Configuring Security on Avaya Virtual Services Platform 9000

Release 4.1
NN46250-601
Issue 08.02
October 2015

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/LicenseInfo 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, UNDER THE LINK “AVAYA SOFTWARE LICENSE TERMS (Avaya Products)” 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 is the software contained within the list of Heritage Nortel Products located at http://support.avaya.com/LicenseInfo under the link “Heritage Nortel Products” or such successor site as designated by Avaya. For Heritage Nortel Software, Avaya grants Customer a license to use Heritage Nortel Software provided hereunder solely to the extent of the authorized activation or authorized usage level, solely for the purpose specified in the Documentation, and solely as embedded in, for execution on, or 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. The open source software license terms provided as Third Party Terms are consistent with the license rights granted in these Software License Terms, and may contain additional rights benefiting You, such as modification and distribution of the open source software. The Third Party Terms shall take precedence over these Software License Terms, solely with respect to the applicable Third Party Components to the extent that these Software License Terms impose greater restrictions on You than the applicable Third Party Terms.

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.

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

Purpose of this document

Security documentation provides procedures and conceptual information that you can use to administer and configure the security features for the Avaya Virtual Services Platform 9000.

The security function includes tasks related to product security; for example, the management and protection of resources from unauthorized or detrimental access and use. Security documents include information that supports the configuration and ongoing management of

- Communications
- Data security
- User security
- Access

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.

Chapter 2: New in this release

The following sections detail what is new in *Configuring Security on Avaya Virtual Services Platform 9000,* NN46250-601, for Release 4.1.

Features

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

Media Access Control Security (MACsec)

Release 4.1 adds support for Media Access Control Security (MACsec) on the Avaya Virtual Services Platform 9000 9048XS-2 Input/Output (I/O) module. MACsec is based on the IEEE 802.1ae standard that allows authorized systems in a network to transmit data confidentially and to protect against data transmitted or modified by unauthorized devices.

In addition to host level authentication, MACsec capable LANs provide data origin authentication, data confidentiality and data integrity between authenticated hosts or systems. MACsec protects data from external hacking while the data passes through the public network to reach a receiver host.

MACsec provides security at the data link layer or the physical layer. It provides enhancements at the MAC service sub layer for its operation and services to the upper layer.

For more information, see:

- MACsec on page 167.

For more information on MACsec, see: *Troubleshooting Avaya Virtual Services Platform 9000,* NN46250-700, and *Monitoring Performance on Avaya Virtual Services Platform 9000,* NN46250-701.

Reverse Path Checking for IPv6

Release 4.1 adds the ability to use reverse path checking for IPv6. After you enable reverse path checking on a port or VLAN, the switch can determine if the IPv4 address associated with a packet is verifiable. If you enable reverse path checking, the switch can reduce the problems caused by spoofed IPv6 addresses into a network. For more information, see

- Configuring reverse path checking for IPv6 on a port on page 35.
- Configuring reverse path checking for IPv6 on a VLAN on page 37.

Other changes

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

Downloading the software

Release 4.1 changes the location of software downloads. For more information, see Downloading the software on page 145.

Chapter 3: Security fundamentals

This section provides conceptual content to help you configure and customize the security services on Avaya Virtual Services Platform 9000.

Security overview

Security is a critical attribute of networking devices, such as the Virtual Services Platform 9000. Security features are split into two main areas:

- Control path—protects the access to the device from a management perspective.
- Data path—protects the network from malicious users by controlling access authorization to the network resources (such as servers and stations). This protection is primarily accomplished by using filters or access lists.

You can protect the control path using:

- logon and passwords
- access policies, in which you specify the network and address that can use a service or daemon
- secure protocols, such as Secure Shell (SSH), Secure Copy (SCP), and the Simple Network Management Protocol version 3 (SNMPv3)
- the Message Digest 5 Algorithm (MD5), which protects routing updates, Open Shortest Path First (OSPF), and Border Gateway Protocol (BGP)

You can protect the data path using:

- Media Access Control (MAC) address filtering
- Layer 3 filtering, such as Internet Protocol (IP) and User Datagram Protocol (UDP)/ Transmission Control Protocol (TCP) filtering
- routing policies, which prevent users from accessing restricted areas of the network
- mechanisms to prevent denial-of-service (DOS) attacks

ACLI passwords

Virtual Services Platform 9000 ships with default passwords assigned for access to Avaya Command Line Interface (ACLI) through a console or management session. If you have read/write/all access authority, and you are using SNMPv3, you can change passwords that are in an encrypted format. If you are using Enterprise Device Manager (EDM), you can also specify the number of available Telnet sessions and rlogin sessions.

Important:

The default passwords are documented and well known. Avaya strongly recommends that you change the default passwords and community strings immediately after you first log on.

For security purposes, if you fail to log on correctly on the master Control Processor (CP) module in three consecutive instances, the CP module locks for 60 seconds.

Port Lock feature

You can use the Port Lock feature to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. You cannot modify locked ports until the ports are first unlocked.

Access policies for services

You can create an access policy to control access to Virtual Services Platform 9000. An access policy specifies the hosts or networks that can access the device through various services, such as Telnet, SNMP, Trivial File Transfer Protocol (TFTP), File Transfer Protocol (FTP), Hypertext Transfer Protocol (HTTP), Remote Shell (RSH), and remote login (rlogin). You can enable or disable access services by setting flags from ACLI.

You can define network stations that can explicitly access Virtual Services Platform 9000 or stations that cannot access it. For each service you can also specify the level of access, such as read-only or read-write-all.

Important:

A third-party security scan shows Virtual Services Platform 9000 service ports open and in the listen state. No connections are accepted on these ports unless you enable the particular daemon. Avaya does not dynamically start and stop the daemons at runtime and needs to keep them running from system startup.

For more information about configuring access policies, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

User-based policy support

You can set up a user-based policy (UBP) system by using Avaya Enterprise Policy Manager (EPM), a RADIUS server, and a Virtual Services Platform 9000 with EAP enabled.

EPM is an application designed to manage the traffic prioritization and network access security for business applications. It provides centralized control of advanced packet classification and the ability to priority mark, police, meter, or block traffic.

EPM 5.0 supports UBPs, which allow security administrators to establish and enforce roles and conditions for each user for all access ports in the network. The UBP feature in EPM works in conjunction with Extensible Access Protocol (EAP) technology to enhance the security of the network. Users log on to the networks and the system authenticates users as the network connection establishes.

The UBP feature works as an extension to the roles feature in EPM. In a UBP environment, role objects are linked directly to specific users (as RADIUS attributes), as opposed to being linked simply to device interfaces. The role object then links the user to specific policies that control the user's access to the network.

When the RADIUS server successfully authenticates a user, the device sends an EAP session start event to the EPM policy server. The policy server then sends user-based policy configuration information for the new user roles to the interface, based on the role attribute assigned to that user on the RADIUS server.

Reverse path checking

When you enable the reverse path checking feature, the feature prevents packet forwarding for incoming IP packets that have incorrect or forged (spoofed) IP addresses. Reverse path checking guarantees that traffic received on one interface was sent by a station from this interface, which prevents address spoofing. When you enable this mode, Virtual Services Platform 9000 performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

You configure reverse path checking for each IP interface. When you enable reverse path checking, Virtual Services Platform 9000 checks all routing packets that come through that interface. It ensures that the source address and source interface appear in the routing table, and that it matches the interface, on which the packet was received.

You can use one of two modes for reverse path checking:

- Exist-only mode: In this mode, reverse path checking checks whether the source IP address of the incoming packet exists in the routing table. If the routing engine finds the source IP entry, the packet forwards as usual; otherwise, the system discards the packet.
- Strict mode: In this mode, reverse path checking checks whether the source IP address of the incoming packet exists in the routing table. If the routing engine does not find the source IP entry, the system drops the packet. If the routing engine finds the source IP entry, reverse path

checking further checks if the source IP interface matches the incoming interface of the packet. If they match, the system forwards the packet as usual, otherwise, the system discards the packet.

The following figure illustrates how strict mode reverse path checking works.

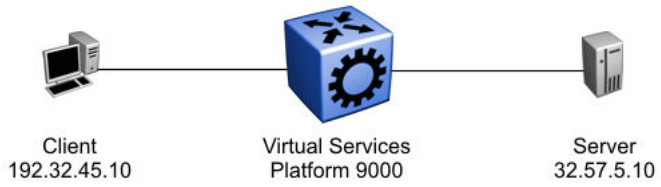

Figure 1: Reverse path checking network configuration

Consider the following parameters:

- Virtual Services Platform 9000 connects a server (32.57.5.10) to a client (192.32.45.10).
- Virtual Services Platform 9000 has reverse path checking enabled.
- The following table details the routing table entries of Virtual Services Platform 9000.

Table 1: Virtual Services Platform 9000 example routing table

Destination address	Next hop address	Forward through port
32.57.5.10	173.56.42.2	3/7
192.32.45.10	145.34.87.2	7/2
192.32.46.10	145.34.88.2	7/1

When a legitimate packet is sent, the following actions occur:

1. The client sends a packet to the server. The packet has a source IP address of 192.32.45.10 and a destination IP address of 32.57.5.10.
2. The packet arrives to Virtual Services Platform 9000 on port 7/2 (brouter); the routing engine performs a destination IP address lookup and finds the destination port is 3/7.
3. Reverse path checking operations begin. The routing engine performs a lookup for the source IP address of 192.32.45.10. The routing engine finds an entry in the routing table that specifies that the next-hop port is 7/2, which matches the incoming port of the packet. Because the address and port information matches, the packet forwards as usual.

When a spoofed packet is sent, the following actions occur:

1. The client sends a packet to the server with a forged IP address of 192.32.46.10 through port 7/2.
2. Reverse path checking finds that the source IP address next-hop port is 7/1, which does not match the incoming port of the packet of 7/2. In this case, the system discards the packet.

You can also think of reverse path checking as follows. If A sends packets to B through route X ingress port Y, then B sends the return packets to A through egress X through the same port Y. If returning packets take a different path, the system discards them.

hsecure mode

Avaya Virtual Services Platform 9000 supports a flag called high secure (hsecure). Hsecure introduces the following behaviors for passwords:

- 10-character enforcement
- aging time
- limitation of failed logon attempts
- protection mechanism to filter certain IP addresses.

After you enable the hsecure flag, the software enforces the 10-character rule for all passwords. This password must contain a minimum of two uppercase characters, two lowercase characters, two numbers, and two special characters.

After you enable hsecure, the system requires you to save the configuration file and reboot the system for hsecure to take effect. If the existing password does not meet the minimum requirements for hsecure, the system prompts you to change the password during the first login.

The default username is rwa and the default password is rwa. In hsecure, the system prompts you to change these during first login because they do not meet the minimum requirements for hsecure.

When you enable hsecure, the system disables Simple Network Management Protocol (SNMP) v1, SNMPv2 and SNMPv3. If you want to use SNMP, you must re-enable SNMP, with the command **`no boot config flag block-snmp`**.

Aging enforcement

After you enable the hsecure flag, you can configure a duration after which you must change your password. You configure the duration by using the aging parameter.

For SNMP and File Transfer Protocol (FTP), after a password expires, the system denies access. Before you access the system, you must change a community string to a new string consisting of more than eight characters.

Important:

Consider the following after you enable the hsecure flag:

- You cannot enable the Web server for Enterprise Device Manager (EDM) access.
- You cannot enable the Secure Shell (SSH) password authentication.

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

Filtering mechanism

Incorrect IP source addresses as network or broadcast addresses are filtered at the virtual router interface. Source addresses 192.168.168.0 and 192.168.168.255 are discarded.

This change is valid for all IP subnets, not only for /24.

You can filter addresses only if you enable the hsecure mode.

Denial-of-service attack prevention with hsecure

To protect Virtual Services Platform 9000 against IP packets with an illegal source address of 255.255.255.255 from being routed (according to RFC1812 Section 4.2.2.11 and RFC971 Section 3.2), Virtual Services Platform 9000 supports a configurable flag, called high secure (hsecure). High secure mode introduces a protection mechanism to filter certain IP addresses, and two restrictions on passwords: 10-character enforcement and aging time.

If the device starts in hsecure mode with default factory settings, and no previously configured password, the system prompts you to change the password. The new password must follow the rules high secure mode mandates. After you enable hsecure and restart the system, if you have an invalid-length password you must change the password.

If you enable hsecure for the first time and the password file does not exist, then the device creates a normal default username (rwa) and password (rwa). In this case, the password does not meet the minimum requirements for hsecure and as a result the system prompts you to change the password.

The following information describes hsecure mode operations:

- When you enable the hsecure flag, after a certain duration the system asks you to change your password. If you do not configure the aging parameter, the aging parameter defaults to 90 days.
- For SNMP and FTP, access is denied when a password expires. You must change the community strings to a new string made up of more than eight characters before accessing the system.
- You cannot enable the Web server at any time.
- You cannot enable the SSH password-authentication feature at any time.
- Incorrect IP source addresses as network or broadcast addresses are filtered at the virtual router interface. Source addresses 192.168.168.0 and 192.168.168.255 are discarded.

Hsecure is disabled by default. When you enable hsecure, the desired behavior (not routing source packets with an IP address of 255.255.255.255) applies to all ports.

Important:

The command to enable high-secure mode only applies to packets going to the CP, not to data path traffic.

DoS protection mechanisms

Several internal mechanisms and features protect Virtual Services Platform 9000 against Denial-of-Service (DoS) attacks.

Broadcast and multicast rate limiting

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

For more information about how to configure the rate limits for broadcast or multicast packets on a port, see *Configuring QoS and ACL-Based Traffic Filtering on Avaya Virtual Services Platform 9000,* NN46250-502.

Directed broadcast suppression

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

To prevent the flooding of other networks with DoS attacks, such as the Smurf attack, Virtual Services Platform 9000 is protected by directed broadcast suppression. This feature is enabled by default. Avaya recommends that you not disable it.

For more information about directed broadcast suppression, see *Configuring Security on Avaya Virtual Services Platform 9000,* NN46250-601.

Prioritization of control traffic

Virtual Services Platform 9000 uses a sophisticated prioritization scheme to schedule control packets on physical ports. This scheme involves two levels with both hardware and software queues to guarantee proper handling of control packets regardless of the switch load. In turn, this guarantees the stability of the network. Prioritization also guarantees that applications that use many broadcasts are handled with lower priority.

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

ARP request threshold recommendations

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

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

For more information about how to configure the ARP threshold, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.

Multicast Learning Limitation

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

For more information, see *Configuring IP Multicast Routing Protocols on Avaya Virtual Services Platform 9000,* NN46250-504.

Configuration considerations

Use the information in this section to understand the limitations of some security functions such as BSAC RADIUS servers and Layer 2 protocols before you attempt to configure security.

Single profile enhancement for BSAC RADIUS servers

Before enabling Remote Access Dial-In User Services (RADIUS) accounting on the device, you must configure at least one RADIUS server.

Virtual Services Platform 9000 software supports Avaya Identity Engines Ignition server. To use these servers, you must first obtain the software for the server. You must also make changes to one or more configuration files for these servers.

Single Profile is a feature that is specific to BSAC RADIUS servers. In a BSAC RADIUS server, when you create a client profile, you can specify all the returnable attributes. When you use the same profile for different products (Virtual Services Platform 9000 and Baystack 450, for example) you specify all the returnable attributes in the single profile.

Attribute format for a third-party RADIUS server

If you use a third-party RADIUS server and need to modify the dictionary files, you must use the following vendor-specific attribute format for ACLI commands:

```
1    1         2                          2+x +----+----+----------------
+---------------------------------------+ |type|len | Vendor-Id      |        value
(string)                    | |    |    |
|                                        | +----+----+----------------
+---------------------------------------
+
|                                            |                              1
1            v   x                          +----+----+-------------------------------
+                             |type|len |     value (cli-command)
|                             |    |    |
|                             +----+----+--------------------------------+
```

RADIUS on management ports

The management port supports the RADIUS protocol. When RADIUS packets are sent out of the management port, the SRC-IP address is properly entered in the RADIUS header. You can hold and synchronize the status of the UDP SRC by a virtual IP flag.

For more information about the supported RADIUS servers, see the documentation of the RADIUS server.

SNMP cloned user considerations

If the user, from which you are cloning has authentication, you can choose for the new user to either have the same authentication protocol as the user, from which it was cloned, or no authentication. If you choose authentication for the new user, you must provide a password for that user. If you want a new user to have authentication, you must indicate that at the time you create the new user. You can assign a privacy protocol only to a user that has authentication.

If the user, from which you are cloning has no authentication, then the new user has no authentication.

Layer 2 redundancy and High Availability clarifications

Layer 2 (L2) redundancy supports the synchronization of VLAN and QoS software parameters. High Availability (CPU-HA) mode, which is an extension to and includes the Layer 2 redundancy software feature, supports the synchronization of VLAN and QoS software parameters, static and default route records, ARP entries, and LAN virtual interfaces. Specifically, CPU-HA mode passes table information and Layer 3 protocol-specific control packets to the standby CP module.

Layer 2 redundancy and CPU-HA mode saves the boot configuration file onto both the master and the secondary CP modules. The secondary CP module resets automatically. You must manually reset the master CP module.

CPU-HA mode limitations and considerations

The following section describes the limitations and considerations of the CPU-HA feature:

- The CPU-HA mode is not compatible with the Packet Capture Tool (PCAP). Be sure to disable the CPU-HA mode before using PCAP.
- You can use the CPU-HA mode to configure redundant ARP and IP static route tables. For more information about creating ARP and IP static routes, see *Configuring IP Routing on Avaya Virtual Services Platform 9000,* NN46250-505.
- Enable CPU-HA mode to disable the brouter port capability; you cannot assign IP addresses to Ethernet ports. To assign an IP address, you must create a VLAN, add ports to that VLAN, and then assign the IP address to it.

Interoperability configuration

Avaya Virtual Services Platform 9000 is compatible with RADIUS servers and EAP servers. For more information about Avaya Virtual Services Platform 9000 RADIUS and EAP compatibility, see *Data Networking — Ignition Server PEAP Active Directory Authentication TCG* (Identity Engines Ignition Server — Ethernet Routing Switch 8800 8300 1600 5500 5600 4500 2500), NN48500–626.

You can search the InSite Knowledge Base on the Avaya Support site at www.avaya.com/support. Use the Advanced Search option to narrow your search to specific categories (products) and document types.

Security configuration using ACLI

Configure security information used on the control and data paths to protect the network from uncontrolled access to network resources.

For more information about how to configure passwords and access policies, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

Enabling hsecure

Before you begin

- For more information about DoS prevention using hsecure, see Denial-of-service attack prevention with hsecure on page 18.
- You must log on to Global Configuration mode in ACLI.

About this task

Use the boot configuration flag hsecure (high security mode) to prevent denial-of-service (DoS) attacks.

The hsecure flag is disabled by default. When you enable it, the software enforces the 10 character rule for all passwords.

When you upgrade from a previous release, if the password does not have at least 10 characters, you receive a prompt to change your password to the mandatory 10-character length.

If you enable hsecure for the first time and the password file does not exist, then the device creates a normal default username (rwa) and password (rwa). In this case, the password does not meet the minimum requirements for hsecure and as a result the system prompts you to change the password.

Procedure

1. Enable or disable hsecure mode:

   ```
   boot config flags hsecure
   ```

   The following warning messages appear:

   ```
   Warning: For security purposes, all unsecure services - TFTP, FTP, Rlogin, Telnet,
   SNMP are disabled. Individually enable the required services.
   Warning: Please save boot configuration and reboot the switch for this to take
   effect.
   ```

2. Save the configuration and restart the device for the change to take effect.

Example

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

Enable hsecure mode:

```
VSP-9012:1(config)#boot config flags hsecure
Warning: For security purposes, all unsecure services - TFTP, FTP,
Rlogin, Telnet, SNMP are disabled. Individually enable the required
services. Warning: Please save boot configuration and reboot the switch
for this to take effect.
```

Save the configuration:

```
VSP-9012:1(config)#save config
```

Restart the switch:

```
VSP-9012:1(config)#reset
Are you sure you want to reset the switch (y/n)? y
```

Changing an invalid-length password

Before you begin

Important:

When you enable hsecure, passwords must contain a minimum of 10 characters or numbers with a maximum of 64. The password must contain a minimum of: two uppercase characters, two lowercase characters, two numbers, and two special characters.

About this task

After you enable **hsecure** and restart the system, change your password if you have an invalid-length password.

