions

Use the data in following table to use the `auto-negotiate` command.

Variable	Value
<code>port {slot/port[-slot/port] [,...]}</code>	Specifies the port or ports that you want to configure.

Table continues...

Variable	Value
enable]	<p>Enables Autonegotiation for the port or other ports of the module.</p> <p>The default form of this command is <code>default auto-negotiate [port {slot/port[-slot/port][, ...]] [enable]</code>.</p> <p>The no form of this command is <code>no auto-negotiate [port {slot/port[-slot/port][, ...]] [enable]</code>.</p>

Configuring 100/1000 Mbps Ethernet ports

Configure 100/1000 Mbps Ethernet ports so they operate optimally for your network conditions. These ports use the small form-factor pluggable (SFP) transceivers.

About this task

Important:

Ensure that all ports that belong to the same MLT or LACP group use the same port speed.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port][, ...]}
```

2. Enable Autonegotiation:

```
auto-negotiate [port {slot/port[-slot/port][, ...]] enable
```

3. Enable flow control as required:

```
tx-flow-control [port {slot/port[-slot/port][, ...]] enable
```

4. **(Optional)** Create a name for the port or group of ports:

```
name [port {slot/port[-slot/port][, ...]] WORD<0-42>
```

5. **(Optional)** Disable the interface:

```
shutdown [port {slot/port[-slot/port][, ...]]
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 3/48
VSP-9012:1(config-if)#auto-negotiate enable
VSP-9012:1(config-if)#tx-flow-control enable
```

Variable definitions

Use the data in the following table to use the `auto-negotiate` command.

Variable	Value
port {slot/port[-slot/port][,...]}	Specifies the port or ports to configure.
enable	<p>Enables Auto-Negotiation for the selected ports.</p> <p>The default form of this command is <code>default auto-negotiate [port {slot/port[-slot/port][,...]}] [enable]</code>.</p> <p>To disable Auto-Negotiation, use the <code>no</code> form of this command: <code>no auto-negotiate [port {slot/port[-slot/port][,...]}] [enable]</code>.</p> <p>* Note:</p> <p>The 10 Gigabit Ethernet I/O module ports can operate at either 1 Gigabit per second (Gbps) or 10 Gbps, dependent upon the capabilities optical transceiver that you install.</p> <p>This presents an ambiguity with respect to the autonegotiation settings of the port, while 1 Gigabit Ethernet (GbE) ports require autonegotiation; autonegotiation is not defined and is non-existent for 10 GbE ports.</p> <p>For a 10 GbE I/O module, you have the capability to swap back-and-forth between 1 GbE and 10 GbE operation by simply swapping transceivers. To help with this transition between 1 GbE and 10 GbE port operation, Avaya allows you to configure autonegotiation when you install a 10 GbE transceiver, even though autonegotiation is not defined for 10GbE.</p> <p>You can do this in anticipation of a port changeover from 10 GbE to 1 GbE. In this manner, you could essentially preconfigure a port in 1 GbE mode while the 10 GbE transceiver is still installed. The port is ready to go upon the changeover to the 1 GbE transceiver.</p> <p>In addition, you can use a saved configuration file with autonegotiation enabled to boot a system with either 10 GbE or 1 GbE transceivers installed. If you install a 1 GbE</p>

Table continues...

Variable	Value
	transceiver, the system applies autonegotiation. If you install a 10 GbE transceiver, the system does not remove the autonegotiation settings from the configuration, but the system simply ignores the configuration because autonegotiation settings are irrelevant to a 10 GbE transceiver. The system preserves the saved configuration for autonegotiation when resaved no matter which speed of transceiver you install.

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

Variable	Value
port {slot/port[-slot/port][,...]}	Specifies the port or ports to configure.
<WORD 0-42>	Configures a name for the port or ports.

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

Variable	Value
port {slot[-slot][,...]}	Specifies the port or ports to configure. The no form of this command is <code>no shutdown [port {slot[-slot][,...]}]</code> .

Use the data in the following table to use the **tx-flow-control** command.

Variable	Value
port {slot/port[-slot/port][,...]}	Specifies the port or ports to configure.
enable	Enables flow control. After you enable tx-flow-control, the TX port transmits MAC control PAUSE frames to indicate congestion on the receive side of the port interface. The link partner ceases transmission based on the MAC control frame pause-timer value. The default is disabled. You can enable flow control only on 1 Gbps and 10 Gbps ports. You cannot enable flow control for ports that run at less than 1 Gbps. The default form of this command is <code>default tx-flow-control [port {slot/port[-slot/port][,...]}] [enable]</code> . The no form of this command is <code>no tx-flow-control [port {slot/port[-slot/port][,...]}] [enable]</code> .

Configuring 10 Gbps and 40 Gbps Ethernet ports

Configure 10 Gbps and 40 Gbps Ethernet ports so they operate optimally for your network conditions.

Important:

- You cannot configure Autonegotiation for 10 Gbps ports.
- Ensure that all ports that belong to the same MLT or LACP group use the same port speed.

About this task

These ports use the small form-factor pluggable (SFP) and small form-factor pluggable plus (SFP+) transceivers. For more information about configuring SFP transceivers, see [Configuring 100/1000 Mbps Ethernet ports](#) on page 38.

SFP transceivers are also supported on the 9024XL.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port][, ...]}
```

2. Enable flow control:

```
tx-flow-control [port {slot/port[-slot/port][, ...]] enable
```

Example

```
VSP-switch:1>enable
VSP-switch:1#configure terminal
VSP-switch:1(config)#interface gigabitethernet 4/24
VSP-switch:1(config-if)#tx-flow-control enable
```

Variable definitions

Use the data in the following table to use the `tx-flow-control` command.