Procedure

1. At the ACLI prompt, log on to the system.
2. Enter the password.

   When you have an invalid-length password, the following message appears:

   ```
   Your password is valid but less than mandatory 10 characters.
   Please change the password to continue.
   ```

3. When prompted, enter the new password.
4. When prompted, reenter the new password.

Example

Log on to the switch:

```
Login: rwa
```

Enter the password:

```
Password:***
Your password is valid but less than mandatory 10 characters. Please
chnage the password to continue.
```

Enter the new password:

```
Enter the new password: **********
```

Re-enter the new password:

```
Re-enter the new password: **********
Password successfully changed.
```

Changing passwords

Before you begin

- You must use an account with read-write-all privileges to change passwords. For security, the switch saves passwords to a hidden file.
- You must log on to the Global Configuration mode in ACLI.

About this task

Configure new passwords for each access level, or change the logon or password for the different access levels of the switch. After you receive Avaya Virtual Services Platform 9000, use default passwords to initially access ACLI. If you use Simple Network Management Protocol version 3 (SNMPv3), you can change encrypted passwords.

If you enable the hsecure flag, after the aging time expires, the system prompts you to change your password. If you do not configure the aging time, the default is 90 days.

In hsecure mode, the master Control Processor (CP) module synchronizes the password aging time with the secondary CP module. After the password expires, you must change the password in the master CP module to log on to the secondary CP module.

Procedure

1. Change a password:

   ```
   cli password WORD<1-20> {layer1|layer2|layer3|read-only|read-write|
   read-write-all}
   ```

2. Enter the old password.
3. Enter the new password.
4. Enter the new password a second time.
5. Configure password options:

   ```
   password [access-level WORD<2-8>] [aging-time day <1-365>] [default-
   lockout-time <60-65000>] [lockout WORD<0-46> time <60-65000>] [min-
   passwd-len <10-20>] [password-history <3-32>]
   ```

Example

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

Change a password:

```
VSP-9012:1(config)#cli password smith read-write-all
```

Enter the old password:

```
VSP-9012:1(config)#rwa
```

Enter the new password:

```
VSP-9012:1(config)#TestKey1
```

Enter the new password a second time:

```
VSP-9012:1(config)#TestKey1
```

Set password to an access level of read-write-all and the expiration period for the password to 60 days:

```
VSP-9012:1(config)#password access-level rwa aging-time 60
```

Variable definitions

Use the data in the following table to use the **cli password** command.

Table 2: Variable definitions

Variable	Value
layer1|layer2|layer3|read-only|read-write|read-write-all	Changes the password for the specific access level.
WORD<1–20>	Specifies the user logon name.

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

Table 3: Variable definitions

Variable	Value
access level WORD<2–8>	Permits or blocks this access level. The available access level values are as follows: • l1 • l2 • l3 • ro • rw • rwa
aging-time day *<1-365>*	Configures the expiration period for passwords in days, from 1–365. The default is 90 days.
default-lockout-time *<60-65000>*	Changes the default lockout time after three invalid attempts. Configures the lockout time, in seconds, and is in the 60–65000 range. The default is 60 seconds. To configure this option to the default value, use the default operator with the command.
lockout *WORD<0–46> time <60-65000>*	Configures the host lockout time. • *WORD<0–46>* is the host IP address in the format a.b.c.d.

Table continues…

Variable	Value
	• *<60-65000>* is the lockout-out time, in seconds, in the 60–65000 range. The default is 60 seconds.
min-passwd-len *<10-20>*	Configures the minimum length for passwords in high-secure mode. The default is 10 characters. To configure this option to the default value, use the default operator with the command.
password-history *<3-32>*	Specifies the number of previous passwords the switch stores. You cannot reuse a password that is stored in the password history. The default is 3. To configure this option to the default value, use the default operator with the command.

Synchronizing the master and standby CP module passwords

Before you begin

- You must log on to the Global Configuration mode in ACLI.

Important:

This procedure does not apply to HA mode, where the system automatically synchronizes the passwords on the master and secondary CP module.

About this task

Synchronize the master and secondary CP module passwords.

The secondary CP module does not use the RADIUS protocol to authenticate users who log on to the secondary CP module. The command **`save standby`** saves only the configuration file to the secondary CP module, and does not change the runtime configuration on the secondary CP module.

Procedure

1. Change the password on the master CP module.
2. Save the configuration file to the secondary CP module:

   ```
   save config standby WORD<1-99>
   ```

3. Log on to the secondary CP module.
4. Begin a password change:

   ```
   cli password WORD<1-20> {layer1|layer2|layer3|read-only|read-write|
   read-write-all}
   ```

5. At the `Enter the old password :` prompt, enter the old password.
6. At the `Enter the New password :` prompt, enter the new password.
7. At the `Re-enter the New password :` prompt, enter the same new password.

Variable definitions

Use the data in the following table to use the **cli password** command.

Table 4: Variable definitions

Variable	Value
WORD<1-20>	Specify the user name you require after the password change.
{layer1|layer2|layer3|read-only|read-write|read-write-all}	Specify the access level you require for that user name.

Use the data in the following table to use the **save standby** command.

Table 5: Variable definitions

Variable	Value
<WORD 1-99>	Specify the name of the configuration file on the secondary CP module you must save the current configuration to.

Configuring directed broadcast

Before you begin

- You must log on to VLAN Interface Configuration mode in ACLI.

About this task

A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. When you disable (or suppress) directed broadcasts on an interface, all frames sent to the subnet broadcast address for a local router interface are dropped. Disabling directed broadcasts protects hosts from possible denial-of-service (DOS) attacks. By default, this feature is enabled on the device.

Procedure

Configure Avaya Virtual Services Platform 9000 to forward directed broadcasts for a VLAN:

```
ip directed-broadcast enable
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface vlan 2
VSP-9012:1(config-if)#ip directed-broadcast enable
```

Variable definitions

Use the data in the following table to use the **ip directed-broadcast** command.

Table 6: Variable definitions

Variable	Value
enable	Enables the device to forward directed broadcast frames to the specified VLAN. The default setting for this feature is enabled.

Preventing certain types of DOS attacks

Before you begin

- You must log on to GigabitEthernet Interface Configuration mode in ACLI.

About this task

Protect Avaya Virtual Services Platform 9000 against IP packets with illegal IP addresses such as loopback addresses or a source IP address of ones, or Class D or Class E addresses from being routed. Virtual Services Platform 9000 supports high-secure configurable flag.

Important:

After you enable this flag, the desired behavior (not routing source packets with an IP address of 255.255.255.255) applies to all ports that belong to the same module.

Important:

The command to enable high-secure mode only applies to packets going to the CP, not to datapath traffic.

Procedure

Enable high-secure mode:

```
high-secure [port {slot/port[-slot/port][,...]}] enable
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface GigabitEthernet 4/16
VSP-9012:1(config-if)#high-secure enable
```

Variable definitions

Use the data in the following table to use the **high-secure** command.

Table 7: Variable definitions

Variable	Value
port {slot/port[-slot/port][,...]}	Specifies the port on which you want to enable high-secure mode.
enable	Enables the high-secure feature that blocks packets with illegal IP addresses. This flag is disabled by default. Use the no operator to remove this configuration. To configure this option to the default value, use the default operator with the command.

Configuring reverse path checking on a port

Before you begin

- You must log on to the GigabitEthernet Interface Configuration mode in ACLI.

About this task

You can use the unicast reverse path checking feature to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. When you enable reverse path checking, Avaya Virtual Services Platform 9000 performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Two modes for reverse path checking exist:

- exist-only mode
- strict mode

Procedure

Configure reverse path checking on a port:

```
ip rvs-path-chk mode {exist-only|strict}
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface GigabitEthernet 4/16
```

Check whether the source IP address of the incoming packet exists in the routing table:

```
VSP-9012:1(config-if)#ip rvs-path-chk mode strict
```

Variable definitions

Use the data in the following table to use the **`ip rvs-path-chk mode`** command.

Table 8: Variable definitions

Variable	Value
mode{exist-only|strict}	Specifies the mode for reverse path checking. In exist-only mode, reverse path checking checks whether the source IP address of the incoming packet exists in the routing table. In strict mode, reverse path checking checks whether the source IP address of the incoming packet exists in the routing table. To set this option to the default value, use the default operator with the command. The default is exist-only.

Configuring reverse path checking on a VLAN

Before you begin

- You must log on to VLAN Interface Configuration mode in ACLI.

Important:

You must assign a valid IP address to the selected port.

About this task

You can use the unicast reverse path checking feature to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. When you enable reverse path checking, Virtual Services Platform 9000 performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Two modes for reverse path checking exist:

- exist-only mode
- strict mode

Procedure

Configure reverse path checking on a VLAN:

```
ip rvs-path-chk mode {exist-only|strict}
```

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#interface vlan 2
```

Check whether the source IP address of the incoming packet exists in the routing table:

```
VSP-9012:1(config-if)#ip rvs-path-chk mode strict
```

Variable definitions

Use the data in the following table to use the **`ip rvs-path-chk`** command.

Table 9: Variable definitions

Variable	Value
mode {exist-only|strict}	Specifies the mode for reverse path checking. In exist-only mode, reverse path checking checks whether the source IP address of the incoming packet exists in the routing table. In strict mode, reverse path checking checks whether the source IP address of the incoming packet exists in the routing table. To set this option to the default value, use the default operator with the command. The default is exist-only.

Configuring port lock

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Configure port lock to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. You cannot modify a locked port until you unlock the port.

Procedure

1. Enable port lock globally:

   ```
   portlock enable
   ```

2. Log on to the GigabitEthernet Interface Configuration mode:

   ```
   interface gigabitethernet {slot/port[-slot/port][,...]}
   ```

3. Lock a port:

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

Example

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

Log on to GigabitEthernet Interface Configuration mode:

```
VSP-9012:1(config)#interface GigabitEthernet 4/1
```

Lock port 4/1:

```
VSP-9012:1(config-if)#lock port 4/1 enable
```

Unlock port 4/1:

```
VSP-9012:1(config-if)#no lock port 4/1 enable
```

Variable definitions

Use the data in the following table to use the **`interface gigabitethernet`** command.

Table 10: Variable definitions

Variable	Value
{slot/port[-slot/port][,...]}	Specifies the port you want to configure.

Use the data in the following table to use the **`lock port`** command.

Table 11: Variable definitions

Variable	Value
{slot/port[-slot/port][,...]}	Specifies the port you want to lock. Use the no form of this command to unlock a port: `no lock port {slot/port[-slot/port][,...]}`. The default is disabled.

Security configuration using Enterprise Device Manager

Configure security information used on the control and data paths to protect the network from uncontrolled access to network resources.

For more information about how to configure passwords and access policies, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

Enabling port lock

About this task

Use the port lock feature to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. You cannot modify locked ports until you first unlock the port.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **General**.
3. Click the **Port Lock** tab.
4. To enable port lock, select the **Enable** check box.
5. Click **Apply**.

Port Lock field descriptions

Use the data in the following table to use the **Port Lock** tab.

Name	Description
Enable	Activates the port lock feature. Clear this check box to unlock ports. The default is disabled.
LockedPorts	Lists the locked ports. Click the ellipsis (...) button to select the ports you want to lock or unlock.

Locking a port

Before you begin

- You must enable port lock before you lock or unlock a port.

About this task

Use the port lock feature to administratively lock a port or ports to prevent other users from changing port parameters or modifying port action. You cannot modify locked ports until you first unlock the port.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **General**.
3. Click the **Port Lock** tab.
4. In the **LockedPorts** box, click the ellipsis **(...)** button.
5. Click the desired port or ports.
6. Click **Ok**.
7. In the Port Lock tab, click **Apply**.

Port Lock field descriptions

Use the data in the following table to use the **Port Lock** tab.

Name	Description
Enable	Activates the port lock feature. Clear this check box to unlock ports. The default is disabled.
LockedPorts	Lists the locked ports. Click the ellipsis (...) button to select the ports you want to lock or unlock.

Configuring reverse path checking on a port

Configure reverse path checking on a port to determine if a packet IPv4 address is verifiable. Use reverse path checking to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. After you enable reverse path checking, the

VSP switch performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Reverse path checking operates in one of two modes:

- exist-only mode
- strict mode

Before you begin

- The system supports reverse path checking only on ports that have a valid IPv4 address.

Procedure

1. In the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **IP**.
4. Click the **Reverse Path Checking** tab.
5. Select the **Enable** check box to enable reverse path checking.
6. Select **exist-only** or **strict**.
7. Click **Apply**.

Reverse Path Checking field descriptions

Use the data in the following table to use the **Reverse Path Checking** tab.

Name	Description
Enable	Enables reverse path checking on the selected port. The default is disabled.
Mode	Specifies the mode for reverse path checking. The modes are: • exist-only—Configures reverse path checking to check whether the incoming packet source IPv4 address exists in the routing table. If reverse path checking finds the source IP entry, the system forwards the packet; otherwise the system discards the packet. • strict—Configures reverse path checking to check whether the incoming packet source IPv4 address exists in routing table. If reverse path checking does not find the source IP entry, the system drops the packet; otherwise, reverse path checking further checks if the source IP interface matches the incoming interface of the packet. If they match, the system forwards the packet; otherwise the system discards the packet. The default is exist-only.

Configuring reverse path checking for IPv6 on a port

Configure reverse path checking on a port to determine if a packet IPv6 address is verifiable. Use reverse path checking to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. After you enable reverse path checking, the VSP switch performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Reverse path checking operates in one of two modes:

- exist-only mode
- strict mode

Before you begin

- The system supports reverse path checking only on ports that have a valid IPv6 address.

Procedure

1. n the Device Physical View tab, select a port.
2. In the navigation tree, expand the following folders:**Configuration** > **Edit** > **Port**.
3. Click **IPv6**.
4. Click the **Reverse Path Checking** tab.
5. Select **Enable** from the drop-down list box to enable reverse path checking.
6. Select **exist-only** or **strict** in the Mode field from the drop-down list box.
7. Click **Apply**.

Reverse Path Checking field descriptions

Use the data in the following table to use the **Reverse Path Checking** tab.

Name	Description
Enable	Enables reverse path checking on the selected port. The default is disabled.
Mode	Specifies the mode for reverse path checking. The modes are: • exist-only—Configures reverse path checking to check whether the incoming packet source IP address exists in the routing table. If reverse path checking finds the source IPv6 entry, the system forwards the packet; otherwise the system discards the packet. • strict—Configures reverse path checking to check whether the incoming packet source IPv6 address exists in routing table. If reverse path checking

Table continues…

Name	Description
	does not find the source IP entry, the system drops the packet; otherwise, reverse path checking further checks if the source IP interface matches the incoming interface of the packet. If they match, the system forwards the packet; otherwise the system discards the packet. The default is exist-only.

Configuring reverse path checking on a VLAN

Before you begin

- Before you can configure reverse path checking on a VLAN, you must assign a valid IPv4 address to the selected VLAN.

About this task

Configure reverse path checking on a VLAN to determine if a packet IPv4 address is verifiable. Use reverse path checking to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. After you enable reverse path checking, the VSP switch performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Reverse path checking operates in one of two modes:

- exist-only mode
- strict mode

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the VLAN on which you want to configure reverse path checking.
4. In the toolbar, click **IP**.
5. Click the **Reverse Path Checking** tab.
6. Select the **Enable** box to enable reverse path checking.
7. Select **exist-only** or **strict**.
8. Click **Apply**.

Reverse Path Checking field descriptions

Use the data in the following table to use the **Reverse Path Checking** tab.

Name	Description
Enable	Enables reverse path checking on the selected VLAN. The default is disabled.
Mode	Specifies the mode for reverse path checking. The modes are: • exist-only—Configures reverse path checking to check whether the incoming packet source IPv4 address exists in the routing table. If reverse path checking finds the source IP entry, the system forwards the packet; otherwise, the system discards the packet. • strict—Configures reverse path checking to check whether the incoming packet source IPv4 address exists in routing table. If reverse path checking does not find the source IP entry, then the system drops the packet. Otherwise, reverse path checking further checks if the source IP interface matches the incoming interface of the packet. If they match, then the system forwards the packet. Otherwise, the system discards the packet. The default is exist-only.

Configuring reverse path checking for IPv6 on a VLAN

Configure reverse path checking on a VLAN to determine if a packet IPv6 address is verifiable. Use reverse path checking to reduce the problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network. After you enable reverse path checking, the VSP switch performs a reverse path check to determine if the packet IP address is verifiable. If the address is not verifiable, the system discards the packet.

Reverse path checking operates in one of two modes:

- exist-only mode
- strict mode

Before you begin

- Before you can configure reverse path checking on a VLAN, you must assign a valid IP address to the selected VLAN.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **VLAN**.
2. Click **VLANs**.
3. Click the **Basic** tab.
4. Click the VLAN on which you want to configure reverse path checking.
5. In the toolbar, click **IPv6**.
6. Click the **Reverse Path Checking** tab.

7. Select **Enable** from the drop-down list box to enable reverse path checking.
8. Select **exist-only** or **strict** in the Mode field from the drop-down list box.
9. Click **Apply**.

Reverse Path Checking field descriptions

Use the data in the following table to use the **Reverse Path Checking** tab.

Name	Description
Enable	Enables reverse path checking on the selected VLAN. The default is disabled.
Mode	Specifies the mode for reverse path checking. The modes are: • exist-only—Configures reverse path checking to check whether the incoming packet source IPv6 address exists in the routing table. If reverse path checking finds the source IP entry, the system forwards the packet; otherwise, the system discards the packet. • strict—Configures reverse path checking to check whether the incoming packet source IPv6 address exists in routing table. If reverse path checking does not find the source IP entry, then the system drops the packet. Otherwise, reverse path checking further checks if the source IP interface matches the incoming interface of the packet. If they match, then the system forwards the packet. Otherwise, the system discards the packet. The default is exist-only.

Changing passwords

About this task

Configure new passwords for each access level, or change the logon or password for the different access levels of the system to prevent unauthorized access. After you receive an Avaya Virtual Services Platform 9000, use default passwords to initially access CLI. If you use Simple Network Management Protocol version 3 (SNMPv3), you can change passwords in encrypted format.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **General**.
3. Click the **CLI** tab.

4. Specify the username and password for the appropriate access level.
5. Click **Apply**.

CLI field descriptions

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

Name	Description
RWAUserName	Specifies the user name for the read-write-all CLI account.
RWAPassword	Specifies the password for the read-write-all CLI account.
RWEnable	Activates the read-write access level.
RWUserName	Specifies the user name for the read-write CLI account.
RWPassword	Specifies the password for the read-write CLI account.
RWL3Enable	Activates the read-write Layer 3 access level.
RWL3UserName	Specifies the user name for the Layer 3 read-write CLI account.
RWL3Password	Specifies the password for the Layer 3 read-write CLI account.
RWL2Enable	Activates the read-write Layer 2 access level.
RWL2UserName	Specifies the user name for the Layer 2 read-write CLI account.
RWL2Password	Specifies the password for the Layer 2 read-write CLI account.
RWL1Enable	Activates the read-write Layer 1 access level.
RWL1UserName	Specifies the user name for the Layer 1 read-write CLI account.
RWL1Password	Specifies the password for the Layer 1 read-write CLI account.
ROEnable	Activates the read/only CLI account level.
ROUserName	Specifies the user name for the read-only CLI account.
ROPassword	Specifies the password for the read-only CLI account.
MaxTelnetSessions	Indicates the maximum number of concurrent Telnet sessions (0–8). The default is 8.
MaxRloginSessions	Indicates the maximum number of concurrent Rlogin sessions (0–8). The default is 8.

Table continues…

Name	Description
Timeout	Indicates the number of seconds of inactivity for a Telnet or Rlogin session before automatic timeout and disconnect (30–65535 seconds). The default is 900.
NumAccessViolations	Indicates the number of CLI access violations detected by the system. This field is a read-only field.

Chapter 4: Extensible Authentication Protocol over LAN

Extensible Authentication Protocol over LAN (EAPoL) is a port-based network access control protocol. EAPoL provides security by preventing users from accessing network resources before they are authenticated. The EAPoL authentication feature prevents users from accessing a network to assume a valid identity and access confidential material or launch denial-of-service attacks.

You can use EAPoL to set up network access control on internal LANs and to exchange authentication information between an end station or server that connects to a VSP switch and an authentication server (such as a RADIUS server). This security feature extends the benefits of remote authentication to internal LAN clients. For example, if a new client PC fails the authentication process, EAPoL prevents the new client PC from accessing the network.

Important:

The current release supports only one EAP Supplicant for each port. If the device receives frames from different MAC addresses on the same port, that system disables the port. Avaya is currently working on a solution to support multiple Supplicants. For more information, contact your local representative.

EAPoL terminology

The following section lists some components and terms used with EAPoL-based security.

- Supplicant—a device, such as a PC, that applies for access to the network.
- Authenticator—software on VSP switch that authorizes or rejects a Supplicant attached to the other end of a LAN segment.
  - Port Access Entity (PAE)—software that controls each port on the device. The PAE, which resides on the VSP switch, supports the Authenticator functionality.
  - Controlled Port—any port on the device with EAPoL enabled.
- Authentication Server—a RADIUS server that provides AAA services to the authenticator.

EAPoL configuration considerations

The following section lists EAPoL configuration considerations.

- You must configure at least one EAPoL RADIUS server and shared secret fields.
- You cannot configure EAPoL on ports that are currently configured for the following:
  - Shared segments
  - MultiLink Trunking
  - Port mirroring

- Change the authentication status to auto for each port that you want to control. The *auto* setting automatically authenticates the port according to the results of the RADIUS server. The default authentication setting for each port is *force-authorized*.
- You can connect only a single client on each port that is configured for EAPoL. (If you attempt to add additional clients on the EAPoL authorized port, the port goes to force-unauthorized mode).

Configuration process

The Authenticator facilitates the authentication exchanges that occur between the Supplicant and the Authentication Server. The Authenticator PORT ACCESS ENTITY (PAE) encapsulates the EAPoL message into a RADIUS packet, and then sends the packet to the Authentication Server.

The Authenticator manages the access to controlled port. At system initialization, or when a Supplicant initially connects to one of the controlled ports on the device, the system blocks data traffic of the Supplicant until gets authenticated. After the Authentication Server notifies the Authenticator PAE about the success or failure of the authentication, the Authenticator decides whether to permit/deny the traffic of client on controlled port.

The following figure illustrates how the VSP switch, configured with EAPoL, reacts to a new network connection.

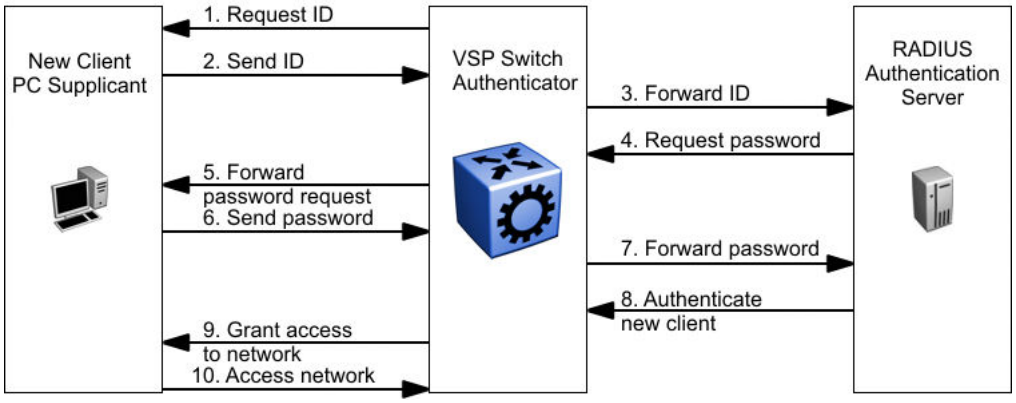

Figure 2: EAPoL configuration example

In Figure 2: EAPoL configuration example on page 42, the switch uses the following steps to authenticate a new client:

1. The VSP switch detects a new connection on one of its EAPoL-enabled ports and requests a user ID from the new client PC.
2. The new client sends its user ID to the VSP switch.
3. The VSP switch uses RADIUS to forward the user ID to the RADIUS server.
4. The RADIUS server responds with a request for the password of the user.
5. The VSP switch forwards the request from the RADIUS server to the new client.
6. The new client sends an encrypted password to the VSP switch, within the EAPoL packet.
7. The VSP switch forwards the EAPoL packet to the RADIUS server.
8. The RADIUS server authenticates the password.
9. The VSP switch grants the new client access to the network.

10. The new client accesses the network.

If the RADIUS server cannot authenticate the new client, it denies the new client access to the network.

The following figure shows the Ethernet frames and the corresponding codes for EAPoL as specified by 802.1x.

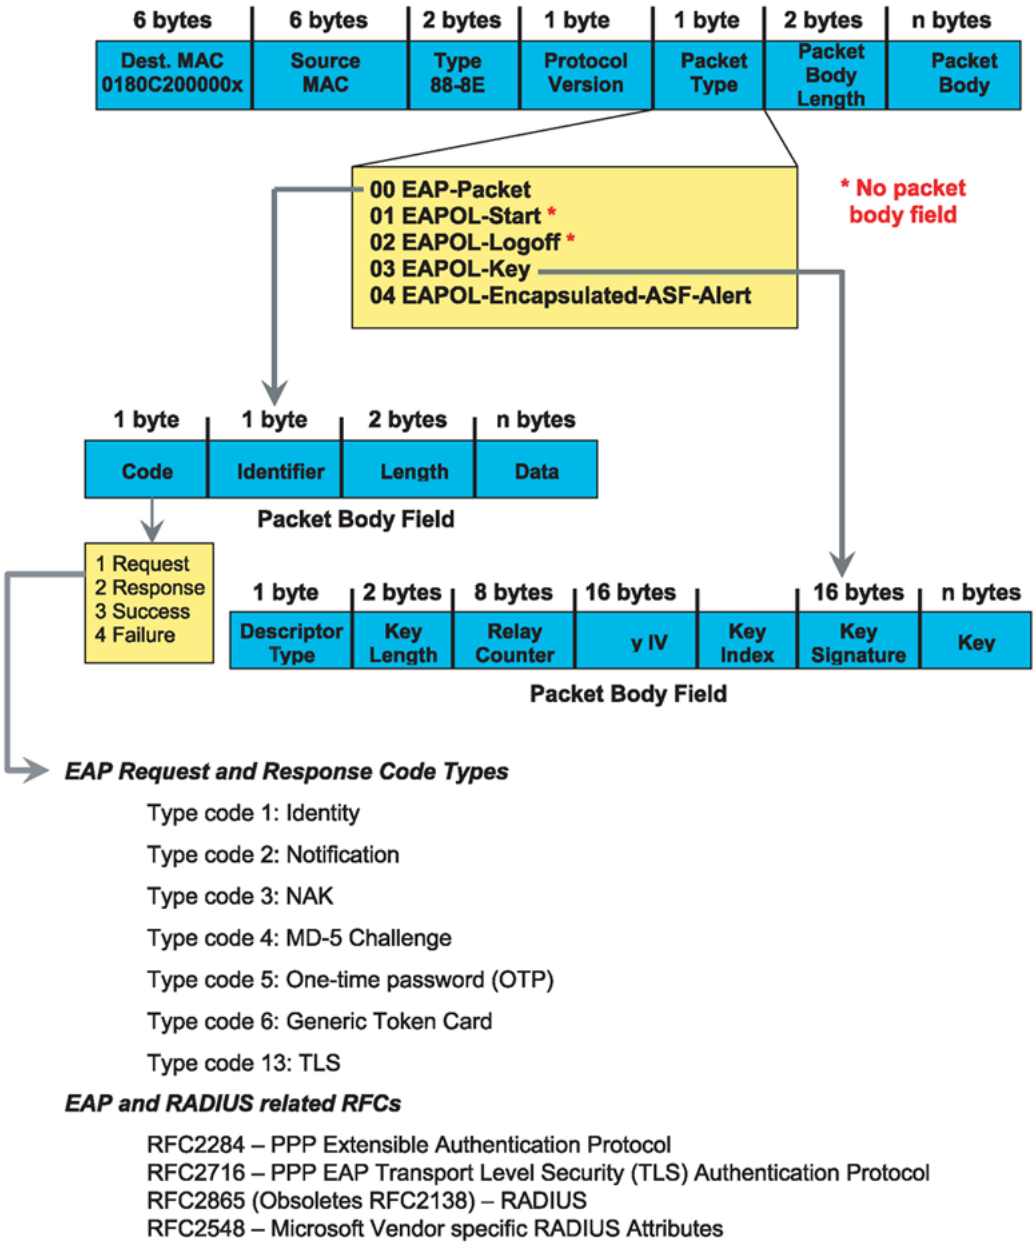

Figure 3: 802.1x Ethernet frame

The following figure shows the flow diagram for EAPoL on a VSP switch.

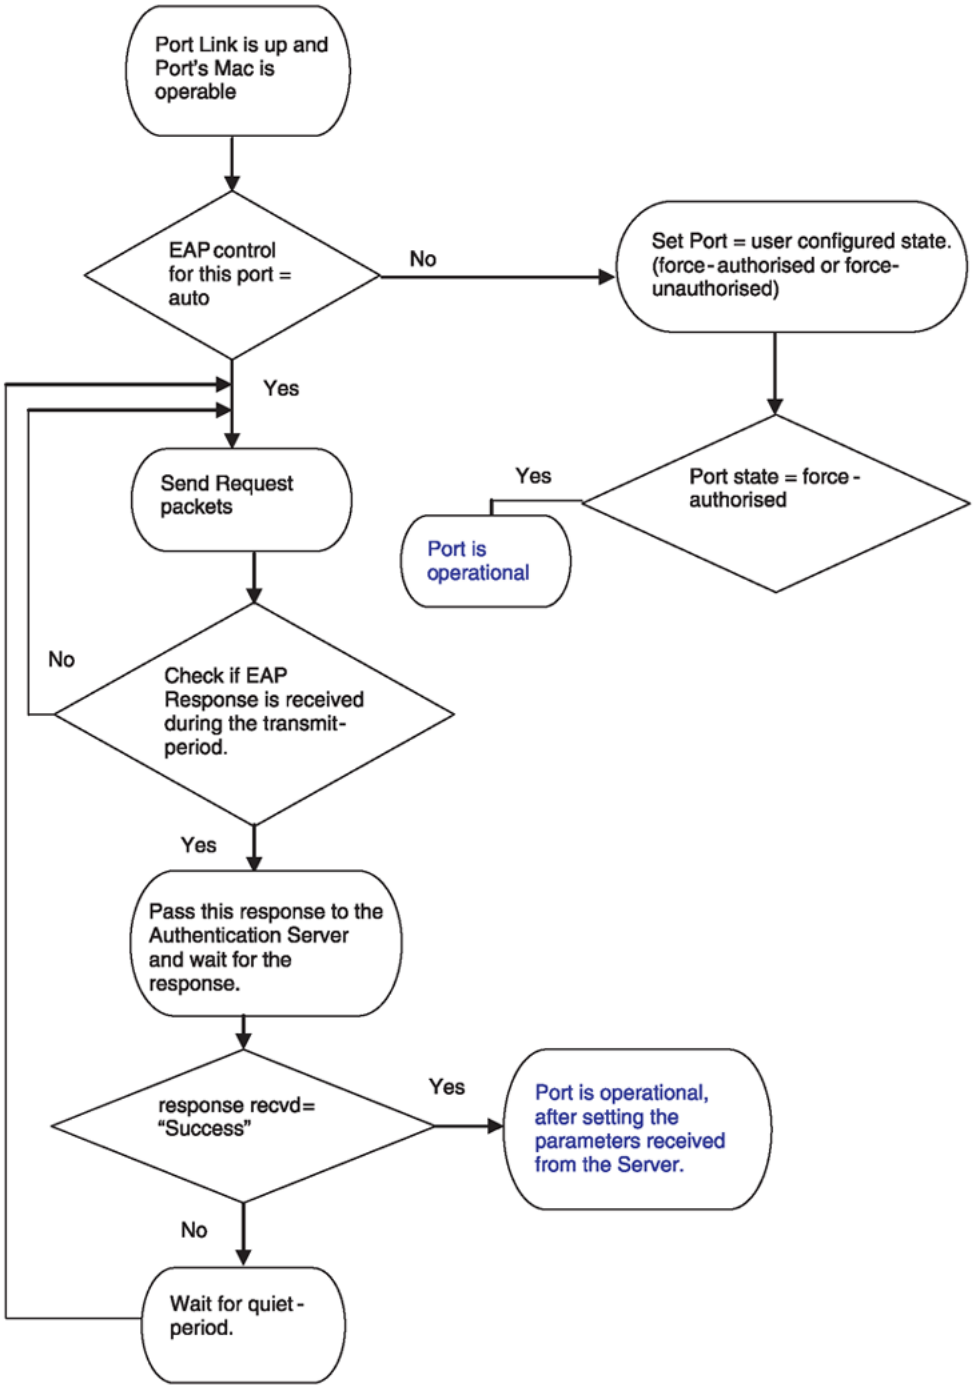

Figure 4: Virtual Services Platforms EAPoL flow diagram

System requirements

The following are the minimum system requirements for EAPoL:

- RADIUS server

- Client software that supports EAPoL

You must specify the RADIUS server that supports EAP as the primary RADIUS server for VSP switch systems. You must configure your VSP switch for VLANs and EAPoL security.

If you configure EAPoL on a port, the following limitations apply:

- You cannot enable EAPoL on tagged ports.
- You cannot enable EAPoL on ports that belong to an MLT group.
- You cannot enable tagging on EAPoL enabled ports.

  Note:

  This includes Switched UNI ports because Switched UNI requires that the port be tagged.
- You cannot add EAPoL-enabled ports to an MLT group.
- You can only configure one Supplicant for each EAPoL-enabled port.

EAPoL dynamic VLAN assignment

If you configure a RADIUS server to send a VLAN ID in the Access-Accept response, the EAPOL feature dynamically changes the VLAN configuration of the port by adding the port to the specified VLAN.

EAPoL dynamic VLAN assignment affects the following VLAN configuration values:

- Port membership
- Port priority

When you disable EAPoL on a port that was previously authorized, VLAN configuration values for that port are restored directly from the nonvolatile random access memory (NVRAM) of the device.

The following exception applies to dynamic VLAN assignments:

- The dynamic VLAN configuration values assigned by EAPoL are not stored in the VSP switch NVRAM.

You can set up your Authentication Server (RADIUS server) for EAPoL dynamic VLAN assignments. You can use the Authentication Server to configure user-specific settings for VLAN memberships and port priority.

When you log on to a system that is configured for EAPoL authentication, the Authentication Server recognizes your user ID and notifies the device to assign preconfigured (user-specific) VLAN membership and port priorities to the device. The configuration settings are based on configuration parameters that were customized for your user ID and previously stored on the Authentication Server.

RADIUS return attributes supported for EAPoL

The VSP switch uses the RADIUS tunnel attributes to place a port into a particular VLAN to support dynamic VLAN switching based on authentication.

The RADIUS server indicates the desired VLAN by including the tunnel attribute within the Access-Accept message. RADIUS uses the following tunnel attributes:

- Tunnel-Type = VLAN (13)
- Tunnel-Medium-Type = 802
- Tunnel-Private-Group-ID = VLAN ID

The VLAN ID is 12 bits, uses a value from <*1-4084*>, and is encoded as a string.

In addition, you can set up the RADIUS server to send a vendor-specific attribute to configure port priority. You can assign the VSP switch Supplicant port a QoS value from 0 to 6.

The following figure shows the RADIUS vendor-specific frame format.

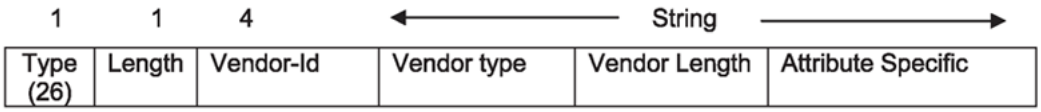

Figure 5: RADIUS vendor-specific frame format

VSP switch Port Priority frame format

- vendor specific type = 26
- length = 12
- vendor-id = 1584
- string = vendor type = 1 + vendor length = 6 + attribute specific = priority

The following figure shows the port priority frame format.

Figure 6: Port priority frame format

RADIUS configuration prerequisites for EAPoL

Connect the RADIUS server to a force-authorized port. This ensures that the port is always available and not tied to whether or not the device is EAPoL-enabled. To set up the Authentication Server, set the following Return List attributes for all user configurations (for more information, see your Authentication Server documentation):

- VLAN membership attributes
  - Tunnel-Type: value 13, Tunnel-Type-VLAN
  - Tunnel-Medium-Type: value 6, Tunnel-Medium-Type-802
  - Tunnel-Private-Group-ID: ASCII value 1 to 4094 (this value identifies the specified VLAN)
- Port priority (vendor-specific) attributes
  - Vendor ID: value 1584, Bay Networks Vendor ID
  - Attribute Number: value 1, Port Priority
  - Attribute Value: value 0 (zero) to 6 (this value indicates the port priority value assigned to the specified user)

Important:

You need to configure these attributes only if you require Dynamic VLAN membership or Dynamic Port priority.

RADIUS accounting for EAPoL

The VSP switch provides the ability to account EAPoL sessions using the RADIUS accounting protocol. A user session is defined as the interval between the instance at which a user is successfully authenticated (port moves to authorized state) and the instance at which the port moves out of the authorized state.

The following table summarizes the accounting events and information logged.

Table 12: Summary of accounting events and information logged

Event	Radius attributes	Description
User is authenticated by EAPoL and port enters authorized state	Acct-Status-Type	Start
	Nas-IP-Address	IP address to represent the VSP switch
	Nas-Port	Port number on which the user is EAPoL authorized
	Acct-Session-ID	Unique string representing the session
	User-Name	EAPoL user name
User logs off and port enters unauthorized state	Acct-Status-Type	Stop
	Nas-IP-Address	IP address to represent the VSP switch
	Nas-Port	Port number on which the user is EAPoL unauthorized
	Acct-Session-ID	Unique string representing the session
	User-Name	EAPoL user name
	Acct-Input-Octets	Number of octets input to the port during the session
	Acct-Output-Octets	Number of octets output to the port during the session
	Acct-Terminate-Cause	Reason for terminating user session. For more information about the mapping of 802.1x session termination cause to RADIUS accounting attribute, see Table 13: 802.1x session termination mapping on page 47.
	Acct-Session-Time	Session interval

The following table describes the mapping of the causes of 802.1x session terminations to the corresponding RADIUS accounting attributes.

Table 13: 802.1x session termination mapping

IEEE 802.1Xdot1xAuthSessionTerminateCause Value	RADIUSAcct-Terminate-Cause Value
supplicantLogoff(1)	User Request (1)
portFailure(2)	Lost Carrier (2)

Table continues…

IEEE 802.1Xdot1xAuthSessionTerminateCause Value	RADIUSAcct-Terminate-Cause Value
supplicantRestart(3)	Supplicant Restart (19)
reauthFailed(4)	Reauthentication Failure (20)
authControlForceUnauth(5)	Admin Reset (6)
portReInit(6)	Port Reinitialized (21)
portAdminDisabled(7)	Port Administratively Disabled (22)
notTerminatedYet(999)	—

Related links

EAPoL configuration using ACLI on page 48
EAPoL configuration using Enterprise Device Manager on page 58

EAPoL configuration using ACLI

Extensible Authentication Protocol over LAN (EAPoL) is a port-based network access-control protocol. EAPoL provides security to your network by preventing users from accessing network resources before they receive authentication.

You can use EAPoL to set up network access control on internal LANs and to exchange authentication information between any end station or server that connects to the VSP switch and an authentication server (such as a RADIUS server). This security feature extends the benefits of remote authentication to internal LAN clients. For example, if a new client PC fails the authentication process, EAPoL prevents the PC from accessing the network.

EAPoL uses RADIUS protocol for EAPoL-authorized logons. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.

Before configuring your device, you must configure at least one EAPoL RADIUS server and shared secret fields.

You cannot configure EAPoL on ports that are currently configured for:

- Shared segments
- MultiLink Turnking (MLT)
- Port mirroring

Change the status of each port that you want to be controlled to auto. The auto setting automatically authenticates the port according to the results of the RADIUS server. The default authentication setting for each port is force-authorized.

You can connect only a single client on each port configured for EAPoL. If you attempt to add additional clients on the EAPoL authorized port, the port goes to force-unauthorized mode.

Globally enabling EAPoL on the device

Enable EAPoL globally on the VSP switch before you enable it on a port or interface.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Globally configure EAPoL:

   ```
   eapol enable
   ```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# eapol enable
```

Configuring EAPoL on an interface

Configure EAPoL on the VSP switch.

Before you begin

- EAPoL must be globally enabled.

About this task

When you configure a port with the EAP status of auto(Authorization depends on result of EAP authentication), only one supplicant is allowed on this port. Multiple EAP supplicants are not allowed on the same physical VSP switch port.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

2. Enable EAPoL on an interface:

   ```
   eapol status {authorized|auto|unauthorized}
   ```

3. Disable EAPoL on on interface:

   ```
   no eapol status
   ```

Example

Enable EAPoL on an interface:

Disable EAPoL on an interface:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 1/2
Switch:1(config-if)# no eapol status
```

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 4/17
Switch:1(config-if)# eapol status authorized
```

Variable definitions

Use the data in the following table to use the **`eapol status`** command.

Variable	Value
authorized	Specifies that the port is always authorized. The default value is authorized.
auto	Specifies that port authorization depends on the results of the EAPoL authentication by the RADIUS server. The default value is authorized.
unauthorized	Specifies that the port is always unauthorized. The default value is authorized.

Configuring EAPoL on a port

Configure EAPoL on a specific port when you do not want to apply EAPoL to all of the VSP switch ports.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

2. Configure the maximum EAP requests sent to the supplicant before timing out the session:

   ```
   eapol port {slot/port[-slot/port][,...]} max-request <1-10>
   ```

3. Configure the time interval between authentication failure and the start of a new authentication:

   ```
   eapol port {slot/port[-slot/port][,...]} quiet-interval <1-65535>
   ```

4. Enable reauthentication:

   ```
   eapol port {slot/port[-slot/port][,...]} re-authentication enable
   ```

5. Configure the time interval between successive authentications:

   ```
   eapol port {slot/port[-slot/port][,...]} re-authentication-period
   <1-2147483647>
   ```

6. Configure the timer for waiting for a RADIUS response:

   ```
   eapol port {slot/port[-slot/port][,...]} server-timeout <1-65535>
   ```

7. Enable an external device to manage the session:

   ```
   eapol port {slot/port[-slot/port][,...]} sess-manage-mode enable
   ```

8. Configure which port to open immediately after 802.1x authentication:

   ```
   eapol port {slot/port[-slot/port][,...]} sess-manage-open-immediate
   enable
   ```

9. Configure the EAP authentication status:

   ```
   eapol port {slot/port[-slot/port][,...]} status {authorized|auto|
   unauthorized}
   ```

10. Configure the wait for supplicant response timer for all EAP packets except EAP Request/ Identity:

    ```
    eapol port {slot/port[-slot/port][,...]} supplicant-timeout
    <1-65535>
    ```

11. Configure the traffic control level:

    ```
    eapol port {slot/port[-slot/port][,...]} traffic-control {in|in-out}
    ```

12. Configure wait time for supplicant :

    ```
    eapol port {slot/port[-slot/port][,...]} transmit-interval <1-65535>
    ```

Example

Configure the maximum EAP requests sent to the supplicant before timing out the session:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface GigabitEthernet 4/17
Switch:1(config-if)#eapol max-request 10
Switch:1(config-if)#eapol port 4/17 quiet-interval 500
```

Variable definitions

Use the data in the following table to use the **`eapol port`** command.

Variable	Value
{slot/port[-slot/port][,...]} c	Specifies the port or list of ports used by EAPoL. Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port).

Table continues…

Variable	Value
max-request <1-10>	Specifies the maximum EAP requests sent to the supplicant before timing out the session. The default is 2.
quiet-interval <1-65535>	Specifies the time interval in seconds between the authentication failure and start of a new authentication. The default is 60.
re-authentication enable	Enables reauthentication of an existing supplicant at a specified time interval.
re-authentication-period <1-2147483647>	Specifies the time interval in seconds between successive reauthentications. The default is 3600 (1 hour).
server-timeout <1-65535>	Specifies the time in seconds to wait for a response from the RADIUS server. The default is 30.
sess-manage-mode enable	Enables an external device to manage the port session.
sess-manage-open-immediate enable	Specifies the port to be opened immediately after 802.1x authentication.
status {authorized|auto|unauthorized}	Specifies the desired EAP authentication status for this port.
supplicant-timeout <1-65535>	Specifies the time in seconds to wait for a response from the supplicant for all EAP packets except EAP Request/Identity.
traffic-control {in|in-out}	Specifies the desired level of traffic control of the port.
transmit-interval <1-65535>	Specifies the time in seconds to wait for a response from the supplicant for EAP Request/Identity packets.

Configuring an EAPoL-enabled RADIUS server

The switch uses RADIUS servers for authentication and accounting services. Use the no form to delete a RADIUS server.

Before you begin

- You must enable EAPoL globally.

About this task

The RADIUS server uses the secret key to validate users.

RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Add an EAPoL-enabled RADIUS server:

```
radius server host WORD <0-46> used-by eapol [key WORD<0-20>] [port
1-65536] [priority <1-10>] [retry <0-6>] [timeout <1-20>] [enable]
[acct-port <1-65536>] [acct-enable] [source-ip WORD <0-46>]
```

By default, the switch uses RADIUS UDP port 1812 for authentication, and port 1813 for accounting. You can change the port numbers or other RADIUS server options.

Example

```
Switch:1> enable

Switch:1# configure terminal
```

Add an EAPoL RADIUS server:

```
Switch:1(config)# radius server host fe80:0:0:0:21b:4fff:fe5e:73fd key
radiustest used-by eapol
```

Variable definitions

Use the data in the following table to configure an EAPoL-enabled RADIUS server with the **`radius server host`** command.

Table 14: Variable definitions

Variable	Value
host *WORD<0–46>*	Specifies the IP address of the selected server. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.
WORD<0-20>	Specifies the secret key, which is a string of up to 20 characters.

Use the data in the following table to use optional arguments of the **`radius server host`** command.

Variable	Value
port *<1-65535>*	Specifies the port ID number.
priority *<1-10>*	Specifies the priority number. The lowest number is the highest priority.
retry *<0-6>*	Specifies the retry count of the account.
timeout *<1-10>*	Specifies the timeout of the server. The default is 30.
enable	Enables the functions used by the RADIUS server host.
acct-port *<1-65536>*	Specifies the port account.
acct-enable	Enables the account.
source-ip *WORD<0–46>*	Specifies the IP source. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.

Configuring the VSP switch for EAPoL and RADIUS

Perform the following procedure to configure the switch for EAPoL and RADIUS.

About this task

You must configure the VSP switch, through which UBP users connect to communicate with the RADIUS server to exchange EAPoL authentication information, as well as user role information. You must specify the IP address of the RADIUS server, as well as the shared secret (a password that authenticates the device with the RADIUS server as an EAPoL access point). You must enable EAPoL globally on each device, and you must configure EAPoL authentication on each device port, through which EAPoL/UBP users connect.

RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.

For more information about EPM and UBP, see the user documentation for your Avaya Enterprise Policy Manager (EPM) application.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Create a RADIUS server that is used by EAPoL:

   ```
   radius server host WORD <0-46> key WORD<0-20> used-by eapol
   ```

3. Log on to the Interface Configuration mode:

   ```
   interface vlan <1-4084>
   ```

4. Enable the device to communicate through EAPoL:

   ```
   eapol enable
   ```

5. Globally enable session management:

   ```
   eapol sess-manage enable
   ```

   Important:

   When EPM learns interfaces on the device, it configures the **eapol sess-manage-mode** command to **enable** on individual interfaces.

6. Exit from VLAN interface mode:

   ```
   exit
   ```

7. Enter Interface Configuration mode:

   ```
   interface GigabitEthernet <slot/port>
   ```

8. Enable device ports for EAPoL authentication:

```
eapol port <slot/port> status auto
```

9. Enable periodic supplicant re-authenticating:

```
eapol port {slot/port[-slot/port][,...]} re-authentication enable
```

10. Save your changes:

```
save config
```

Example

```
Switch:1> enable
Switch:1# configure terminal
```

Create a RADIUS server that is used by EAPoL:

```
Switch:1(config)# radius server host fe90:0:0:0:21b:4eee:fe5e:75fd key
radiustest used-by eapol
Switch:1(config)# interface vlan 2
```

Enable the device to communicate through EAPoL:

```
Switch:1(config-if)# eapol enable
```

Save your changes:

```
Switch:1(config-if)# save config
```

Variable definitions

Use the data in the following table to use the **`radius server host WORD<0–46> usedby eapol`** command.

Table 15: Variable definitions

Variable	Value
host *WORD<0–46>*	Specifies the IP address of the selected server. This address tells the device where to find the RADIUS server, from which it obtains EAPoL authentication and user role information. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.
key *WORD<0-20>*	Specifies the shared secret key that you use for RADIUS authentication. The shared secret is held in common by the RADIUS server and all EAPoL-enabled devices in your network. It authenticates each device with the RADIUS server as an EAPoL access point. When you configure your RADIUS server, you must configure the same shared secret value as you specify here.

Changing the authentication status of a port

The VSP switch authorizes ports by default, which means that the ports are always authorized and are not authenticated by the RADIUS server.

You can also make the ports controlled so that they are dependent on being authorized by the Radius Server when you globally enable EAPoL (auto).

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

2. Configure the authorization status of a port:

   ```
   eapol status {unauthorized|authorized|auto}
   ```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# interface GigabitEthernet 3/1
```

Configure the authorization status of a port:

```
Switch:1(config-if)# eapol status auto
```

Variable definitions

Use the data in the following table to use the **`eapol status`** command.

Variable	Value
authorized	Specifies that the port is always authorized. The default value is authorized.
auto	Specifies that port authorization depends on the results of the EAPoL authentication by the RADIUS server. The default value is authorized.
unauthorized	Specifies that the port is always unauthorized. The default value is authorized.

Deleting an EAPoL-enabled RADIUS server

Delete an EAPoL-enabled RADIUS server if you want to remove the server.

About this task

RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Delete an EAPoL-enabled RADIUS server:

   ```
   no radius server host WORD<0-46> used-by eapol
   ```

Example

```
Switch:1> enable
Switch:1# configure terminal
Switch:1(config)# no radius server host fe79:0:0:0:21d:4fdf:fe5e:73fd
used-by eapol
```

Variable definitions

Use the data in the following table to use the **`radius server host WORD<0-46> usedby eapol`** command.

Table 16: Variable definitions

Variable	Value
host *WORD<0–46>*	Specifies the IP address of the selected server. This address tells the device where to find the RADIUS server, from which it obtains EAPoL authentication and user role information. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.
key *WORD<0-20>*	Specifies the shared secret key that you use for RADIUS authentication. The shared secret is held in common by the RADIUS server and all EAPoL-enabled devices in your network. It authenticates each device with the RADIUS server as an EAPoL access point. When you configure your RADIUS server, you must configure the same shared secret value as you specify here.

EAPoL configuration using Enterprise Device Manager

Extensible Authentication Protocol over LAN (EAPoL) is a port-based network access-control protocol. EAPoL provides security to your network by preventing users from accessing network resources before they receive authentication.

You can use EAPoL to set up network access control on internal LANs and to exchange authentication information between any end station or server that connects to the VSP switch and an authentication server (such as a RADIUS server). This security feature extends the benefits of remote authentication to internal LAN clients. For example, if a new client PC fails the authentication process, EAPoL prevents the PC from accessing the network.

EAPoL uses RADIUS protocol for EAPoL-authorized logons. RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration in all but the following case. When adding a RADIUS server in Enterprise Device Manager (EDM) or modifying a RADIUS configuration in EDM, you must specify if the address type is an IPv4 or an IPv6 address.

Before you begin

- Before configuring your device, you must configure at least one EAPoL RADIUS server and shared secret fields.
- You cannot configure EAPoL on ports that are currently configured for:
  - Shared segments
  - MultiLink Trunking (MLT)
  - Port mirroring
- Change the status of each port that you want to be controlled to auto. For more information on changing the status, see Configuring EAPoL on a port on page 59. The auto setting automatically authenticates the port according to the results of the RADIUS server. The default authentication setting for each port is force-authorized.
- You can connect only a single client on each port configured for EAPoL. If you attempt to add additional clients on the EAPoL authorized port, the port goes to force-unauthorized mode.

Globally configuring EAPoL on the server

About this task

Use SystemAuthControl to globally enable or disable EAPoL on the server. By default, EAPoL is disabled. This feature sets all controlled ports on the server as EAPoL-enabled.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Data Path**.
2. Click **802.1x - EAPOL**.
3. Click the **Global** tab.
4. From the SystemAuthControl, select **enabled**.
5. Click **Apply**.

Global field descriptions

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

Name	Description
SystemAuthControl	Enables system authentication control. EAPoL is disabled by default.

Configuring EAPoL on a port

About this task

Configure EAPoL or change the authentication status on one or more ports.

Ports are force-authorized by default. Force-authorized ports are always authorized and are not authenticated by the RADIUS server. You can change this setting so that the ports are always unauthorized.

You can also make the ports controlled so that they are automatically authenticated when you globally enable EAPoL.

Procedure

1. In the Device Physical View tab, select the port you need to configure.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **General**.
4. Click the **EAPOL** tab.
5. **(Optional)** Select the **PortInitialize** check box to initialize EAPoL authentication on this port.
6. Select the **PortReauthenticate** check box if you want.
7. Select the **AdminControlledDirections** option you want.
8. Select the **AuthControlledPortControl** option you want.
9. In the **QuietPeriod** field, type the time interval.
10. In the **TxPeriod** field, type the time.
11. In the **SuppTimeout** field, type the response time.
12. In the **ServerTimeout** field, type the time.
13. In the **MaxReq** field, type the number of times.
14. In the **ReAuthPeriod** field, type the time between reauthentications.
15. Select the **ReAuthEnabled** field if you want.
16. Click **Apply**.

EAPoL field descriptions

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

Name	Description
PortProtocolVersion	Displays the protocol version number of the EAPoL implementation supported by the port.
PortCapabilities	Displays the capabilities of the Port Access Entity (PAE) associated with the port. This parameter indicates whether Authenticator functionality, supplicant functionality, both, or neither, is supported by the PAE of the port.
PortInitialize	Initializes EAPoL authentication on this port. After the port initializes, this field reverts to its default, which is disabled.
PortReauthenticate	Reauthenticates the supplicant connected to this port immediately. The default is disabled.
PaeState	Displays the current Authenticator PAE state. The possible states are: • initialized • disconnected • connecting • authenticating • authenticated • aborting • held • forceAuth • forceUnauth The default state is forceAuth.
BackendAuthState	Displays the current state of Backend Authentication. The possible states are: • request • response • success • fail • timeout • idle • initialize The default state is idle.

Table continues…

Name	Description
AdminControlDirections	Determines whether the port exerts control over communication in both directions (both incoming and outgoing) or only in the incoming direction. The default value is both.
OperControlledDirections	Displays the current direction of control over communications exerted on the port.
AuthControlledPortStatus	Displays the current state of the port: • unauthorized • auto • authorized The default value is authorized.
AuthControlledPortControl	Configures the authentication status for this port. The default is forceAuthorized. forceUnauthorized—port is always unauthorized. auto—configures the port to match the global EAPoL authentication setting. forceAuthorized—port is always authorized. The default value is forceAuthorized.
QuietPeriod	Configures the time interval (in seconds) between authentication failure and the start of a new authentication. The allowed range is 1–65535; the default is 60.
TxPeriod	Configures the time in seconds to wait for a response from a supplicant for EAP Request/Identity packets. The allowed range is 1–65535; the default is 30.
SuppTimeout	Configures the time (in seconds) to wait for a response from a supplicant for all EAP packets except EAP Request/Identity packets. The allowed range is 1–65535; the default is 30.
ServerTimeout	Configures the time (in seconds) to wait for a response from the RADIUS server. The allowed range is 1–65535; the default is 30.
MaxReq	Configures the maximum number of times to retry sending packets to the supplicant. The allowed range is 1–10; the default is 2.
ReAuthPeriod	Configures the time interval (in seconds) between successive reauthentications.

Table continues…

Name	Description
	The allowed range is 1–2147483647; the default is 3600 (1 hour).
ReAuthEnabled	Reauthenticates an existing supplicant at the time interval specified in ReAuthPeriod.
SessionId	Displays a unique identifier for the session, in the form of a printable ASCII string of at least three characters.
SessionAuthenticMethod	Displays the authentication method used to establish the session.
SessionTime	Displays the duration of the session in seconds.
SessionTerminateCause	Displays the reason for the session termination.
SessionUserName	Displays the user name representing the identity of the supplicant PAE.
LastEapolFrameVersion	Displays the protocol version number carried in the most recently received EAPoL frame.
LastEapolFrameSource	Displays the source MAC address carried in the most recently received EAPoL frame.

Showing the Port Access Entity Port table

About this task

Use the Port Access Entity (PAE) Port Table to display system-level information for each port the PAE supports. An entry appears in this table for each port of this system.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Data Path**.
2. Click **802.1x - EAPOL**.
3. Click the **EAP Security** tab.

EAP Security field descriptions

Use the data in the following table to use the **EAP Security** tab.

Name	Description
PortNumber	Indicates the port number associated with this port.
PortProtocolVersion	Indicates the protocol version associated with this port.
PortCapabilities	Indicates the PAE functionality that this port supports and that can be managed through this MIB. • dot1PaePortAuthCapable(0)—Authenticator functions are supported.

Table continues…

Name	Description
	• dot1xPaeSuppCapable(1)—Supplicant functions are supported.
PortInitialize	Indicates the initialization control for this port. Configure this attribute true to initialize the port. The attribute value reverts to false when initialization is complete.
PortReauthenticate	Specifies the reauthentication control for this port. Setting this attribute true causes the Authenticator PAE state machine for the port to reauthenticate the Supplicant. Setting this attribute false has no effect. This attribute always returns false when it is read.

Showing EAPoL Authentication

About this task

Use the Authenticator Configuration table to display configuration objects for the Authenticator PAE associated with each port.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Data Path**.
2. Click **802.1x - EAPOL**.
3. Click the **Authentication** tab.

Authentication field descriptions

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

Name	Description
PortNumber	Indicates the number associated with this port.
PAEState	Indicates the current value of the authenticator Port Access Entity (PAE) state machine.
BackendAuthState	Indicates the current state of the Backend Authentication state machine.
AdminControlledDirections	Indicates the current value of the administrative controlled directions parameter for the port.
OperControlledDirections	Indicates the current value of the operational controlled directions parameter for the port.
AuthControlledPortStatus	Indicates the current value of the controlled port status parameter for the port.
AuthControlledPortControl	Indicates the current value of the controlled port control parameter for the port.

Table continues…

Name	Description
QuietPeriod	Indicates the value, in seconds, of the QuietPeriod constant currently in use by the Authenticator PAE state machine. The range is 1 to 65535. The default is 60 seconds.
TxPeriod	Indicates the value, in seconds, of the TxPeriod constant currently in use by the Authenticator PAE state machine. The range is 1 to 65535. The default is 30 seconds.
SuppTimeout	Indicates the value, in seconds, of the SuppTimeout constant currently in use by the Backend Authentication state machine. The range is 1 to 65535.
ServerTimeout	Indicates the server timeout value, in seconds, currently in use by the Backend Authentication state machine. The range is 1 to 65535. The default is 30 seconds.
MaxReq	Indicates the value of the maxReq constant currently in use by the Backend Authentication state machine. The range is 1 to 10. The default is 2.
ReAuthPeriod	Indicates the value, in seconds, of the reauthentication interval currently in use by the Reauthentication Timer state machine (8.5.5.1). The default is 3600 seconds.
ReAuthEnabled	Indicates whether reauthentication is enabled (true) or disabled (false). The default is false.

Chapter 5: RADIUS

Remote Access Dial-In User Services (RADIUS) is a distributed client and server system that assists in securing networks against unauthorized access, allowing a number of communication servers and clients to authenticate the identity of users through a central database. The database within the RADIUS server stores information about clients, users, passwords, and access privileges including the use of shared secret.

RADIUS is a fully open and standard protocol, defined by two Requests for Comments (RFC) (Authentication: RFC2865, Accounting: RFC2866). With Virtual Services Platform 9000, you use RADIUS authentication to get secure access to the system (console, Telnet, SSH, EDM), and RADIUS accounting to track the management sessions (ACLI only).

RADIUS support for IPv6

RADIUS supports both IPv4 and IPv6 addresses on Virtual Services Platform 9000 with no differences in functionality or configuration in all but the following case. When you add or update a RADIUS server in Enterprise Device Manager (EDM) you must specify if the address type is an IPv4 or an IPv6 address.