Variable	Value
port {slot/port[-slot/port][, ...]}	Specifies the port or ports to configure.
enable	Enables flow control. After you enable tx-flow-control, the TX port transmits MAC control PAUSE frames to indicate congestion on the receive side of the port interface. The link partner ceases transmission based on the MAC control frame pause-timer value. The default is disabled.

Table continues...

Variable	Value
	<p>You can enable flow control only on 1 Gbps, 10 Gbps and 40 Gbps ports.</p> <p>The default form of this command is <code>default tx-flow-control [port {slot/port[-slot/port]}[,...]] [enable]</code>.</p> <p>The no form of this command is <code>no tx-flow-control [port {slot/port[-slot/port]}[,...]] [enable]</code>.</p>

Associating a port to a VRF instance

Associate a port to a Virtual Router Forwarding (VRF) instance so that the port becomes a member of the VRF instance.

Before you begin

- The VRF instance must exist. For more information about the creation of VRFs, see *Configuring IP Routing on Avaya Virtual Services Platform 9000*, NN46250-505.

About this task

You can assign a VRF instance to a port after you configure the VRF. The system assigns ports to the Global Router, VRF 0, by default.

Procedure

- Enter GigabitEthernet Interface Configuration mode:


```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port]}[,...]
```
- Associate a VRF instance with a port:


```
vrf <WORD 1-16>
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface gigabitethernet 3/21
VSP-9012:1(config-if)#vrf red
```

Variable definitions

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

Variable	Value
<WORD 1-16>	Configures the VRF name.

Chapter 5: Ethernet module configuration using EDM

Configure an Ethernet port to ensure proper operation and optimum performance of the port.

Enabling a module

Enable an Ethernet module to allow traffic to flow through it.

Procedure

1. In the **Device Physical View** tab, select a module.
2. In the navigation pane, expand the following folders: **Configuration > Edit**.
3. Click **Card**.
4. Select **up** in the **AdminStatus** field.
5. Select the power priority in the **PowerManagementPriority** field.
6. Click **Apply**.

Card field descriptions

Use the data in the following table to use the **Card** tab.

Name	Description
CardType	The card type.
CardDescription	A description of the card.
CardSerialNo	The card serial number.
CardPartNo	The card part number.
CardAssemblyDate	The assembly date of the card.
CardHWConfig	The card hardware configuration.
AdminStatus	The administrative status of the card. Use this field to enable, disable, or reset the card.

Table continues...

Name	Description
OperStatus	The current operational status of the card.
ModuleType	The module type.
ModuleDescription	The module description.
ModuleSerialNo	The module serial number.
ModulePartNo	The module part number.
ModuleDateCode	The module date code.
ModuleHWConfig	The module hardware configuration.
PowerManagementPriority	The power management priority setting.
Slot Power	The status of slot power.

Disabling a module

Disable an Ethernet module before you remove it from the chassis to minimize traffic loss. Traffic does not flow on a disabled module.

Procedure

1. In the **Device Physical View** tab, select a module.
2. In the navigation tree, open the following folders: **Configuration > Edit**.
3. Click **Card**.
4. Select **down**.
5. Click **Apply**.

Configuring the minimum module generation support

Configure the minimum generation of modules the Virtual Services Platform 9000 supports. Both Virtual Services Platform 9010 and Virtual Services Platform 9012 support first and second generation modules.

The default is gen1.

About this task

In gen1 mode, VSP 9000 supports both first and second generation modules.

In gen2 mode, VSP 9000 supports only second generation modules. After reboot in second generation-mode, the system disables all first generation modules, and provides full fastpath routing for one million IPv4 FIB routes.

*** Note:**

When you change the linecard-compatibility-mode the system saves the configuration and reboots.

Procedure

1. In the navigation tree, expand the following folders: **Configuration > Edit > Chassis**.
2. Click the **Boot Config** tab.
3. In the **LinecardCompatibilityMode** field, select **gen1** or **gen2**.

Boot Config field descriptions

Use the data in the following table to use the **Boot Config** tab.

Name	Description
Slot	Specifies the slot number of the boot device.
SwVersion	Specifies the software version that currently runs on the chassis.
LastRuntimeConfigSource	Specifies the last source for the run-time configuration.
LinecardCompatibilityMode	<p>Specifies the minimum module generation support for VSP 9000.</p> <ul style="list-style-type: none"> • gen1 specifies VSP 9000 supports both first and second generation modules. When you specify gen1 the device configures the maximum FIB route entries to 400,000 IPv4 FIB routes. • gen2 specifies VSP 9000 supports only second generation modules. In second generation mode, the device configures the maximum FIB route entries to one million IPv4 FIB routes. After reboot in second generation-mode, the system disables all first generation modules, and provides full fastpath routing for one million IPv4 FIB routes. <p>The default is gen1.</p>

Configuring an Ethernet port

Configure an Ethernet port so it operates optimally for your network.

Before you begin

! Important:

- You cannot configure Auto-Negotiation for 10 Gbit/s ports.

- Ensure that all ports that belong to the same MLT or LACP group use the same port speed.

About this task

* Note:

Because Virtual Services Platform 9000 does not negotiate the MTU for IPv4 interfaces, the interface MTU is the maximum-sized packet that the CP transmits. Virtual Services Platform 9000 receives and processes any packet less than the system MTU. In the fastpath, Virtual Services Platform 9000 receives and sends packets less than, or equal to, the system MTU.