How RADIUS works

A RADIUS application has two components:

• RADIUS server	A computer equipped with server software (for example, a UNIX workstation) that is located at a central office or campus. The server has authentication and access information in a form that is compatible with the client. Typically, the database in the RADIUS server stores client information, user information, password, and access privileges, including the use of a shared secret. A network can have one server for both authentication and accounting, or one server for each service.
• RADIUS client	A device, router, or a remote access server, equipped with client software, that typically resides on the same local area network (LAN) segment as the server. The client is the network access point between the remote users and the server.

The two RADIUS processes are:

- RADIUS authentication—Identifies remote users before you give them access to a central network site.
- RADIUS accounting—Performs data collection on the server during a dial-in session of the remote user with the client.

Configuration of the RADIUS server and client

For more information about how to configure a RADIUS server, see the documentation that came with the server software. See RADIUS configuration example on page 96 and Identity Engine configuration example on page 98 for a configuration of VSP 9000 and an Avaya Identity Engines Igntion Server.

Virtual Services Platform 9000 software supports Avaya Identity Engines Ignition server. To use these servers, you must first obtain the software for the server you will use. Also, you must make changes to one or more configuration files for these servers.

RADIUS authentication

You can use RADIUS authentication to use a remote server to authenticate logons. The RADIUS server also provides access authority. RADIUS assists network security and authorization by managing a database of users. The device uses this database to verify user names and passwords as well as information about the type of access priority available to the user.

Important:

VSP 9000 supports RADIUS Authentication in HA mode.

When the RADIUS client sends an authentication request requesting additional information such as a SecurID number, it sends it as a challenge-response. Along with the challenge-response, it sends a reply-message attribute. The reply-message is a text string, such as "Please enter the next number on your SecurID card:". The RFC-defined maximum length of each reply-message attribute is 253 characters. If you have multiple instances of reply-message attributes that together form a large message that displays to the user, the maximum length is 2000 characters.

You can use additional user names to access the device, in addition to the six existing user names of ro, L1, L2, L3, rw, and rwa. The RADIUS server authenticates the user name and assigns one of the existing access priorities to that name. Unauthenticated user names are denied access to the device. You must add user names ro, L1, L2, L3, rw, and rwa to the RADIUS server if you enable authentication. Users not added to the server are denied access.

The following list shows the user configurable options of the RADIUS feature:

- Up to 10 RADIUS servers in each device for fault tolerance (each server is assigned a priority and is contacted in that order).
- A secret key for each server to authenticate the RADIUS client
- The server UDP port
- Maximum retries allowed
- Time-out period for each attempt

RADIUS authentication on secondary CP modules

Secondary CP modules support RADIUS authentication. To connect to a secondary CP module using RADIUS, you must configure the management port on the secondary CP module with an out-of-band IP address, and a route must exist from the management port to the RADIUS server. In addition, you must configure an entry on the RADIUS server that contains the IP address of the secondary CP module.

However, if you configure the RADIUS source-ip option to use a CLIP address or the management virtual IP address on Virtual Services Platform 9000, then you cannot connect to the secondary CP module with RADIUS.

Use of RADIUS to modify user access to ACLI commands

Virtual Services Platform 9000 provides ACLI command access based on the configured access level of a user. However, you can use RADIUS to override ACLI command access that Virtual Services Platform 9000 provides.

To override user access to ACLI commands, you must configure the command-access-attribute on Virtual Services Platform 9000 and on the RADIUS server. (Virtual Services Platform 9000 uses decimal value 194 as the default for this parameter.) On the RADIUS server, you can then define the commands that the user can or cannot access.

Regardless of the RADIUS server configuration, you must configure the access of the user on Virtual Services Platform 9000 based on the six platform access levels.

RADIUS accounting

RADIUS accounting logs all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Session-IDs for each RADIUS account generate as 12-character strings. The first four characters in the string form a random number in hexadecimal format. The last eight characters in the string indicate the number of user sessions started since the last restart, in hexadecimal format.

The Network Address Server (NAS) IP address for a session is the address of the device interface, to which the remote session connects over the network. For a console session, modem session, and sessions running on debug ports, this value is set to 0.0.0.0, as is the case with RADIUS authentication.

The following table summarizes the events and associated accounting information logged at the RADIUS accounting server.

Table 17: Accounting events and logged information

Event	Accounting information logged at server
Accounting is turned on at router	• Accounting on request: NAS IP address
Accounting is turned off at router	• Accounting off request: NAS IP address
User logs on	• Accounting start request: NAS IP address • Session ID • User name
More than 40 ACLI commands are executed	• Accounting interim request: NAS IP address • Session ID • ACLI commands • User name
User logs off	• Accounting stop request: NAS IP address • Session ID • Session duration • User name • Number of input octets for session

Table continues…

Event	Accounting information logged at server
	• Number of octets output for session • Number of packets input for session • Number of packets output for session • ACLI commands

When the device communicates with the RADIUS accounting server, the following actions occur:

1. If the server sends an invalid response, the response is silently discarded and the server does not make an attempt to resend the request.
2. User-specified number of attempts are made if the server does not respond within the user-configured timeout interval. If a server does not respond to any of the retries, requests are sent to the next priority server (if configured). You can configure up to 10 RADIUS servers for redundancy.

Related links

RADIUS configuration using ACLI on page 68

RADIUS configuration using ACLI

You can configure Remote Access Dial-In User Services (RADIUS) to secure networks against unauthorized access, and allow communication servers and clients to authenticate users identity through a central database.

The database within the RADIUS server stores client information, user information, password, and access privileges, including the use of shared secret.

RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration using ACLI.

RADIUS is a fully open and standard protocol, defined by RFCs (Authentication: RFC2865, accounting RFC2866). With Avaya Virtual Services Platform 9000, you use RADIUS authentication to secure access to the device (console/Telnet/SSH), and RADIUS accounting to track the management sessions for Avaya Command Line Interface (ACLI) only.

RADIUS authentication allows the remote server to authenticate logons. RADIUS accounting logs all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Direct RADIUS authentication is supported on secondary CP modules and in High-Availability (HA) mode.

Configuring RADIUS attributes

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Configure RADIUS to authenticate user identity through a central database.

Procedure

1. Configure RADIUS access priority:

   ```
   radius access-priority-attribute <192-240>
   ```

2. Configure RADIUS accounting:

   ```
   radius accounting {attribute-value <192-240>|enable|include-cli-commands}
   ```

3. Configure the RADIUS authentication info attribute value:

   ```
   radius auth-info-attr-value <0-255>
   ```

4. Clear RADIUS statistics:

   ```
   radius clear-stat
   ```

5. Configure the value of the CLI commands:

   ```
   radius cli-commands-attribute <192-240>
   ```

6. Configure the value of the command access attribute:

   ```
   radius command-access-attribute <192-240>
   ```

7. Configure the maximum number of servers allowed:

   ```
   radius maxserver <1-10>
   ```

8. Configure the multicast address attribute:

   ```
   radius mcast-addr-attr-value <0-255>
   ```

Example

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

Configure RADIUS access priority:

```
VSP-9012:1(config)# radius access-priority-attribute 192
```

Configure RADIUS accounting to include CLI commands:

```
VSP-9012:1(config)# radius accounting include-cli-commands
```

Variable definitions

Use the data in the following table to use the **`radius`** command.

Table 18: Variable definitions

Variable	Value
access-priority-attribute <192-240>	Specifies the value of the access priority attribute in the range of 192 to 240. The default is 192.
accounting {attribute-value <192-240>| enable|include-cli-commands}	Configures the accounting attribute value, enable accounting, or configure if accounting includes CLI commands. The default is false. Use the no option to disable the accounting attribute value: **`no radius accounting enable`**.
auth-info-attr-value <0-255>	Specifies the value of the authentication information attribute in the range of 0 to 255.The default is 91.
clear-stat	Clears RADIUS statistics.
cli-cmd-count <1–40>	Specifies how many ACLI commands, from 1 to 40, before the system sends a RADIUS accounting interim request. The default value is 40.
cli-commands-attribute <192-240>	Specifies the value of ACLI commands attribute in the range of 192 to 240. The default is 195.
cli-profile	Enable RADIUS CLI profiling. ACLI profiling grants or denies access to users being authenticated by way of the RADIUS server. You can add a set of ACLI commands to the configuration on the RADIUS server, and you can specify the command-access more for these commands. The default is false.
command-access-attribute <192-240>	Specifies the value of the command access attribute in the range of 192 to 240. The default is 194.
enable	Enable RADIUS authentication globally on Avaya Virtual Services Platform 9000.
maxserver <1-10>	Specific to RADIUS authentication, configures the maximum number of servers allowed for the device. The range is between 1 and 10. The default is 10.
mcast-addr-attr-value <0-255>	Specifies the value of the multicast address attribute in the range of 0 to 255. The default is 90.
server host *WORD<0–46>* key *WORD<0–32>*[used-by {cli|snmp|eapol| web} [acct-enable] [acct-port *<1–65536>*] [enable] [port *<1–65536>*] [priority *<1–10>*] [retry *<0–6>*] [source-ip *WORD<0–46>*] [timeout *<1–60>*]	• host *WORD<0–46>* Creates a host server. WORD<0–46> signifies an IP address. • key *WORD<0–32>* Specifies a secret key in the range of 0–32 characters. • used-by *{cli|snmp|eapol|web}* Specifies how the server functions. Configures the server for authentication for: - cli - snmp

Table continues…

Variable	Value
	- eapol - web • acct-enable Enables RADIUS accounting on this server. The system enables RADIUS accounting by default. • acct-port *<1–65536>* Specifies a UDP port of the RADIUS accounting server (1 to 65536). The default value is 1816. The UDP port value set for the client must match the UDP value set for the RADIUS server. • enable Enables the server. The default is true. • port *<1–65536>* Specifies a UDP port of the RADIUS server. The default value is 1812. • priority *<1–10>* Specifies the priority value for this server. The default is 10. • retry *<0–6>* Specifies the maximum number of authentication retires. The default is 3. • source-ip *WORD<0–46>* Specifies a configured IP address as the source address when transmitting RADIUS packets. WORD<0–46> signifies an IP address. • timeout *<1–60>* Specifies the number of seconds before the authentication request times out. The default is 3.
sourceip-flag	Enable the source IP so Avaya Virtual Services Platform 9000 uses a configured source IP address. If the outgoing interface on Avaya Virtual Services Platform 9000 fails, a different source IP address is used — requiring that you make configuration changes to define the new RADIUS client on the RADIUS server. To simplify RADIUS server configuration, you can configure Avaya Virtual Services Platform 9000 to use a Circuitless IP (CLIP) address as the source IP and NAS IP address when transmitting RADIUS packets. A CLIP is not associated with a physical interface and is always in an active and operational state. You can configure Avaya Virtual Services Platform 9000 with multiple CLIP interfaces.

Table continues…

Variable	Value
	Note: If you configure the RADIUS source-ip option to use a CLIP address on Virtual Services Platform 9000, then you cannot connect to the secondary CP module using RADIUS. By default, Avaya Virtual Services Platform 9000 uses the IP address of the outgoing interface as the source IP, and the NAS IP address for RADIUS packets that it transmits.

Configuring RADIUS profile

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Use RADIUS ACLI profiling to grant or deny ACLI command access to users being authenticated by way of the RADIUS server. You can add a set of ACLI commands to the configuration file on the radius server, and you can specify the command-access mode for these commands. The default is false.

Procedure

Enable RADIUS ACLI profiling:

```
radius cli-profile
```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# radius cli-profile
```

Enabling RADIUS authentication

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Enable or disable RADIUS authentication globally on the device to allow further configuration to take place. Use the no option to disable RADIUS authentication globally. The default is false or disabled.

Procedure

1. Enable RADIUS authentication globally on Avaya Virtual Services Platform 9000:

```
radius enable

default radius enable
```

2. Disable RADIUS authentication globally on Avaya Virtual Services Platform 9000:

   ```
   no radius enable
   ```

3. Configure RADIUS authentication globally to its default on Avaya Virtual Services Platform 9000:

   ```
   default radius enable
   ```

Enabling the source IP flag for the RADIUS server

Before you begin

- You must log on to the Global Configuration mode in ACLI.
- To configure the CLIP address as the source IP address, you must enable the global RADIUS sourceip-flag. You can then configure the source-ip address parameter while defining the RADIUS server on Virtual Services Platform 9000. The source IP address must be a CLIP address, and that you can configure a different CLIP address for each RADIUS server.

Important:

Use the source IP option only for the RADIUS servers connected to the in-band network.

About this task

By default, Avaya Virtual Services Platform 9000 uses the IP address of the outgoing interface as the source IP, and the NAS IP address for RADIUS packets that it transmits. Enable the source IP so Virtual Services Platform 9000 uses a configured source IP address instead. Therefore, if the outgoing interface on Virtual Services Platform 9000 fails, a different source IP address is used—requiring that you make configuration changes to define the new RADIUS Client on the RADIUS server.

RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration in ACLI.

To simplify RADIUS server configuration, you can configure Virtual Services Platform 9000 to use a CLIP address as the source IP and NAS IP address when transmitting RADIUS packets. A CLIP address is not associated with a physical interface and is always in an active and operational state. You can configure Virtual Services Platform 9000 with multiple CLIP interfaces.

Note:

If you configure the RADIUS source-ip option to use a CLIP address on Virtual Services Platform 9000, then you cannot connect to the secondary CP module using RADIUS.

The default for **`radius sourceip-flag`** is false.

Procedure

Enable the RADIUS packet source IP flag:

```
radius sourceip-flag
```

Enabling RADIUS accounting

Before you begin

- You must configure a RADIUS server before you can enable RADIUS accounting.
- You must log on to the Global Configuration mode in ACLI.

About this task

Enable Remote Access Dial-in User Services (RADIUS) accounting to log all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Procedure

1. Enable RADIUS accounting globally:

   ```
   radius accounting enable
   ```

2. Include or exclude CLI commands in RADIUS accounting updates:

   ```
   radius accounting include-cli-commands
   ```

3. Specify the integer value of the CLI commands attribute:

   ```
   radius accounting attribute-value <192-240>
   ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:(config)# radius accounting enable
VSP-9012:(config)# radius accounting include-cli-commands
```

Variable definitions

Use the data in the following table to use the **`radius accounting`** command.

Table 19: Variable definitions

Variable	Value
enable	Enable RADIUS globally.
include-cli-commands	Include or exclude CLI commands in RADIUS accounting updates.
attribute-value *<192–240>*	Specify the integer value of the CLI commands attribute.

Enabling RADIUS-SNMP accounting

Before you begin

- You must configure a RADIUS server before you can enable RADIUS-SNMP accounting.
- You must log on to the Global Configuration mode in ACLI.

About this task

Enable Remote Access Dial-in User Services (RADIUS) Simple Network Managing Protocol (SNMP) accounting globally. Use SNMP to remotely collect management data. An SNMP agent is a software process that monitors the UDP port 161 for SNMP messages. Each SNMP message sent to the agent contains a list of management objects.

Procedure

1. Enable RADIUS Simple Network Management Protocol (SNMP) accounting globally:

   ```
   radius-snmp acct-enable
   ```

2. Set a timer to send a stop accounting message for RADIUS Simple Network Management Protocol (SNMP):

   ```
   radius-snmp abort-session-timer <30-65535>
   ```

3. Set the timer for re-authentication of the SNMP session:

   ```
   radius-snmp re-auth-timer <30-65535>
   ```

4. Specify the user name for SNMP access:

   ```
   radius-snmp user WORD <0-20>
   ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:(config)# radius-snmp acct-enable
VSP-9012:(config)# radius-snmp abort-session-timer 30
```

Variable definitions

Use the data in the following table to use the **radius-snmp** command.

Table 20: Variable definitions

Variable	Value
acct-enable	Enables RADIUS accounting globally. You cannot enable RADIUS accounting before you configure a valid server. The system disables RADIUS accounting by default. The default is false. Use the no option to disable RADIUS accounting globally: **no radius-snmp acct-enable**
abort-session-timer *<30–65535>*	Set the timer, in seconds, to send a stop accounting message. The default is 180.
re-auth-timer *<30–65535>*	Sets timer for re-authentication of the SNMP session. The timer value ranges from 30 to 65535 seconds. The default is 180.
user *WORD <0–20>*	Specifies the user name for SNMP access. WORD <0–20> specifies the user name in a range of 0 to 20 characters. The default is snmp_user.

Configuring RADIUS accounting interim request

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Configure RADIUS accounting interim requests to create a log whenever a user executes more than the number of ACLI commands you specify.

If the packet size equals or exceeds 1.8 KB, an interim request packet is sent even if the configured limit is not reached. Therefore, the trigger to send out the interim request is either the configured value or a packet size greater than, or equal to 1.8 KB, whichever happens first.

Procedure

1. Configure RADIUS accounting interim requests:

   ```
   radius cli-cmd-count <1-40>
   ```

2. Include or exclude CLI commands in RADIUS accounting:

   ```
   radius accounting include-cli-commands
   ```

   Important:

   You must configure the **`radius accounting include-cli-commands`** command for accounting interim requests to function.

Example

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#radius cli-cmd-count 30
VSP-9012:1(config)#radius accounting include-cli-commands
```

Variable definitions

Use the data in the following table to use the **`radius cli-cmd-count`** command.

Table 21: Variable definitions

Variable	Value
<1-40>	Specifies how many ACLI commands, from 1 to 40, before the system sends a RADIUS accounting interim request. The default value is 40.

Configuring RADIUS authentication and RADIUS accounting attributes

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Configure RADIUS authentication and RADIUS accounting attributes to determine the size of the packets received.

Procedure

1. Configure the RADIUS authentication attribute value:

   ```
   radius command-access-attribute <192-240>
   ```

2. Configure the RADIUS accounting attribute value:

   ```
   radius accounting attribute-value <192-240>
   ```

Example

```
VSP-9012:1> enable
VSP-9012:1# configure terminal
VSP-9012:1(config)# radius command-access-attribute 192
VSP-9012:1(config)# radius accounting attribute-value 192
```

Variable definitions

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

Table 22: Variable definitions

Variable	Value
access-priority-attribute <192-240>	Specifies the value of the access priority attribute in the range of 192 to 240. The default is 192.
accounting {attribute-value <192-240>| enable|include-cli-commands}	Configures the accounting attribute value, enable accounting, or configure if accounting includes CLI commands. The default is false. Use the no option to disable the accounting attribute value: **no radius accounting enable**.
auth-info-attr-value <0-255>	Specifies the value of the authentication information attribute in the range of 0 to 255.The default is 91.
clear-stat	Clears RADIUS statistics.
cli-cmd-count <1–40>	Specifies how many ACLI commands, from 1 to 40, before the system sends a RADIUS accounting interim request. The default value is 40.

Table continues…

Variable	Value
cli-commands-attribute <192-240>	Specifies the value of ACLI commands attribute in the range of 192 to 240. The default is 195.
cli-profile	Enable RADIUS CLI profiling. ACLI profiling grants or denies access to users being authenticated by way of the RADIUS server. You can add a set of ACLI commands to the configuration on the RADIUS server, and you can specify the command-access more for these commands. The default is false.
command-access-attribute <192-240>	Specifies the value of the command access attribute in the range of 192 to 240. The default is 194.
enable	Enable RADIUS authentication globally on Avaya Virtual Services Platform 9000.
maxserver <1-10>	Specific to RADIUS authentication, configures the maximum number of servers allowed for the device. The range is between 1 and 10. The default is 10.
mcast-addr-attr-value <0-255>	Specifies the value of the multicast address attribute in the range of 0 to 255. The default is 90.
server host *WORD<0–46>* key *WORD<0–32>*[used-by {cli|snmp|eapol| web} [acct-enable] [acct-port *<1–65536>*] [enable] [port *<1–65536>*] [priority *<1–10>*] [retry *<0–6>*] [source-ip *WORD<0–46>*] [timeout *<1–60>*]	• host *WORD<0–46>* Creates a host server. WORD<0–46> signifies an IP address. • key *WORD<0–32>* Specifies a secret key in the range of 0–32 characters. • used-by *{cli|snmp|eapol|web}* Specifies how the server functions. Configures the server for authentication for: - cli - snmp - eapol - web • acct-enable Enables RADIUS accounting on this server. The system enables RADIUS accounting by default. • acct-port *<1–65536>* Specifies a UDP port of the RADIUS accounting server (1 to 65536). The default value is 1816. The UDP port value set for the client must match the UDP value set for the RADIUS server. • enable Enables the server. The default is true.

Table continues…

Variable	Value
	• port *<1–65536>* Specifies a UDP port of the RADIUS server. The default value is 1812. • priority *<1–10>* Specifies the priority value for this server. The default is 10. • retry *<0–6>* Specifies the maximum number of authentication retires. The default is 3. • source-ip *WORD<0–46>* Specifies a configured IP address as the source address when transmitting RADIUS packets. WORD<0–46> signifies an IP address. • timeout *<1–60>* Specifies the number of seconds before the authentication request times out. The default is 3.
sourceip-flag	Enable the source IP so Avaya Virtual Services Platform 9000 uses a configured source IP address. If the outgoing interface on Avaya Virtual Services Platform 9000 fails, a different source IP address is used — requiring that you make configuration changes to define the new RADIUS client on the RADIUS server. To simplify RADIUS server configuration, you can configure Avaya Virtual Services Platform 9000 to use a Circuitless IP (CLIP) address as the source IP and NAS IP address when transmitting RADIUS packets. A CLIP is not associated with a physical interface and is always in an active and operational state. You can configure Avaya Virtual Services Platform 9000 with multiple CLIP interfaces. **Note:** If you configure the RADIUS source-ip option to use a CLIP address on Virtual Services Platform 9000, then you cannot connect to the secondary CP module using RADIUS. By default, Avaya Virtual Services Platform 9000 uses the IP address of the outgoing interface as the source IP, and the NAS IP address for RADIUS packets that it transmits.

Adding a RADIUS server

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Add a RADIUS server to allow RADIUS service on Avaya Virtual Services Platform 9000.

RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration using ACLI.

Procedure

Add a RADIUS server:

```
radius server host WORD <0-46> key WORD<0-32> [used-by {cli|snmp|eapol|
web}] [acct-enable][acct-port <1-65536>] [enable] [port <1-65536>]
[priority <1-10>][retry <0-6>] [source-ip WORD <0-46>] [timeout <1-60>]
```

Example

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

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

Add a RADIUS server:

```
VSP-9012:1(config)# radius server host
4717:0000:0000:0000:0000:0000:7933:0001 key testkey1 used-by snmp port 12
retry 5 timeout 10 enable
```

Variable definitions

Use the data in the following table to use the **radius server** command.

Table 23: Variable definitions

Variable	Value
host WORD *<0–46>*	Creates a host server. WORD <0–46> signifies an IPv4 address in the format A.B.C.D or an IPv6 address in the format x:x:x:x:x:x:x:x. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration using ACLI.
key *WORD<0-32>*	Specifies a secret key in the range of 0–32 characters.
used-by {cli|snmp|eapol|web}	Specifies how the server functions • cli—configure the server for CLI authentication. • snmp—configure the server for SNMP authentication. • eapol—configure the server for EAPoL authentication. • web—configure the server for http(s) authentication Use the no option to remove a host server: **no radius server host WORD<0–46> used-by {cli|snmp|eapol|web}**. The default is cli. The default command is: **default radius server**

Table continues…

Variable	Value
	host WORD<0-46> used-by {cli|snmp|eapol|web}
acct-enable	Enables RADIUS accounting on this server. The system enables RADIUS accounting by default.
acct-port *<1-65536>*	Specifies a UDP port of the RADIUS accounting server (1 to 65536). The default value is 1816. **Important:** The UDP port value set for the client must match the UDP value set for the RADIUS server.
enable	Enables this server. The default is true.
port *<1-65536>*	Specifies a UDP port of the RADIUS server. The default value is 1812.
priority *<1-10>*	Specifies the priority value for this server. The default is 10.
retry *<0-6>*	Specifies the maximum number of authentication retries. The default is 3.
source-ip WORD*<0–46>*	Specifies a configured IP address as the source address when transmitting RADIUS packets. WORD <0–46> signifies an IPv4 address in the format A.B.C.D or an IPv6 address in the format x:x:x:x:x:x:x:x. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration using ACLI.
timeout *<1-60>*	Specifies the number of seconds before the authentication request times out. The default is 3.

Modifying RADIUS server settings

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Change a specified RADIUS server value without having to delete the server and recreate it again.

RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration using ACLI.

Procedure

Modify a RADIUS server:

```
radius server host WORD <0-46> [used-by {cli|eapol|snmp|web}] [key
WORD<0-20>] [port 1-65536] [priority <1-10>] [retry <0-6>] [timeout
```

```
<1-20>] [enable] [acct-port <1-65536>] [acct-enable] [source-ip WORD <0-
46>]
```

Example

```
VSP-9012:1> enable

VSP-9012:1# configure terminal
```

Modify a RADIUS server:

```
VSP-9012:1(config)# radius server host
4717:0000:0000:0000:0000:0000:7933:0001 used-by snmp port 12 retry 5
timeout 10 enable
```

Variable definitions

Use the data in the following table to use the **`radius server host`** command.

Table 24: Variable definitions

Variable	Value
used-by {cli|eapol|snmp| web}	Specifies how the server functions: • cli—Configures the server for CLI authentication. • snmp—Configures the server for SNMP authentication. • eapol—Configures the server for EAPoL authentication. • web—Configures the server for Web authentication. Use the no option to remove a host server: **`no radius server host WORD<0-46> used-by {cli|snmp|eapol|web}`**. The default is cli. The default command is: **`default radius server host WORD<0-46> used-by {cli|snmp|eapol|web}`**.
host *WORD <0–46>*	Configures a host server. WORD <0–46> signifies an IPv4 address in the format A.B.C.D or an IPv6 address in the format x:x:x:x:x:x:x:x. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration using ACLI.
acct-enable	Enables RADIUS accounting on this server. The system enables RADIUS accounting by default.
acct-port *<1-65536>*	Configures the UDP port of the RADIUS accounting server (1 to 65536). The default value is 1813. **Important:** The UDP port value set for the client must match the UDP value set for the RADIUS server.
enable	Enables the RADIUS server. The default is true.
key *WORD <0–20>*	Configures the secret key of the authentication client.
port *<1-65536>*	Configures the UDP port of the RADIUS authentication server (1 to 65536). The default value is 1812.

Table continues…

Variable	Value
priority *<1–10>*	Configures the priority value for this server (1 to 10). The default is 10.
retry *<0–6>*	Configures the number of authentication retries the server accepts (0 to 6). The default is 3.
source-ip *WORD <0–46>*	Specifies a configured IP address as the source address when transmitting RADIUS packets. To use this option, you must have the global RADIUS sourceip-flag set to true. RADIUS supports IPv4 and IPv6 addresses, with no difference in functionality or configuration using ACLI.
timeout *<1–20>*	Configures the number of seconds before the authentication request times out (1 to 20). The default is 3.

Showing RADIUS information

About this task

Display the global status of RADIUS information to ensure you configured the RADIUS feature according to the needs of the network.

Procedure

Display the global status of RADIUS information:

```
show radius
```

Example

```
VSP-9012:1>show radius
             acct-attribute-value : 193
                      acct-enable : false
       acct-include-cli-commands : false
       access-priority-attribute : 192
            auth-info-attr-value : 91
        command-access-attribute : 194
          cli-commands-attribute : 195
                   cli-cmd-count : 40
              cli-profile-enable : false
                          enable : false
                       maxserver : 10
           mcast-addr-attr-value : 90
                   sourceip-flag : false
```

Displaying RADIUS server information

About this task

If your system is configured with a RADIUS server you can display the RADIUS server information.

Procedure

To display the RADIUS server information enter the following command:

```
show radius-server
```

Note:

If no RADIUS server is configured, the system displays the following message:

```
no RADIUS server configured
```

Example

```
VSP-9012:1>show radius-server

================================================================================
                              Radius Server Entries
================================================================================
                                                                      ACCT
Name                USED                          TIME EN-   ACCT EN-    SOURE
                    BY   SECRET PORT PRIO RETRY OUT ABLED PORT ABLED IP
1.1.1.1             cli ****** 1812 10   1     3   true  1813 true  0.0.0.0
1000:0:0:0:0:0:0:1 cli ****** 1812 10   1     3   true  1813 true  0:0:0:0:0:0:0:0
10.10.10.10         cli ****** 1812 10   1     3   true  1813 true  0.0.0.0
4000:0:0:0:0:0:0:1 cli ****** 1812 10   1     3   true  1813 true  0:0:0:0:0:0:0:0
```

Showing RADIUS SNMP configurations

About this task

Display current RADIUS SNMP configurations.

Procedure

Display the current RADIUS server SNMP configurations:

```
show radius snmp
```

Example

```
VSP-9012:1>show radius snmp
              abort-session-timer : 180
                      acct-enable : false
                             user : snmp_user
                           enable : false
                    re-auth-timer : 180
```

RADIUS configuration using Enterprise Device Manager

You can configure Remote Access Dial-In User Services (RADIUS) to assist in securing networks against unauthorized access, and allow communication servers and clients to authenticate the identity of users through a central database.

The database within the RADIUS server stores client information, user information, password, and access privileges, including the use of a shared secret.

RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration in all but the following case. When adding a RADIUS server in Enterprise Device Manager (EDM) or modifying a RADIUS configuration in EDM, you must specify if the address type is an IPv4 or an IPv6 address.

RADIUS is a fully open and standard protocol, defined by RFCs (Authentication: RFC2865, accounting RFC2866). With Avaya Virtual Services Platform 9000, you use RADIUS authentication to secure access to the device (console/Telnet/SSH), and RADIUS accounting to track the management sessions for Avaya Command Line Interface (ACLI) only.

RADIUS authentication allows the remote server to authenticate logons. RADIUS accounting logs all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Direct RADIUS authentication is supported on secondary CP modules and in High-Availability (HA) mode.

Enabling RADIUS authentication

About this task

Enable RADIUS authentication globally to allow all features and functions of RADIUS to operate with the RADIUS server.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. In the **RADIUS Global** tab, select the **Enable** check box.
4. In the **MaxNumberServer** field, type a value for the maximum number of servers.
5. In the **AccessPriorityAttrValue** field, type an access policy value (by default, this value is 192).
6. Configure the rest of the parameters in the RADIUS global tab.
7. Click **Apply**.

RADIUS Global field descriptions

Use the data in the following table to use the **RADIUS Global** tab.

Name	Description
Enable	Enables the RADIUS authentication feature globally.
MaxNumberServer	Specifies the maximum number of servers to be used, between 1 and 10, inclusive.
AccessPriorityAttrValue	Specific to RADIUS authentication. Specifies the vendor-specific attribute value of the access-priority attribute to match the type

Table continues…

Name	Description
	value set in the dictionary file on the RADIUS server. The valid values are 192 through 240. Avaya recommends the default setting of 192 for Virtual Services Platform 9000.
AcctEnable	Enables RADIUS accounting.
AcctAttriValue	Specific to RADIUS accounting. Specifies the vendor-specific attribute value of the CLI-command attribute to match the type value set in the dictionary file on the RADIUS server. This value must be different from the access-priority attribute value configured for authentication. The valid values are 192 through 240. The default value is 193.
AcctIncludeCli	Specifies whether you want CLI commands included in RADIUS accounting requests.
ClearStat	Clears RADIUS statistics from the device.
McastAttributeValue	Specifies the value of the Mcast attribute. The valid values are 0 through 255. The default value is 90.
AuthInfoAttrValue	Specifies the value of the authentication information attribute. The valid values are 0 through 255. The default value is 91.
CommandAccessAttrValue	Specifies the value of the command access attribute. The valid values are 192 through 240. The default value is 194.
CliCommandAttrValue	Specifies the value of the CLI command attribute. The valid values are 192 through 240. The default value is 195.
AuthInvalidServerAddress	Displays the number of access responses from unknown or invalid RADIUS servers.
SourceIpFlag	Includes a configured IP address as the source address in RADIUS packets. The default is false. RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration.
CliCmdCount	Gives the value for the CLI command count. Specify an integer from 1 to 40. The default is 40.
CliProfEnable	Enables RADIUS CLI profiling.

Enabling RADIUS accounting

Before you begin

- You must set up a RADIUS server and add it to the configuration file of the device before you can enable RADIUS accounting on the device. Otherwise, the system displays an error message.

About this task

Enable RADIUS accounting to log all of the activity of each remote user in a session on the centralized RADIUS accounting server.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. In the **RADIUS Global** tab, select the **AcctEnable** check box.
4. In the **AcctAttrValue** field, type an access policy value (by default, this value is 193).
5. Click **Apply**.

RADIUS Global field descriptions

Use the data in the following table to use the **RADIUS Global** tab.

Name	Description
Enable	Enables the RADIUS authentication feature globally.
MaxNumberServer	Specifies the maximum number of servers to be used, between 1 and 10, inclusive.
AccessPriorityAttrValue	Specific to RADIUS authentication. Specifies the vendor-specific attribute value of the access-priority attribute to match the type value set in the dictionary file on the RADIUS server. The valid values are 192 through 240. Avaya recommends the default setting of 192 for Virtual Services Platform 9000.
AcctEnable	Enables RADIUS accounting.
AcctAttriValue	Specific to RADIUS accounting. Specifies the vendor-specific attribute value of the CLI-command attribute to match the type value set in the dictionary file on the RADIUS server. This value must be different from the access-priority attribute value configured for authentication. The valid values are 192 through 240. The default value is 193.
AcctIncludeCli	Specifies whether you want CLI commands included in RADIUS accounting requests.
ClearStat	Clears RADIUS statistics from the device.
McastAttributeValue	Specifies the value of the Mcast attribute. The valid values are 0 through 255. The default value is 90.
AuthInfoAttrValue	Specifies the value of the authentication information attribute. The valid values are 0 through 255. The default value is 91.
CommandAccessAttrValue	Specifies the value of the command access attribute. The valid values are 192 through 240. The default value is 194.
CliCommandAttrValue	Specifies the value of the CLI command attribute. The valid values are 192 through 240. The default value is 195.
AuthInvalidServerAddress	Displays the number of access responses from unknown or invalid RADIUS servers.

Table continues…

Name	Description
SourceIpFlag	Includes a configured IP address as the source address in RADIUS packets. The default is false. RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration.
CliCmdCount	Gives the value for the CLI command count. Specify an integer from 1 to 40. The default is 40.
CliProfEnable	Enables RADIUS CLI profiling.

Disabling RADIUS accounting

Before you begin

- You cannot globally disable RADIUS accounting unless a server entry exists.

About this task

Disabling RADIUS accounting removes the accounting function from the RADIUS server.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. In the **RADIUS Global** tab, disable RADIUS accounting by clearing the **AcctEnable** check box.
4. Click **Apply**.

Enabling RADIUS accounting interim request

About this task

Enable the RADIUS accounting interim request feature to create a log whenever more than the specified number of CLI commands are executed.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. In the **RADIUS Global** tab, type the number of CLI commands in the **CliCmdCount** field.
4. Click **Apply**.

RADIUS Global field descriptions

Use the data in the following table to use the **RADIUS Global** tab.

Name	Description
Enable	Enables the RADIUS authentication feature globally.
MaxNumberServer	Specifies the maximum number of servers to be used, between 1 and 10, inclusive.
AccessPriorityAttrValue	Specific to RADIUS authentication. Specifies the vendor-specific attribute value of the access-priority attribute to match the type value set in the dictionary file on the RADIUS server. The valid values are 192 through 240. Avaya recommends the default setting of 192 for Virtual Services Platform 9000.
AcctEnable	Enables RADIUS accounting.
AcctAttriValue	Specific to RADIUS accounting. Specifies the vendor-specific attribute value of the CLI-command attribute to match the type value set in the dictionary file on the RADIUS server. This value must be different from the access-priority attribute value configured for authentication. The valid values are 192 through 240. The default value is 193.
AcctIncludeCli	Specifies whether you want CLI commands included in RADIUS accounting requests.
ClearStat	Clears RADIUS statistics from the device.
McastAttributeValue	Specifies the value of the Mcast attribute. The valid values are 0 through 255. The default value is 90.
AuthInfoAttrValue	Specifies the value of the authentication information attribute. The valid values are 0 through 255. The default value is 91.
CommandAccessAttrValue	Specifies the value of the command access attribute. The valid values are 192 through 240. The default value is 194.
CliCommandAttrValue	Specifies the value of the CLI command attribute. The valid values are 192 through 240. The default value is 195.
AuthInvalidServerAddress	Displays the number of access responses from unknown or invalid RADIUS servers.
SourceIpFlag	Includes a configured IP address as the source address in RADIUS packets. The default is false. RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration.
CliCmdCount	Gives the value for the CLI command count. Specify an integer from 1 to 40. The default is 40.
CliProfEnable	Enables RADIUS CLI profiling.

Configuring the source IP option for the RADIUS server

Before you begin

- To configure the CLIP as the source IP address, you must configure the global RADIUS **sourceip-flag** parameter as true. You can configure the **source-ip** address parameter while you define the RADIUS Server on Virtual Services Platform 9000. The source IP address must be a CLIP address, and you can configure a different CLIP address for each RADIUS server.

For more information about configuring the source IP address, see Adding a RADIUS server on page 91.

Important:

Use the source IP option only for the RADIUS servers connected to the in-band network.

About this task

By default, Avaya Virtual Services Platform 9000 uses the IP address of the outgoing interface as the source IP and NAS IP address for RADIUS packets that it transmits. When you configure the RADIUS server, this IP address is used when defining the RADIUS Clients that communicate with it. Therefore, if the outgoing interface on Virtual Services Platform 9000 fails, a different source IP address is used—requiring that you make configuration changes to define the new RADIUS client on the RADIUS server.

To simplify RADIUS Server configuration, you can configure Virtual Services Platform 9000 to use a Circuitless IP Address (CLIP) as the source IP and NAS IP address when transmitting RADIUS packets. A CLIP is not associated with a physical interface and is always in an active and operational state. You can configure Virtual Services Platform 9000 with multiple CLIP interfaces.

RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration in all but the following case. When adding a RADIUS server in Enterprise Device Manager (EDM) or modifying a RADIUS configuration in EDM, you must specify if the address type is an IPv4 or an IPv6 address.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. In the **RADIUS Global** tab, select the **SourceIpFlag** check box.
4. Click **Apply**.

RADIUS Global field descriptions

Use the data in the following table to use the **RADIUS Global** tab.

Name	Description
Enable	Enables the RADIUS authentication feature globally.
MaxNumberServer	Specifies the maximum number of servers to be used, between 1 and 10, inclusive.
AccessPriorityAttrValue	Specific to RADIUS authentication. Specifies the vendor-specific attribute value of the access-priority attribute to match the type value set in the dictionary file on the RADIUS server. The valid values are 192 through 240. Avaya recommends the default setting of 192 for Virtual Services Platform 9000.
AcctEnable	Enables RADIUS accounting.
AcctAttriValue	Specific to RADIUS accounting. Specifies the vendor-specific attribute value of the CLI-command attribute to match the type

Table continues…

Name	Description
	value set in the dictionary file on the RADIUS server. This value must be different from the access-priority attribute value configured for authentication. The valid values are 192 through 240. The default value is 193.
AcctIncludeCli	Specifies whether you want CLI commands included in RADIUS accounting requests.
ClearStat	Clears RADIUS statistics from the device.
McastAttributeValue	Specifies the value of the Mcast attribute. The valid values are 0 through 255. The default value is 90.
AuthInfoAttrValue	Specifies the value of the authentication information attribute. The valid values are 0 through 255. The default value is 91.
CommandAccessAttrValue	Specifies the value of the command access attribute. The valid values are 192 through 240. The default value is 194.
CliCommandAttrValue	Specifies the value of the CLI command attribute. The valid values are 192 through 240. The default value is 195.
AuthInvalidServerAddress	Displays the number of access responses from unknown or invalid RADIUS servers.
SourceIpFlag	Includes a configured IP address as the source address in RADIUS packets. The default is false. RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration.
CliCmdCount	Gives the value for the CLI command count. Specify an integer from 1 to 40. The default is 40.
CliProfEnable	Enables RADIUS CLI profiling.

Adding a RADIUS server

About this task

Add a RADIUS server to allow RADIUS service on Avaya Virtual Services Platform 9000.

Remote Dial-In User Services (RADIUS) supports both IPv4 and IPv6 addresses, with no differences in functionality or configuration in all but the following case. When adding a RADIUS server or updating a RADIUS server in Enterprise Device Manager (EDM) you must specify if the address type is an IPv4 or an IPv6 address.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. Click the **RADIUS Servers** tab.
4. Click **Insert**.
5. In the **AddressType** box, select IPv4 or IPv6.