For more information about the system MTU, see *Administering Avaya Virtual Services Platform 9000*, NN46250-600.

Procedure

1. In the **Device Physical View** tab, select a port.
2. In the navigation pane, expand the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **Interface** tab.
5. Complete the fields, as required.
6. Click **Apply**.

Interface field descriptions

Use the data in the following table to use the **Interface** tab.

Name	Description
Index	Displays the index of the port, written in the slot/port format.
Name	Configures the name of the port.
Descr	Displays the description of the port. A textual string containing information about the interface. This string should include the name of the manufacturer, the product name and the version of the hardware interface.
Type	Displays the type of connector plugged in the port.
Mtu	Displays the Maximum Transmission Unit (MTU) for the port. The size of the largest datagram that the interface can send or receive in octets. For interfaces used for transmitting network datagrams, this is the size of the largest network datagram that the interface can send.

Table continues...


Name	Description
PhysAddress	Displays the physical address of the port. The address of the interface at the protocol layer immediately `below' the network layer in the protocol stack. For interfaces which do not have such an address (for instance, a serial line), this object should contain an octet string of zero length.
VendorDescr	Displays the vendor of the connector plugged in the port. This option is only applicable to ports on GBIC cards.
AdminStatus	Configures the port as enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
OperStatus	Displays the current status of the port. The status includes enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.
LastChange	Displays the timestamp of the last change.
LinkTrap	Enables or disables link trapping.
AutoNegotiate	<p>Enables or disables Autonegotiation for this port.</p> <p> Note:</p> <p>The 10 GigabitEthernet fiber-based I/O module ports can operate at either 1 Gigabit per second (Gbps) or 10 Gbps, dependent upon the capabilities optical transceiver that you install.</p> <p>This presents an ambiguity with respect to the autonegotiation settings of the port, while 1 Gigabit Ethernet (GbE) ports require autonegotiation; autonegotiation is not defined and is non-existent for 10 GbE ports.</p> <p>For a 10GbE fiber-based I/O module, you have the capability to swap back-and-forth between 1 GbE and 10 GbE operation by simply swapping transceivers. To help with this transition between 1 GbE and 10 GbE port operation, Avaya allows you to configure autonegotiation when you install a 10 GbE transceiver, even though autonegotiation is not defined for 10GbE.</p> <p>You can do this in anticipation of a port changeover from 10 GbE to 1 GbE. In this manner, you could essentially preconfigure a port in 1 GbE mode while the 10 GbE transceiver is still installed. The port is ready to</p>

Table continues...

Name	Description
	<p>go upon the changeover to the 1 GbE transceiver.</p> <p>In addition, you can use a saved configuration file with autonegotiation enabled to boot a system with either 10 GbE or 1 GbE transceivers installed. If you install a 1 GbE transceiver, the system applies autonegotiation. If you install a 10 GbE transceiver, the system does not remove the autonegotiation settings from the configuration, but the system simply ignores the configuration because autonegotiation settings are irrelevant to a 10 GbE transceiver. The system preserves the saved configuration for autonegotiation when resaved no matter which speed of transceiver you install.</p>
AdminDuplex	If AutoNegotiate is false, configures if the port should connect using full duplex or half duplex. The default is half.
OperDuplex	Displays the currently saved AdminDuplex value.
AdminSpeed	If AutoNegotiate is false, configures the speed of the port. The default is 10 Mb/s.
OperSpeed	Displays the currently saved AdminSpeed value.
AutoNegAd	<p>Configures the Custom Autonegotiation Advertisement (CANA) settings of the port.</p> <p>You can only use this option after you enable auto-negotiation on the port. If all bits in this object are disabled, and autonegotiation is enabled on the port, then the physical link process on the port is disabled (if hardware supports this ability.)</p> <p>Any change in the value of this bit map forces the PHY to restart the autonegotiation process, which has the same effect as physically unplugging and reattaching the cable plant attached to this port.</p> <p>The capabilities advertised are either all the capabilities supported by the hardware or the user-configured capabilities, which is a subset of all the capability supported by hardware.</p> <p>The default for this object is all of the capabilities supported by the hardware.</p>
QoSLevel	Selects the Quality of Service (QOS) level for this port. The default is level1.

Table continues...

Name	Description
DiffServ	Enables the Differentiated Service feature for this port. The default is disabled.
Layer3Trust	Configures if the system trusts Layer 3 packets coming from access links or core links only. The default is core.
Layer2Override8021p	Specifies whether Layer 2 802.1p override is enabled (selected) or disabled (cleared) on the port. The default is disabled (clear).
Mltid	Shows the MLT ID associated with this port. The default is 0.
Locked	Shows if the port is locked. The default is disabled.
UnknownMacDiscard	Discards packets that have an unknown source MAC address, and prevents other ports from sending packets with that same MAC address as the destination MAC address. The default is disabled.
AdminRouting	Configures the port as routable or not. The default is enabled.
OperRouting	Displays the currently saved AdminRouting value.
HighSecureEnable	Enables or disables the high secure feature for this port.
IngressRatePeak	Configures the peak rate in Kb/s. The default is 0.
IngressRateSvc	Configures the service rate in Kb/s. The default is 0.
EgressRateLimitState	Enables or disables egress port-based shaping to bind the maximum rate at which traffic leaves the port. The default is disabled.
EgressRateLimit	Configures the egress rate limit in Kb/s. VSP 9000 supports the range 10000 to 10000000. If configured to 0, it means this option is disabled.
Action	<p>Performs one of the following actions on the port:</p> <ul style="list-style-type: none"> • none - Specifies none of the following actions. • flushMacFdb - Flushes the MAC forwarding table. • flushArp - Flushes the ARP table. • flushIp - Flushes the IP route table. • flushAll - Flushes all tables. • triggerRipUpdate - Manually trigger a RIP update. • clearLoopDetectAlarm - Manually enable the port on all the disabled VLANs. <p>The default is none.</p>
Result	Displays result of the selected action. The default is none.

Associating a port to a VRF instance

Associate a port to a Virtual Router Forwarding (VRF) instance so that the port becomes a member of the VRF instance.

You can assign a VRF instance to a port after you configure the VRF. The system assigns ports to the GlobalRouter, VRF 0, by default.

Procedure

1. In the **Device Physical** View tab, select a port.
2. In the navigation tree, open the following folders: **Configuration > Edit > Port**.
3. Click **General**.
4. Click the **VRF** tab.
5. To the right of the **BrouterVrflid** box, click the ellipsis (...) button.
6. In the BrouterVrflid dialog box, select the required VRF.
7. Click **OK**.
8. Click **Apply**.

VRF field descriptions

Use the data in the following table to use the **VRF** tab.

Name	Description
Vrflids	Specifies the ID of the current VRF of this port.
VrfNames	Specifies the name of the current VRF of this port.
VrfCount	Specifies how many VRFs this port is associated to.
BrouterVrflid	Specifies the ID of the new VRF to apply to this port.
BrouterVrfName	Specifies the name of the new VRF to apply to this port.

Chapter 6: Lossless Ethernet

The information in this chapter can help you understand Lossless Ethernet configuration concepts and tasks in relation to Ethernet Modules.

Lossless Ethernet fundamentals

On Virtual Services Platform 9000, you can configure 10 Gigabit Ethernet (GbE) ports to be lossless.

The Lossless Ethernet feature applies to traffic in both directions on a lossless port.

*** Note:**

VSP 9000 supports Lossless Ethernet on first generation modules, which include: 9024XL, 9048GB, and 9048GT modules. VSP 9000 does not support Lossless Ethernet on second generation modules.

Unicast traffic and lossless:

You can configure all unicast traffic on the port to be lossless.

Or you can specify that only tagged unicast traffic, with a specific 802.1p value, is lossless.

Pause types:

If the port becomes congested, the switch performs flow control based on the pause type configuration.

To configure the Lossless Ethernet feature, you enable it on a port and configure the pause type.

The pause type can be one of two values:

- pause (FC) — On the port, all unicast traffic is lossless. FC is the default pause type.
- priority-based flow control (PFC) — On the port, all tagged unicast traffic that has an 802.1p value that matches the lossless-802.1p value is lossless. The switch treats traffic that does not meet this requirement as lossy traffic, and it can drop the lossy traffic.

Lossless 802.1p value:

Following are some limitations that apply to the lossless 802.1p value:

- If traffic enters the device on a lossless PFC-enabled port and exits the device on an lossless FC-enabled port (or vice-versa), only packets with matching lossless-802.1p values are lossless.
- The default lossless 802.1p value is 3 but you can configure it to other values.

*** Note:**

If you change the lossless-802.1p value, the change affects all PFC configured ports. You cannot change this value for individual lossless ports.

- The lossless-802.1p value must be mapped to internal QoS level 3. Avaya recommends that you do not use filters to remark the internal QoS.
- When you enable lossless-PFC on a port, the port does not become Lossless-PFC if the lossless 802.1p value maps to an internal QoS level other than 3, or if the internal QoS level 3 maps to any other 802.1p value.
- You cannot change the lossless 802.1p value to match the configured port QoS value on any of the Lossless-PFC-enabled ports in the system.
- You cannot enable Lossless-PFC on a port if its port QoS value is equal to the lossless 802.1p value.
- On a Lossless-PFC-enabled port, you cannot set the port QoS value to the lossless 802.1p value.
- In a Lossless-PFC (802.1Qbb) domain, the lossless behavior is guaranteed as long as the lossless 802.1p, ingress 1p to QoS map and the egress QoS to 1p maps are consistent.

When you change the lossless 802.1p and ingress 1p to QoS map you must configure the egress QoS to 1p map correctly.

Lossless bandwidth:

For Lossless-PFC, you can configure the guaranteed lossless bandwidth in steps of 1 Gb.

The default lossless bandwidth is 50% of the port bandwidth.

Interface module configuration limitations

The following list identifies configuration limitations for interface module support:

- You can configure Lossless Ethernet only on 10 GbE ports.
 - You cannot configure Lossless Ethernet on a 10 GbE port if you install a 1 GbE SFP.
- You cannot configure Lossless Ethernet on the 9048XS-2, 9048GB, 9012QQ-2, and 9048GT modules.
- Only a 9024XL module with a MAC version of BCM56820-B0 supports Lossless-PFC mode. Use the command `show sys-info ASIC` to view the MAC version of the module.

The ports on the 9024XL module divide into three 8–port clusters:

- ports 1–8
- ports 9–16

- ports 17–24

Clusters and Lossless Ethernet:

A cluster is considered lossless if at least one port in the cluster is lossless-enabled.

On a module you can enable Lossless Ethernet on a maximum of two ports per cluster. The pause type of two lossless ports in the same cluster must be the same.

* **Note:**

If you configure the second port in a cluster with a different pause type, the system prompts you to apply that change to the first lossless port in the cluster.

Within a cluster, one port must be in the top row of the cluster, and one port must be in the bottom row of the cluster.

You cannot enable Lossless Ethernet on two ports in the same row of a cluster.