6. In the **Address** box, type the IP address of the RADIUS server that you want to add.
7. In the **UsedBy** box, select an option for the user logon.
8. In the **SecretKey** box, type a secret key.
9. In the **SourceIpAddr** box, type the IP address to use as the source address in RADIUS packets.
10. Click **Insert**.

RADIUS Servers field descriptions

Use the data in the following table to use the **RADIUS Servers** tab.

Name	Description
AddressType	Specifies either an IPv4 or an IPv6 address. RADIUS supports IPv4 and IPv6 addresses.
Address	Specifies the IP address of the RADIUS server. RADIUS supports IPv4 and IPv6 addresses.
UsedBy	Specifies the user logon: • cli: Specifies cli logon. • snmp: Specifies snmp logon. • eap: Specifies EAP PAE Authenticator. • web: Specifies HTTP(s) access authentication. The default is cli.
Priority	Specifies the priority of each server, or the order of servers to send authentication (1 to 10). The default is 10.
TimeOut	Specifies the time interval in seconds before the client retransmits the packet (1 to 20).
Enable	Enables or disables authentication on the server. The default is true.
MaxRetries	Specifies the maximum number of retransmissions allowed (1 to 6). The default is 1.
UdpPort	Specifies the UDP port that the client uses to send requests to the server (1 to 65536). The default value is 1812. The UDP port value set for the client must match the UDP value set for the RADIUS server.
SecretKey	Specifies the RADIUS server secret key, which is the password used by the client to be validated by the server.
AcctEnable	Enables or disable RADIUS accounting. The default is true.
AcctUdpPort	Specifies the UDP port of the RADIUS accounting server (1 to 65536). The default value is 1813. The UDP port value set for the client must match the UDP value set for the RADIUS server.

Table continues…

Name	Description
SourceIpAddr	Specifies the IP address to use as the source address in RADIUS packets. To use this option, you must set the global RADIUS SourceIpFlag to true. RADIUS supports IPv4 and IPv6 addresses.

Reauthenticating the RADIUS SNMP server session

About this task

Specify the number of challenges that you want the RADIUS SNMP server to send to authenticate a given session.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. Click the **RADIUS SNMP** tab.

   The RADIUS SNMP tab appears.
4. Select the **Enable** check box.
5. In the **ReauthenticateTimer** field, enter a value to specify the interval between RADIUS SNMP server reauthentications.

   The timer for reauthentication of the RADIUS SNMP server session is enabled.

   Important:

   To abort the RADIUS SNMP server session, enter a value for the AbortSessionTimer, and then click Enable.
6. Select the **AcctEnable** check box if desired.
7. Click **Apply**.

RADIUS SNMP field descriptions

Use the data in the following table to use the **RADIUS SNMP** tab.

Name	Description
Enable	Enables or disables timer authentication on the server. The default is true.
AbortSessionTImer	Specifies the allowable time, in seconds, before aborting the RADIUS SNMP server session (30 to 65535). The default is 180.
ReAuthenticateTimer	Specifies the time, in seconds, between reauthentications of the RADIUS SNMP server (30 to 65535). The default is 180.
AcctEnable	Enables or disables the RADIUS SNMP session timer.
UserName	Specifies the user name for the RADIUS SNMP accounting.

Configuring RADIUS SNMP

About this task

Configure RADIUS SNMP parameters for authentication and session times.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS.**
3. Select the **RADIUS SNMP** tab.
4. Select the **Enable** check box to enable RADIUS SNMP.
5. In the **AbortSessionTimer** field, enter the period after which the session expires in seconds.
6. In the **ReAuthenticateTimer** field, enter the period of time the system waits before reauthenticating in seconds.
7. Select the **AcctEnable** check box to enable RADIUS accounting for SNMP.
8. In the **UserName** field, type the RADIUS SNMP user name.
9. Click **Apply**.

RADIUS SNMP field descriptions

Use the data in the following table to use the **RADIUS SNMP** tab.

Name	Description
Enable	Enables or disables timer authentication on the server. The default is true.
AbortSessionTImer	Specifies the allowable time, in seconds, before aborting the RADIUS SNMP server session (30 to 65535). The default is 180.
ReAuthenticateTimer	Specifies the time, in seconds, between reauthentications of the RADIUS SNMP server (30 to 65535). The default is 180.
AcctEnable	Enables or disables the RADIUS SNMP session timer.
UserName	Specifies the user name for the RADIUS SNMP accounting.

Modifying a RADIUS configuration

About this task

Modify an existing RADIUS configuration or single function such as retransmissions and RADIUS accounting.

RADIUS supports IPv4 and IPv6 addresses with no difference in functionality or configuration in all except the following case. When modifying a RADIUS configuration in Enterprise Device Manager (EDM), you must specify if the address type is an IPv4 or an IPv6 address.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. Click the **RADIUS Servers** tab.
4. In the row and field to modify, type the information or use the lists to make a selection. Access the lists by double-clicking in a field.
5. When you are done with modifying the RADIUS configuration, click **Apply**.

RADIUS Servers field descriptions

Use the data in the following table to use the **RADIUS Servers** tab.

Name	Description
AddressType	Specifies either an IPv4 or an IPv6 address. RADIUS supports IPv4 and IPv6 addresses.
Address	Specifies the IP address of the RADIUS server. RADIUS supports IPv4 and IPv6 addresses.
UsedBy	Specifies the user logon: • cli: Specifies cli logon. • snmp: Specifies snmp logon. • eap: Specifies EAP PAE Authenticator. • web: Specifies HTTP(s) access authentication. The default is cli.
Priority	Specifies the priority of each server, or the order of servers to send authentication (1 to 10). The default is 10.
TimeOut	Specifies the time interval in seconds before the client retransmits the packet (1 to 20).
Enable	Enables or disables authentication on the server. The default is true.
MaxRetries	Specifies the maximum number of retransmissions allowed (1 to 6). The default is 1.
UdpPort	Specifies the UDP port that the client uses to send requests to the server (1 to 65536). The default value is 1812. The UDP port value set for the client must match the UDP value set for the RADIUS server.
SecretKey	Specifies the RADIUS server secret key, which is the password used by the client to be validated by the server.
AcctEnable	Enables or disable RADIUS accounting. The default is true.
AcctUdpPort	Specifies the UDP port of the RADIUS accounting server (1 to 65536). The default value is 1813.

Table continues…

Name	Description
	The UDP port value set for the client must match the UDP value set for the RADIUS server.
SourceIpAddr	Specifies the IP address to use as the source address in RADIUS packets. To use this option, you must set the global RADIUS SourceIpFlag to true. RADIUS supports IPv4 and IPv6 addresses.

Deleting a RADIUS configuration

About this task

Delete an existing RADIUS configuration.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **RADIUS**.
3. Click the **RADIUS Servers** tab.
4. Identify the configuration to delete by clicking anywhere in the row.
5. Click **Delete**.

RADIUS configuration examples

This section provides configuration examples to configure the Avaya Virtual Services Platform 9000 and Avaya Identity Engines Ignition Server to use RADIUS.

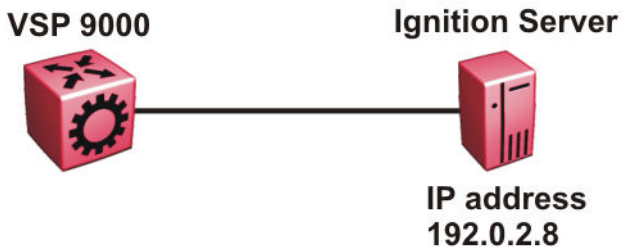

Figure 7: VSP 9000 connects to the Identity Engines Ignition Server

RADIUS configuration on VSP 9000

The following section shows the steps required to configure RADIUS on Avaya Virtual Services Platform 9000.

The example displays how to:

- Configure a key to be used by the RADIUS server and VSP 9000. In the example, the key is configured to the word `secret`.
- Configure an IP address for the RADIUS server. In the example the IP address is 192.0.2.8, which is accessible by the Management Router VRF.
- Configure the RADIUS server to authenticate ACLI and EDM sessions.
- Enable RADIUS.

VSP 9000

```
RADIUS CONFIGURATION

radius server host 192.0.2.8 key ****** used-by cli
radius server host 192.0.2.8 key ****** used-by web
radius enable
```

Verify your configuration

The output for the **`show radius`** command, must have a value for access-priority-attribute of 192 for Avaya Virtual Services Platform 9000 to access the RADIUS Identity Engines Ignition Server. The **`show radius`** output must show as `enable: true` to confirm RADIUS is enabled.

```
VSP-9012:1(config)#show radius
VSP-9012:1(config)#show radius
             acct-attribute-value : 193
                     acct-enable : false
        acct-include-cli-commands : false
        access-priority-attribute : 192
             auth-info-attr-value : 91
         command-access-attribute : 194
           cli-commands-attribute : 195
                    cli-cmd-count : 40
               cli-profile-enable : false
                           enable : true
                 igap-passwd-attr : standard
           igap-timeout-log-fsize : 512
                      maxserver : 10
            mcast-addr-attr-value : 90
                    sourceip-flag : false
```

The output for the **`show radius-server`** command must display the IP address for the RADIUS Identity Engines Ignition Server. The IP address must be accessible to the Management Router VRF on Avaya Virtual Services Platform 9000.

If you want to use the RADIUS server to authenticate sessions in ACLI, under USED BY, the following output must display as cli. If you want to authenticate EDM sessions, under USED BY, the following output must display web.

```
VSP-9012:1(config)#show radius-server
================================================================================
                               Radius Server Entries
```

```
================================================================================
                                                              ACCT
Name                USED                          TIME EN-  ACCT EN-   SOURE
                    BY  SECRET PORT PRIO RETRY OUT ABLED PORT ABLED IP
192.0.2.8           cli ****** 1812 10   1     3   true  1813 true  0.0.0.0
192.0.2.8           web ****** 1812 10   1     3   true  1813 true  0.0.0.0
```

Identity Engine Ignition Server configuration example

The following section shows the steps required to configure RADIUS on Avaya Identity Engines Ignition Server, Release 8.0. Use the preceding information to configure the Virtual Services Platform 9000.

A RADIUS server responds to and audits network access requests. In an Avaya installation, the Identity Engines Ignition Server is the RADIUS server.

The example displays how to do the following:

- Configure outbound attributes
- Create an outbound value
- Configure a user
- Configure the authentication protocol policy
- Create the authorization policy
- Configure RADIUS authenticators

For more information on the Avaya Ignition Server, see the *Avaya Identity Engines Ignition Server Administration*, NN47280–600.

Before you begin

- Configure the Ignition Server appliance and set up its network settings. For more information, see *Avaya Identity Engines Ignition Server Getting Started*, NN47280–300.
- Install Ignition Dashboard on your Windows OS.
- Configure each authenticator (VSP 9000) to recognize the Ignition Server appliance as its RADIUS server.
- Configure your switch to send its packets to the Ignition Server appliance with the appropriate IP address and port 1813 as the RADIUS account port, if you want to use RADIUS accounting.
- Make sure that client machines have VPN client that speaks PAP or MSCHAPv2, if you want to use IPsec for VPN access.
- Ensure licenses are up-to-date.

Procedure

1. If the Ignition Server Dashboard is not connected to your Ignition Server, select **Administration: Login** to connect.
    a. The default login credentials for **User Name** and **Password** are `admin`/`admin`. Avaya recommends you change the default values.

b. In the **Connect to** field enter the IP address of the Ignition Server for RADIUS. In this example, the IP address for the RADIUS server is 192.0.2.8.

2. Configure the outbound attributes to carry your provisioning values.

   a. In the Ignition Server Dashboard, expand the following in the Configuration tree: **Site Configuration** > **Provisioning** > **Outbound Attributes**.

   b. In the Outbound Attributes panel, click **New**.

   c. In the New Outbound Attribute window, type a name for the Ignition Server outbound attribute in the **Outbound Attribute** field. For this example use: `Access Priority-192`.

   d. Click **VSA**.

   e. In the **Vendor** drop down box, select **Bay-Networks**.

   f. In the **VSA** field, the default VSA is **ERS8xxx-AccessPriority**.

   g. Click **OK**. Your new attribute now appears in the list in the Outbound Attributes panel of Avaya Identity Engines.

3. Create an Outbound Value, based on the Outbound Attribute, and give it a priority of 6 to allow Read-Write-All access.

   a. In the Ignition Server Dashboard, expand the following in the Configuration tree: **Site Configuration** > **Provisioning** > **Outbound Values**.

   b. Click **New**.

   c. In the Outbound Value Details window, type an **Outbound Value Name** of `access-priority-attribute-192:RWA` for the outbound value. You later choose this name in your authorization policy to send this value.

   d. Click **New** to add a name-value pair and the Outbound Value Instance window appears.

   e. In the **Choose Global Outbound Attribute** drop down list, select the name of the outbound attribute to carry the value, which in this example is `Access-Priority-192`.

   f. Under **Value**, select **Unsigned-32 bit** and in the field to the right type `6`.

      By entering the value 6, you indicate a priority of 6, which is equal to read-write-all access.

   g. Click **OK**.

4. Configure a user recognized by the RADIUS server.

   a. In the Ignition Server Dashboard, expand the following in the Configuration tree: **Site Configuration** > **Directories** > **Internal Store** > **Internal Users**.

   b. Click **New**.

   c. Fill in the appropriate fields.

      As an example:

      User Name: jsmith

First Name: John

Last Name: Smith

Password: test

Confirm password: test

5. Create the Authentication Policy to allow the RADIUS Ignition Server to communicate with VSP 9000 with the password authentication protocol (PAP).
   a. In the Ignition Server Dashboard, expand the following in the Configuration tree: **Site Configuration** > **Access Policies**.
   b. Select **default-radius-user**.
   c. Select the **Authentication Policy** tab to the right.
   d. The Edit Authentication Policy window appears. Select **PAP**.
   e. Click **OK**.
6. Create the authorization rule to trigger the RADIUS Ignition Server to send this outbound value to the authenticator.
   a. In the Ignition Server Dashboard, expand the following in the Configuration tree: **Site Configuration** > **Access Policies** > **RADIUS**.
   b. Select **default-radius-user**.
   c. Select the **Authorization Policy** tab to the right.
   d. Click **Edit** and the **Edit Authorization Policy** window appears.
   e. In the **Rules** section, select **jsmith** and select **Add**.

      The **Rules** list at the left lets you browse and sort the rules in your policy. Use the up and down arrow buttons at the right to set the rule sequence, and click a rule name in the list to edit that rule. The Selected Rule Details section lets you edit the rule you have selected.
   f. In the Selected Rule Details section, under **Rule Name**, for this example, it reads `jsmith`.
   g. Select **Rule Enabled**.
   h. With jsmith selected in the Rules list, go to the buttons to the right of the **Constraint** list and click **New**.
   i. In the Constraint Details window that opens, create your constraint as follows. In the **Attribute Category** drop down menu, select **User**.
   j. Select **user-id**.
   k. Select **Static value**, type **jsmith**.
   l. Click **Add**.
   m. Click **OK**.

n. In the Selected Rule Details section, under **Action Provisioning (Outbound Values)**, select **Allow** and choose **access-priority-attribute 192:RWA**, which you previously configured.

For this example to function properly, the summary window must display:

IF User: user-id = jsmith THEN Allow

Send Outbound Values: access-priority-attribute-192:RWA

o. Click **OK**.

7. Configure the Ignition Server to connect to authenticators, which is Avaya Virtual Services Platform 9000:

a. In the Ignition Server Dashboard, expand the following folders: **Site Configuration** > **Authenticators** > **default** and the Authenticator Summary window appears.

b. Click **New**, and the Authenticator Details window appears.

c. For this example, type `VSP9000` under name.

d. To the right select **Enable Authenticator**.

e. Type the IP address for the VSP 9000, which is the authenticator. Use the primary CPU address or the management virtual address.

f. In the **Vendor** field, select **Nortel**.

g. In the **Device template** field, select **ers-switches-nortel**.

h. In the **RADIUS Shared Secret** field, type the key value you entered into VSP 9000. In this example, the key is the word `secret`.

To connect using RADIUS, you must use the shared secret for each device. In your switch documentation, the shared secret can also be referred to as a specific key string or an encryption string.

i. Select **Enable RADIUS Access**.

j. Under **Access Policy**, select **default-radius-user**.

k. Click **OK**.

Chapter 6: TACACS+

This chapter provides Terminal Access Controller Access Control Plus (TACACS+) concepts and procedures to complete TACACS+ configuration.

TACACS+ fundamentals

The switch supports the TACACS+ client. TACACS+ is a remote authentication protocol that provides centralized validation of users who attempt to gain access to a router or Network Access Server (NAS).

The TACACS+ feature is a client and server-based protocol that allows the switch to accept a user name and password and send a query to a TACACS+ authentication server, sometimes called a TACACS+ daemon. The TACACS+ server allows access or denies access based on the response by the client.

The TACACS+ feature facilitates the following services:

- Login authentication and authorization for ACLI access through rlogin, Secure Shell (SSH), Telnet, or serial port.
- Login authentication for web access through EDM.
- Command authorization for ACLI through rlogin, SSH, Telnet, or serial port.
- Accounting of ACLI through rlogin, SSH, Telnet, and serial port.

The following figure displays the basic layout of the switch and the TACACS+ server.

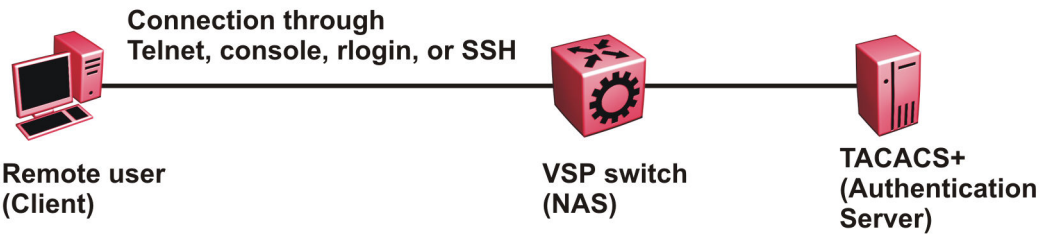

Figure 8: VSP switch and TACACS+ server

The TACACS+ feature uses Transmission Control Protocol (TCP) for its transport to ensure reliable delivery of packets. TACACS+ provides security by encrypting all traffic between the switch, which acts as the Network Access Server, and the TACACS+ server.

TACACS+ is a newer version of TACACS and provides separate authentication, authorization, and accounting (AAA) services. TACACS+ does not support earlier versions of TACACS.

TACACS+ is a base license feature. The TACACS+ feature is disabled by default.

TACACS+ Operation

The switch acts as an NAS to provide a connection to a single user, to a network, subnetwork or interconnected networks. The switch acts as a gateway to guard access to the TACACS+ server and network. Encryption relies on a secret key that is known to the client and the TACACS+ server.

Similar to the Remote Access Dial-In User Services (RADIUS) protocol, TACACS+ provides the ability to centrally manage the users who want to access a remote device. TACACS+ provides management of remote and local users who try to access a device through:

- rlogin
- Secure Shell (SSHv2)
- Telnet
- serial port
- Web management

A TACACS+ daemon, which typically runs on a UNIX or Windows NT workstation, maintains the TACACS+ authentication, authorization, and accounting services. Avaya Identity Engine supports the TACACS+ daemon. Avaya recommends you use the Avaya Identity Engines Ignition server as your TACACS+ server.

You configure users in the TACACS+ server. If you enable authentication, authorization, and accounting services, the following occurs:

- During the logon process, the TACACS+ client initiates the TACACS+ authentication session with the TACACS+ server.
- After successful authentication the TACACS+ client initiates the TACACS+ authorization session with the TACACS+ server. This is transparent to the user. The switch receives the user access level after a successful TACACS+ authorization. The TACACS+ server authorizes every command the user issues if TACACS + command authorization is enabled for that user access level.
- After successful authorization, if you enable TACACS+ accounting, the TACACS+ client sends accounting information to the TACACS+ server.

The TACACS+ feature on Virtual Services Platform 9000 supports partial High Availability (HA) implementation. With partial HA implementation, VSP 9000 synchronizes TACACS+ configuration parameters between the master and standby CP.