If you do not enable Lossless Ethernet on any ports in a lossless cluster half, only the lowest numerical port, which is administratively enabled, is operational. The remaining three ports in the cluster half are operationally down.

A maximum of two ports can be operational in a lossless cluster; either two lossless ports, or one lossless port and one lossy port.

If a port in a cluster half is lossless-enabled, the other three ports in that half will be held operationally down.

After you disable Lossless Ethernet on one half of a cluster, the system will allow one of the four ports in that half to be operationally up in lossy mode.

After you disable Lossless Ethernet on the last port in a cluster, all eight ports return to the operational state that they were in prior to Lossless Ethernet configuration.

The system displays a warning if at least one cluster is in Lossless-PFC mode and you

- attempt to map a non-lossless 802.1p value to internal QoS level 3
- map the lossless 802.1p value to an internal QoS level other than 3

General configuration limitations

The following list identifies configuration limitations for Lossless Ethernet:

- If you configure a port as lossless-PFC, it neither generates nor reacts to FC frames.
- On a Lossless-PFC port, lossless is not guaranteed on IP packets with a DSCP value that is set to one of the Avaya-on-Avaya values (47,41,35, and 27).

+ **Tip:**

Avaya-on-Avaya was formerly known as NT-on-NT.

- If you configure a port as lossless-FC, it neither generates nor reacts to PFC frames.
- You cannot configure an interface policer and Lossless Ethernet on the same port.
- If you configure port mirroring on a lossless port, the system does not guarantee lossless behavior.

- If you configure packet capture (PCAP) on a lossless port, the system does not guarantee lossless behavior.
- The chassis must use at least five Switch Fabric modules to support the Lossless Ethernet feature.

The following limitations apply to all lossless ports, irrespective of pause-type:

- You cannot configure Layer 2 untrusted on a lossless port.

Layer 2 is always trusted for a lossless port. If you enable lossless on a Layer 2 untrusted port, the port is forced to be Layer 2 trusted. If you then disable lossless on that port, Layer 2 remains trusted on that port.

- You cannot configure Layer 3 trusted on a lossless port.

Layer 3 is always untrusted for a lossless port. If you enable lossless on a Layer 3 trusted port, the port is forced to be Layer 3 untrusted. If you disable lossless on that port, Layer 3 becomes trusted.

Configuration limitations exist for a MultiLink Trunking (MLT) with lossless port members. For more information on MLT and PFC configuration limitations, conceptual information on Lossless Ethernet and recommended deployment of Lossless Ethernet, see *Network Design Reference for Avaya Virtual Services Platform 9000*, NN46250-200.

Lossless Ethernet and tx-flow-control

By default, tx-flow-control is disabled on a port.

That is, the port does not transmit flow control (FC), or pause, frames.

* Note:

Avaya recommends that you do not enable tx-flow-control on a lossless port.

If you enable tx-flow-control on a lossless port, Avaya cannot guarantee appropriate lossless behavior.

Lossless Ethernet configuration using ACLI

Configuring Lossless Ethernet

To configure the Lossless Ethernet feature, you enable it on a port and configure the pause type.

Before you begin

- The chassis must use at least five Switch Fabric modules to support Lossless Ethernet feature.

* Note:

- You cannot configure an interface policer and Lossless Ethernet on the same port.
- If you configure port mirroring or packet capture (PCAP) on a lossless port, the system does not guarantee lossless behavior.

- Avaya recommends that you do not enable tx-flow control on a lossless port. If you enable tx-flow control on a lossless port, Avaya cannot guarantee appropriate lossless behavior.

Procedure

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Configure the Lossless 802.1p value to the default value:

```
default qos lossless-802.1p
```

3. Configure the Lossless 802.1p value to a different value:

```
qos lossless-802.1p <0-6>
```

4. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port][,...]}
```

5. Configure the pause-type for lossless mode:

```
lossless-port pause-type {fc|pfc}
```

6. If the pause-type is PFC, configure the lossless PFC guaranteed bandwidth:

```
lossless-port pfc-bandwidth <2-8>
```

7. Configure PFC pause-type and PFC-bandwidth in a single step:

```
lossless-port pause-type pfc pfc-bandwidth <2-8>
```

8. Enable the lossless mode for the port:

```
lossless-port enable
```

Example


```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#default qos lossless-802.1p
VSP-9012:1(config)#interface GigabitEthernet 4/16
VSP-9012:1(config-if)#lossless-port pause-type pfc
VSP-9012:1(config-if)#lossless-port pfc-bandwidth 5
VSP-9012:1(config-if)#lossless-port pause-type pfc pfc-bandwidth 4
VSP-9012:1(config-if)#lossless-port enable
```

Variable definitions

Use the data in the following table to use the `lossless-port` command.

Variable	Value
enable	<p>Enables Lossless Ethernet mode for the port.</p> <p>To disable Lossless Ethernet mode for the port use the <code>no</code> parameter: <code>no lossless-port</code></p> <p>The default is disabled</p>
pause-type	<p>Specifies the pause frame type.</p> <ul style="list-style-type: none"> • pause (FC) — All unicast traffic on the port is lossless. • priority-based flow control (PFC) — All tagged unicast traffic on the port with an 802.1p value that matches the lossless-802.1p value is lossless. The switch treats traffic that does not meet this requirement as lossy traffic. The switch can drop the lossy traffic. <p>The default is FC.</p>
pfc-bandwidth <2–8>	<p>Specifies the bandwidth, in Gbps, in a range from 2 to 8, when the pause-type is set to PFC.</p> <p>The default is 5.</p>

Use the data in the following table to use the `qos lossless-802.1p` command.

Variable	Value
qos lossless-802.1p <0–6>	<p>Specifies the lossless-802.1p value in a range from 0 to 6.</p> <p>While you can configure the Lossless-802.1p value, a change affects all of the Lossless-PFC ports in the system.</p> <p>The default is 3.</p> <p> Note:</p> <p>You must map the Lossless-802.1p value to internal QoS level 3 because, when you configure the Lossless-802.1p value, the system does not apply the Lossless-802.1p change if the lossless 802.1p value is mapped to an Internal QoS level other than 3 or if the internal QoS level 3 is mapped to any other 802.1p value.</p> <p>Avaya recommends that you do not use filters to remark the internal QoS.</p>

Displaying Lossless Ethernet information

Display information on which ports are configured for Lossless Ethernet, the pause-type and priority-based flow control (PFC) bandwidth.

Display the configured Lossless 802.1p value in the system.

Display FC and PFC pause transmit and receive counters.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[-slot/port] [, ...]}
```

2. Display information on a port configured for Lossless Ethernet:

```
show interface GigabitEthernet lossless-config {slot/port[-slot/port] [, ...]}
```

3. Display the 802.1p value used for Lossless Ethernet:

```
show qos lossless-802.1p
```

4. Display PFC pause TX-RX counter:

```
show interface GigabitEthernet error verbose [slot/port[-slot/port] [, ...]]
```

5. Display FC and PFC transmit and receive counters.

```
show interface GigabitEthernet statistics {slot/port[slot/port] [, ...]}
```

Example

Show pause-type and priority-based flow control (PFC) bandwidth.

```
VSP-9012:>show interface gigabitethernet lossless-config 5/1, 5/5, 6/3
```

```
=====
                        Port Lossless Config
=====
PORT      LOSSLESS      PAUSE-TYPE      PFC-BW (Gbps)
-----
5/1       TRUE          PFC             5G
5/5       TRUE          PFC             5G
6/3       TRUE          FC              --
```

Show Lossless 802.1p QOS level.

```
VSP-9012:>show qos lossless-802.1p
```

```
=====
                        QOS Lossless-802.1p
```

```
=====
Lossless-802.1p Level: 3
```

```
VSP-9012:>show interface gigabitethernet error verbose 4/1
```

```
=====
Port Ethernet Error Extended
=====
```

PORT NUM	MAC_RX ERRORS	MAC_TX ERRORS	DEFER TX	PACKET ERRORS	LINK INACTIV	UNKNOWN PROTOS
4/1	0	0	0	0	0	0

PORT NUM	IN FLWCTRL	OUT FLWCTRL	IN PFC	OUT PFC
4/1	0	0	0	0

```
=====
```

```
VSP-9012:>show interface gigabitethernet statistics {slot/port[slot/port] [,...]}
```

```
=====
Port Stats Interface
=====
```

PORT NUM	IN OCTETS	OUT OCTETS	IN PACKET	OUT PACKET
4/4	0	0	0	0

PORT NUM	IN FLOWCTRL	OUT FLOWCTRL	IN PFC	OUT PFC	OUTLOSS PACKETS
4/4	0	0	0	0	0

```
=====
```

Lossless Ethernet configuration using EDM

Configuring Lossless Ethernet

To configure the Lossless Ethernet feature, you enable it on a port and configure the pause type.

*** Note:**

Avaya recommends that you do not enable tx-flow control on a lossless port. If you enable tx-flow control on a lossless port, Avaya cannot guarantee appropriate lossless behavior.

You can change the lossless-802.1p value. For more information on this procedure and configuration limitations, see [Specifying the lossless-802.1p value using EDM](#) on page 60

Before you begin

The chassis must use at least five Switch Fabric modules to support the Lossless Ethernet feature.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **Edit** > **Port**

2. Click **General**.
3. Click the **LosslessEthernet** tab.
4. Select the **Configured** check box.
5. In the **PauseType** box, select **fc** or **pfc**.
6. If the **PauseType** is **pfc**, in the **PfcBandwidth** box, select the appropriate value.
7. Click **Apply**.

Lossless Ethernet field descriptions

Use the data in the following table to use the **Lossless Ethernet** tab.

Name	Description
Configured	Enables lossless Ethernet mode for the port. The default is disabled.
PauseType	PauseType selects the pause frame type. <ul style="list-style-type: none"> • fc — All unicast traffic on the port is lossless. • pfc — With priority-based flow control, all tagged unicast traffic on the port that has an 802.1p value that matches the lossless-802.1p value is lossless. The switch treats traffic that does not meet this requirement as lossy traffic. The switch can drop the lossy traffic. The default is fc.
PfcBandwidth	Specifies the bandwidth, in Gbps, in a range from 2 to 8, when the PauseType is set to pfc . The default is 5.

Specifying the lossless-802.1p value

When you configure a port on your switch to use Lossless-PFC, the system compares the 802.1p value of the traffic on the port to the lossless 802.1p value — if both values match, the switch treats the traffic as lossless.

You can configure a specific lossless 802.1p value.

Before you begin

- Enable Lossless Ethernet on the port.
- Specify the pause -type for the traffic on the port.

About this task

If you change the lossless-802.1p value, the change affects all priority-based flow control (PFC) configured ports. You cannot change this value for individual lossless ports.

Procedure

1. In the navigation pane, expand the following folders: **Configuration > Edit**.
2. Click **Chassis**.
3. Click the **System Flags** tab.
4. In the **Lossless8021p** box, type a number.
5. Click **Apply**.

System Flags field descriptions

Use the data in the following table to use the **System Flags** tab.


Name	Description
EnableAccessPolicy	Activates access policies. The default is disabled.
MrouteStreamLimit	Activates or disables Mroute Stream Limit. The default is disabled.
ForceTrapSender	Configures circuitless IP as a trap originator. The default is disabled.
ForceIpHdrSender	If you enable Force IP Header Sender, the system matches the IP header source address with SNMP header sender networks. The default is disabled.
AuthSuccessTrapEnable	Enables the system to send the authentication success trap, rcnAuthenticationSuccess. The default is disabled.  Note: You must also enable SNMP authentication traps. For more information on configuration of SNMP authentication traps, see <i>Troubleshooting Avaya Virtual Services Platform 9000</i> , NN46250-700.
ForceTopologyIpFlagEnable	Activates or disables the flag that configures the CLIP ID as the topology IP. Values are true or false. The default is disabled.
CircuitlessIpId	Uses the CLIP ID as the topology IP. Enter a value from 1–256.
ProfileType	Configures the system to give preference to one type of traffic over the other in times of over subscription. The values are: <ul style="list-style-type: none"> • balanced • unicastOptimized • multicastOptimized The default is balanced.

Table continues...


Name	Description
Lossless8021p	<p>Specifies the lossless-802.1p value. The range is 0 to 6. The default is 3.</p> <p> Note:</p> <p>The internal QoS level that corresponds to the lossless 802.1p value must be 3. Avaya recommends that you do not use filters to remark the internal QoS level.</p> <p>When you enable lossless-PFC on a port, the port cannot become lossless-PFC if the lossless-802.1p value maps to an internal QoS level other than 3, or if the internal QoS level 3 maps to another 802.1p value.</p> <p>In a Lossless-PFC (802.1Qbb) domain, the lossless behavior is guaranteed as long as the Lossless 802.1p, ingress 1p to QoS map and the egress QoS to 1p map are consistent.</p> <p>When you change the Lossless 802.1p and ingress 1p to QoS map, you must configure the egress QoS to 1p map correctly.</p> <p>For more information about Lossless Ethernet, see <i>Network Design Reference for Avaya Virtual Services Platform 9000</i>, NN46250-200, and <i>Administering Avaya Virtual Services Platform 9000</i>, NN46250-600.</p>
HaCpu	<p>Activates or disables the CPU High Availability feature.</p> <p>If you enable or disable High Availability mode, the secondary CPU resets automatically to load settings from the saved configuration file.</p> <p>The default is enable.</p>
HaCpuState	<p>Indicates the CPU High Availability state.</p> <ul style="list-style-type: none"> • initialization—Indicates the CPU is in this state. • oneWayActive—Specifies modules that need to synchronize register with the framework (either locally or a message received from a remote CPU). • twoWayActive—Specifies modules that need to synchronize register with the framework (either locally or a message received from a remote CPU). • synchronized—Specifies table-based synchronization is complete on the current CPU. • remoteIncompatible—Specifies CPU framework version is incompatible with the remote CPU. • error—Specifies if an invalid event is generated in a specific state the CPU enters Error state.

Table continues...

Name	Description
	<ul style="list-style-type: none"> • disabled—Specifies High Availability is not activated. • peerNotConnected—Specifies no established peer connection. • peerConnected—Specifies peer connection is established. • lostPeerConnection—Specifies a lost connection to peer or standby CPU. • notSynchronized—Specifies table-based synchronization is not complete.
HaEvent	<p>Indicates the High Availability event status.</p> <ul style="list-style-type: none"> • restart—Causes the state machine to restart. • systemRegistrationDone—Causes the CPU to transfer to One Way or Two Way Active state. • tableSynchronizationDone—Causes the CPU to transfer to synchronized state. • versionIncompatible—Causes the CPU to go to remote incompatible state • noEvent—Means no event occurred to date.
StandbyCpu	Indicates the state of the standby CPU.

Chapter 7: Port reference

This chapter contains reference information to help guide you to other documentation containing additional information about Ethernet modules.

Port operation reference

While this document focuses on the configuration of Ethernet modules features, a port can take part in many different types of operations. Use the following table to find documents that contain additional procedures about Ethernet modules and details about various port operations.

Reference to other documents

For more information about other port-related parameters, see the documents referenced in the following table.

Table 6: Documentation references

For more information about	See
<code>access-diffserv</code>	<i>Configuring QoS and ACL-Based Traffic Filtering on Avaya Virtual Services Platform 9000, NN46250-502</i>
<code>action</code>	<i>Configuring IP Routing on Avaya Virtual Services Platform 9000, NN46250-505 and Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
<code>auto-recover-port</code>	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
<code>brouter port</code>	<i>Configuring IP Routing on Avaya Virtual Services Platform 9000, NN46250-505</i>
<code>clear mac-address-table dynamic</code>	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
<code>cp-limit</code>	<i>Administering Avaya Virtual Services Platform 9000, NN46250-600</i>
<code>default-vlan-id</code>	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
<code>eapol</code>	<i>Configuring Security on Avaya Virtual Services Platform 9000, NN46250-601</i>
<code>enable-diffserv</code>	<i>Configuring QoS and ACL-Based Traffic Filtering on Avaya Virtual Services Platform 9000, NN46250-502</i>

Table continues...

For more information about	See
encapsulation dot1q	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
fabric	<i>Administering Avaya Virtual Services Platform 9000, NN46250-600</i>
high-secure	<i>Configuring Security on Avaya Virtual Services Platform 9000, NN46250-601</i>
ip	<i>Configuring IP Routing on Avaya Virtual Services Platform 9000, NN46250-505</i>
ipv6	<i>Configuring IPv6 Routing on Avaya Virtual Services Platform 9000, NN46250-509</i>
lacp	<i>Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000, NN46250-503</i>
lock	<i>Configuring Security on Avaya Virtual Services Platform 9000, NN46250-601</i>
loop-detect	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
mac-security	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
pcap	<i>Troubleshooting Avaya Virtual Services Platform 9000, NN46250-700</i>
policy-vlan-precedence	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
protocol-vlan	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
qos	<i>Configuring QoS and ACL-Based Traffic Filtering on Avaya Virtual Services Platform 9000, NN46250-502</i>
rate-limit	<i>Configuring QoS and ACL-Based Traffic Filtering on Avaya Virtual Services Platform 9000, NN46250-502</i>
remote-mirroring	<i>Troubleshooting Avaya Virtual Services Platform 9000, NN46250-700</i>
routing	<i>Configuring IP Routing on Avaya Virtual Services Platform 9000, NN46250-505 and Configuring IPv6 Routing on Avaya Virtual Services Platform 9000, NN46250-509</i>
slpp	<i>Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000, NN46250-503</i>
snmp	<i>Configuring Security on Avaya Virtual Services Platform 9000, NN46250-601</i>
spanning-tree	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>

Table continues...

Port reference

For more information about	See
spoof-detect	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
tagged-frames-discard	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
untagged-frames-discard	<i>Configuring VLANs and Spanning Tree on Avaya Virtual Services Platform 9000, NN46250-500</i>
vlacp	<i>Configuring Link Aggregation, MLT, and SMLT on Avaya Virtual Services Platform 9000, NN46250-503</i>
vrf	<i>Configuring IP Routing on Avaya Virtual Services Platform 9000, NN46250-505</i>

Glossary

Avaya command line interface (ACLI)	A textual user interface. When you use ACLI, you respond to a prompt by typing a command. After you enter the command, you receive a system response.
Custom AutoNegotiation Advertisement (CANA)	An enhancement of the IEEE 802.3 autonegotiation process on the 10/100/1000 copper ports. Custom AutoNegotiation Advertisement offers improved control over the autonegotiation process. The system advertises all port capabilities that include, for tri-speed ports, 10 Mb/s, 100 Mb/s, 1000 Mb/s speeds, and duplex and half-duplex modes of operation. This advertisement results in autonegotiation between the local and remote end that settles on the highest common denominator. Custom AutoNegotiation Advertisement can advertise a user-defined subset of the capabilities that settle on a lower or particular capability.
Data Communications Equipment (DCE)	A network device (such as a modem) that establishes, maintains, and terminates a session.
far end fault indication (FEFI)	Determines that one of two unidirectional fibers, that form the connection between two switches, fails.
gigabit Ethernet (GbE)	Ethernet technology with speeds up to 10 Gbps.
I/O module	An I/O module is a module that provides network connectivity for various media (sometimes called Layer 0) and protocol types. I/O modules are also called Ethernet modules.
Institute of Electrical and Electronics Engineers (IEEE)	An international professional society that issues standards and is a member of the American National Standards Institute, the International Standards Institute, and the International Standards Organization.
interswitch trunking (IST)	A feature that uses one or more parallel point-to-point links to connect two aggregation switches. The two aggregation switches use this channel to share information and operate as a single logical switch. Only one interswitch trunk can exist on each Split Multilink Trunking (SMLT) aggregation switch.

Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
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.
maximum transmission unit (MTU)	The largest number of bytes in a packet—the maximum transmission unit of the port.
Media Access Control (MAC)	MAC arbitrates access to and from a shared medium.
MultiLink Trunking (MLT)	A method of link aggregation that uses multiple Ethernet trunks aggregated to provide a single logical trunk. A multilink trunk provides the combined bandwidth of multiple links and the physical layer protection against the failure of a single link.
multimode fiber (MMF)	A fiber with a core diameter larger than the wavelength of light transmitted that you can use to propagate many modes of light. Commonly used with LED sources for low speed and short distance lengths. Typical core sizes (measured in microns) are 50/125, 62.5/125 and 100/140.
nanometer (nm)	One billionth of a meter (10^{-9} meter). A unit of measure commonly used to express the wavelengths of light.
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.
port	A physical interface that transmits and receives data.
quality of service (QoS)	QoS features reserve resources in a congested network, allowing you to configure a higher priority for certain devices. For example, you can configure a higher priority for IP deskphones, which need a fixed bit rate and split the remaining bandwidth between data connections if calls in the network are more important than the file transfers.
Routing Information Protocol (RIP)	A distance vector protocol in the IP suite, used by IP network-layer protocol, that enables routers in the same AS to exchange routing information by means of periodic updates. You often use RIP as a very simple interior gateway protocol (IGP) within small networks.
Simple Loop Prevention Protocol (SLPP)	Simple Hello Protocol that prevents loops in a Layer 2 network (VLAN).

Simple Network Management Protocol (SNMP)	SNMP administratively monitors network performance through agents and management stations.
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 10 micrometers), with a cladding that is 10 times the core diameter. These fibers have a potential bandwidth of 50 to 100 gigahertz (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 gigabits per second (Gbps).
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.
unshielded twisted pair (UTP)	A cable with one or more pairs of twisted insulated copper conductors bound in a single plastic sheath.
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	An abstract object managed by the Virtual Router Redundancy Protocol (VRRP) that acts as a default router for hosts on a shared LAN.
virtual router forwarding (VRF)	Provides traffic isolation between customers operating over the same node. Each virtual router emulates the behavior of a dedicated hardware router by providing separate routing functionality, and the network treats each VRF as a separate physical router.