A TACACS+ session establishes with the server in one of two ways:

- Multi-connection mode (also known as per-session): For every authentication, authorization, and accounting (AAA) request the switch establishes a session with the TACACS+ server, and

then once the request finishes, the session is torn down. Multi-connection mode is the default mode.

- Single-connection mode: The first AAA request establishes the session, which is only torn down if TACACS+ is disabled or due to inactivity.

In both multi-connection mode and single-connection mode if failover occurs, the new master CP does not reestablish the session, which makes this partial HA.

For more information on High Availability-CPU mode, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

TACACS+ Architecture

You can connect the TACACS+ server to the VSP switch:

- In-band through one of the data ports.
- Out-of-band through the management port.

Connect the TACACS+ server through a local interface. Management PCs can reside on an out-of-band management Ethernet port, or on the corporate network. Place the TACACS+ server on the corporate network so you can route it to the switch.

Before you configure the switch, you must configure at least one TACACS+ server and a key.

The TACACS+ server and the switch must have the same:

- Encryption key
- Connection mode (single connection or per-session connection. Per-session connection is the same as multi-connection mode.)
- TCP port number

You can configure a secondary TACACS+ server for backup authentication. You specify the primary authentication server when you configure the switch.

Authentication, authorization, and accounting

A fundamental feature of TACACS+ is the separation of authentication, authorization, and accounting (AAA) services, which allows you to selectively implement one or more TACACS + services.

TACACS+ authentication

TACACS+ authentication provides control of authentication through login and password.

Authentication uses a database of users and passwords to determine:

- who a user is
- whether to allow the user access to the NAS

Important:

Prompts for log on and password occur prior to the authentication process. If TACACS+ fails because no valid servers exist, the device uses the user name and password from the local database. If TACACS+ or the local database returns an access denied packet, the authentication process stops. The device attempts no other authentication methods.

The following figure illustrates the authentication process.

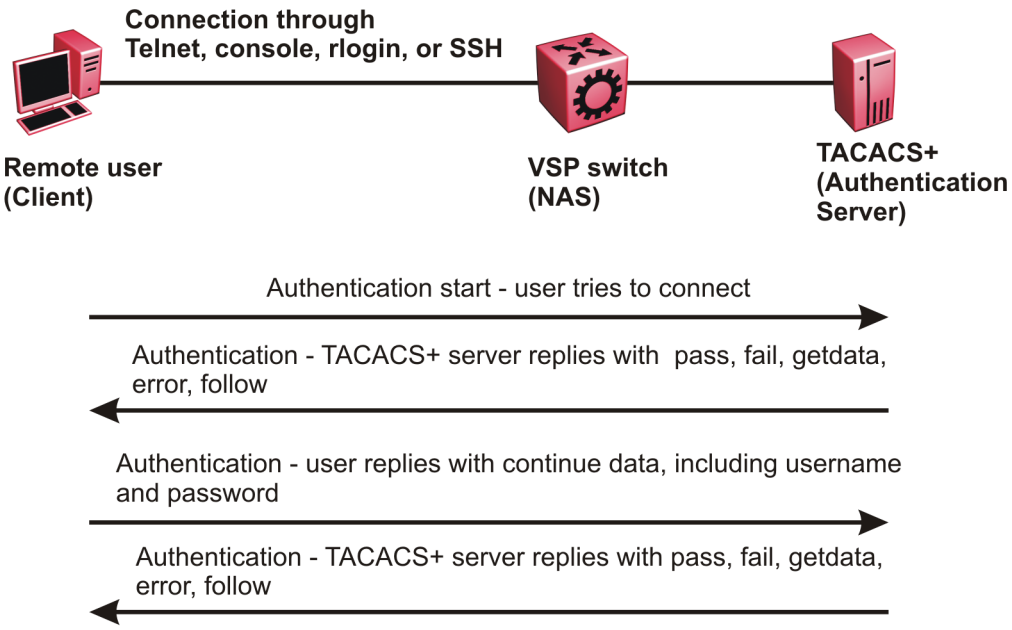

Figure 9: Authentication process

TACACS+ authorization

The transition from TACACS+ authentication to the authorization phase is transparent to the user. After successful completion of the authentication session, an authorization session starts with the authenticated user name. The authorization session provides access level functionality.

Authorization cannot occur without authentication.

Authorization:

- determines what a user can do
- allows administrators fine-grained control over the capabilities of users during sessions

The following figure illustrates the authorization process.

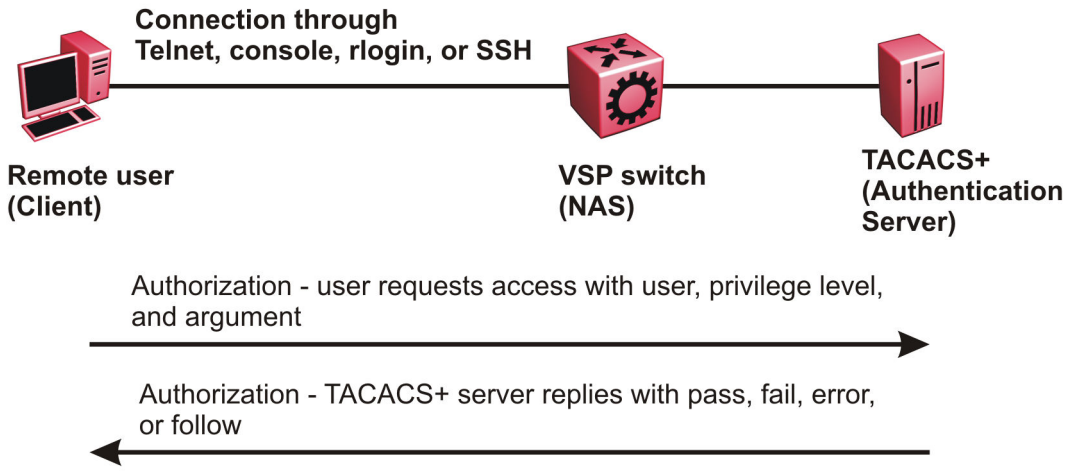

Figure 10: Authorization process

Authorization determines what a user can do. Authorization gives you the ability to limit network services to certain users and to limit the use of certain commands to certain users. The TACACS+ feature enhances the security by tightly policing the command execution for a particular user. After you enable command authorization, all commands, no matter the access level to which they belong, are sent to the TACACS+ server for authorization. Authorization cannot occur without first enabling authentication. You must configure command authorization globally and at individual access levels.

Two kinds of authorization requests exist:

1. Login authorization: Login authorization happens immediately after authentication and is transparent to the user. When the user logs on to the device, authorization provides the user access level. With log on, the device does not send a command to the TACACS+ server. You cannot configure login authorization.
2. Command authorization: When you configure command authorization for a particular level, all commands that you issue are sent to the TACACS+ server for authorization. The device can only issue the commands the TACACS+ server authorizes. You need to configure command authorization globally and at individual access levels, which are visible to the users.

Note:

You must verify that the switch can reach the TACACS+ server and that you configure TACACS + properly before you enable command authorization.

If a user is TACACS+ authenticated and command authorization is enabled for that level, then if the switch cannot reach the TACACS+ server, the switch does not allow the user to issue any

command that has privilege level command authorization enabled. In such a case, the user can only issue logout and exit commands.

If a user tries to log in and the TACACS+ server does not exist or is not reachable, then, as discussed before, a local database in the switch authenticates the user. The switch authorizes a locally authenticated user and a locally authenticated user is not eligible for TACACS+ command authorization.

After the switch requests authorization, the logon credentials are sent to the TACACS+ daemon for authorization. If logon authorization fails, the user receives a permission denied message.

If TACACS+ logon authorization succeeds, the switch uses information from the user profile, which exists in the local user database or on the TACACS+ server, to configure the session for the user.

After you enable TACACS+ command authorization all commands are visible to all users; however, the user can only issue those commands that the TACACS+ server configuration allows.

The switch cannot enforce command access level. The TACACS+ server returns an access level to the switch. The switch allows the user to access the switch according to the access level. The device grants the user access to a command only if the profile for the user allows the access level.

You preconfigure command authorization on the TACACS+ server. You specify a list of regular expressions that match command arguments, and you associate each command with an action to deny or permit.

All members in a group have the same authorization. If you place a user in a group, the daemon looks in the group for authorization parameters if it cannot find them in the user profile.

TACACS+ accounting

TACACS+ accounting enables you to track the services users access and the amount of network resources users consume.

TACACS+ accounting allows you to track:

- what a user does
- when a user does certain actions

The accounting record includes the following information:

- User name
- Date
- Start/stop/elapsed time
- Access server IP address
- Reason

You can use accounting for an audit trail, to bill for connection time or resources used, or for network management. TACACS+ accounting provides information about user sessions using the following connection types: Telnet, rlogin, SSH, and web-based management.

With separation of AAA, accounting can occur independently from authentication and authorization.

The following figure illustrates the accounting process.

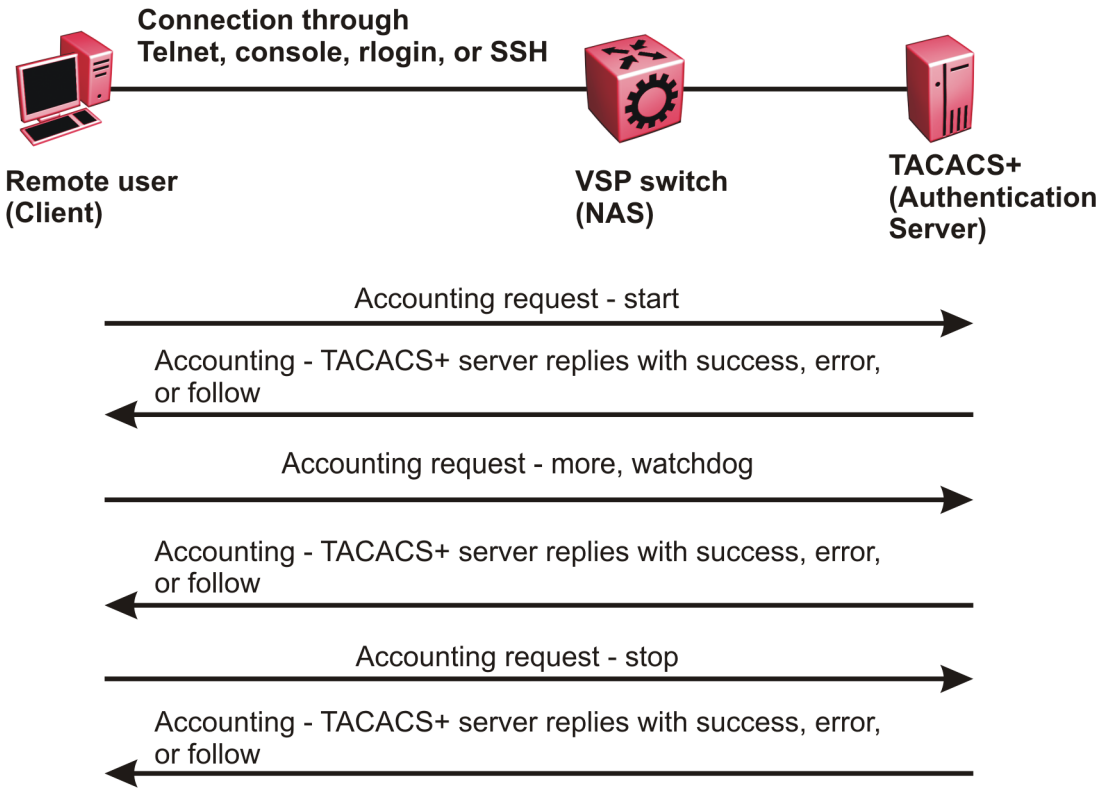

Figure 11: Accounting process

After you enable accounting, the switch reports user activity to the TACACS+ server in the form of accounting records. Each accounting record contains accounting attribute value (AV) pairs. AV pairs are strings of text in the form “attribute-value” sent between the switch and a TACACS+ daemon as part of the TACACS+ protocol. The TACACS+ server stores the accounting records.

You cannot customize the set of events the switch monitors and logs with TACACS+ accounting. TACACS+ accounting logs the following events:

- User logon and logoff
- Logoff generated because of activity timeout
- Unauthorized command
- Telnet session closed (not logged off)

Privilege level changes at runtime

You can change your privilege level at runtime with the **`tacacs switch level`** command.

You need to configure separate profiles in the TACACS+ server configuration file for the switch level. The VSP switch supports only levels 1 to 6 and level 15. The VSP switch uses the profile when you issue the command **tacacs switch level <1–15>**. As part of the profile, you specify a user name, level, and password. To preconfigure a dummy user for that level on the TACACS + daemon, the format of the user name for the dummy user is **$enab<n>$**, where <*n*> is the privilege level to which you want to allow access.

The following is an example of a TACACS+ server profile, which you configure on the TACACS + server:

```
user = $enab6$ {
member = level6
login = cleartext get-me-on-6
}
```

The following table maps user accounts to TACACS+ privilege level.

Switch access level	TACACS+ privilege level	Description
NONE	0	If the TACACS+ server returns an access level of 0, the user is denied access. You cannot log into the device if you have an access level of 0.
READ ONLY	1	Permits you to view only configuration and status information.
LAYER 1 READ WRITE	2	Permits you to view most of the switch configuration and status information and change physical port settings.
LAYER 2 READ WRITE	3	Permits you to view and change configuration and status information for Layer 2 (bridging and switching) functions.
LAYER 3 READ WRITE	4	Permits you to view and change configuration and status information for Layer 2 and Layer 3 (routing) functions.
READ WRITE	5	Permits you to view and change configuration and status information across the VSP switch. This level does not allow you to change security and password settings.
READ WRITE ALL	6	Permits you to have all the rights of read-write access and the ability to change security settings, including Avaya command line interface (ACLI) and web-based

Table continues…

Switch access level	TACACS+ privilege level	Description
		management user names and passwords, and the SNMP community strings.
NONE	7 to 14	If the TACACS+ server returns an access level of 7 to 14, the user is denied access. You cannot log into the device if you have an access level of 7 to 14.
READ WRITE ALL	15	Permits you to have all the rights of read-write access and the ability to change security settings, including Avaya command line interface (ACLI) and Web-based management user names and passwords, and the SNMP community strings. **Note:** Access level 15 is internally mapped to access level 6, which ensures consistency with other vendor implementations. The VSP switch does not differentiate between an access level of 6 and an access level of 15.

TACACS+ command authorization

After you enable TACACS+ authorization, the current privilege-level to command mapping on the VSP switch is no longer relevant because the TACACS+ server has complete responsibility for command authorization. TACACS+ authorization provides access to the system based on username, not based on privilege level.

After you enable TACACS+ command authorization for a particular privilege level, and a user with that privilege level logs on, the user can access commands based on his user name.

TACACS+ switch level and TACACS+ switch back commands

The user can only issue the **`tacacs switch level`** command after TACACS+ authenticates the user. Locally authenticated users, which means users authenticated only by the VSP switch and not by the TACACS+ server, cannot use the **`tacacs switch level`** command.

Consider a user, called X, with a privilege level of 4, who uses the **`tacacs switch level <1-15>`** command to change the privilege level from 4 to 6.

If user X successfully changes the switch level to 6, the user name changes from X to “$enab6$”, and the privilege level changes from 4 to 6. If TACACS+ command authorization is enabled for privilege level 6, then the TACACS+ server authorizes commands issued based on the rules defined for (dummy) user “$enab6$”.

If TACACS+ command authorization is not enabled for privilege level 6, then the VSP switch locally authorizes the user X based on the privilege level of the user.

The user can return to his previous privilege level using the **tacacs switch back** command. In the preceding scenario, if the user issues the **tacacs switch back** command, the user name changes for user X from “$enab6$” to X, and the privilege level changes from 6 to 4.

TACACS+ switch level supports up to eight levels, and TACACS+ switch level allows a user to switch level up to eight times from his original privilege level. The VSP switch stores all of the previous privilege levels in the same order in which the user switches levels. After switching eight times, if the user tries to switch a level the ninth time, the following error message displays:

```
Only allowed to switch level 8 times!
```

The user can switch back to his previous privilege levels using the **tacacs switch back** command. The **tacacs switch back** command switches back in the reverse order in which you issued the **tacacs switch level** command. Consider a user who switched levels from 4 to 5, and then to 6. If the user used the **tacacs switch back** command, the user first moves from 6 to 5, and then using the **tacacs switch back** command again moves from 5 to 4.

Note:

If you want to switch to a privilege level 'X' using **tacacs switch level <1-15>** command, you must create a user "$enabX$" on the TACACS+ server. X is the privilege level that you want to change.

TACACS+ switch level functionality:

The following table explains TACACS+ switch level functionality.

User logs in with	TACACS+ server available	Result
TACACS+ authentication	Yes	The user can issue the **tacacs switch level <1–15>** command.
Local authentication	No	The user cannot issue the **tacacs switch level <1–15>** command.
Local authentication	Yes	Even if a TACACS+ server becomes reachable, the user remains locally authenticated and cannot issue the **tacacs switch level <1–15>** command.

TACACS+ command authorization functionality:

The following table explains TACACS+ command authorization functionality.

User logs in with	Command authorization	Result
Local authentication	—	The VSP switch authorizes the user locally.

Table continues…

User logs in with	Command authorization	Result
TACACS+ authentication	Not enabled for the logged-in level.	The VSP switch authorizes the user locally. If the server connection is lost, the VSP switch authorizes the user locally.
TACACS+ authentication	Enabled for the logged-in level.	The TACACS+ server authorizes the user. If the server connection is lost, the user can only issue **`exit`** and **`logout`** commands.

Note:

A user who configures TACACS+ is locally authenticated and authorized by the VSP switch, so even after the user configures TACACS+, the VSP switch continues to locally authorize the user.

TACACS+ and RADIUS differences

TACACS+ and RADIUS are security protocols that you can use on network devices.

You can enable TACACS+ and RADIUS together. However, TACACS+ has a higher priority. If the TACACS+ server is not available the authentication is sent to RADIUS, if RADIUS is enabled. However, if TACACS+ authentication fails, then requests are not sent to RADIUS.

Following is a list of differences between TACACS+ and RADIUS.

TACACS+	RADIUS
Separates Authorization, Authentication and Accounting (AAA). As a result, you can selectively implement one or more TACACS+ services. With TACACS+ you can use different servers for each service.	Combines authentication and authorization.
Uses TCP. TCP is connection-oriented. TCP immediately indicates if a server crashes or is not running. TCP offers an acknowledgement that a request has been received.	Uses UDP. UDP is best-effort delivery. RADIUS uses re-transmit attempts and timeouts to make up for the support TCP has.
Encrypts the entire body of the packet, which includes the password and username.	Encrypts only the password from the client to the server.
Used for administrator access. Usually used for administrator access to network devices.	Used for subscriber access. Usually used to authenticate remote users to a network.
Can control which access level of commands a user or group can access.	Cannot control which access level of commands can be used.

TACACS+ feature limitations

The current implementation of TACACS+ does not support the following features:

- Point-to-Point Protocol (PPP) authentication and accounting
- IPv6 for TACACS+
- S/KEY (One Time Password) authentication
- PAP/CHAP/MSCHAP authentication methods
- The FOLLOW response of a TACACS+ server, in which the AAA services are redirected to another server. The response is interpreted as an authentication failure.
- User capability to change passwords at runtime over the network. The system administrator must change user passwords locally, on the server.
- TACACS+ command authorization when the user accesses the switch through EDM and SNMP.
- Restriction of command authorization for a specific kind of access. After you enable command authorization, command authorization applies for Telnet, SSH, rlogin, and serial-port access. You cannot restrict command authorization to just one kind of access.

If a user is TACACS+ authenticated and command authorization is enabled for that level, then if the switch cannot reach the TACACS+ server, the switch does not allow the user to execute any command that has privilege level command authorization enabled.

TACACS+ configuration using ACLI

Enabling TACACS+

Enable TACACS+ globally on the switch.

The switch supports the TACACS+ client. TACACS+ is a security application implemented as a client and server-based protocol that provides centralized validation of users who attempt to gain access to a router or network access server (the VSP switch).

By default, TACACS+ is disabled.

Before you begin

- You must have access to and you must configure a TACACS+ server before the TACACS+ features on your switch are available.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable TACACS+ globally:

   ```
   tacacs protocol enable
   ```

3. Disable TACACS+ globally:

   ```
   no tacacs protocol enable
   default tacacs protocol enable
   ```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs protocol enable
```

Adding a TACACS+ server

Add a primary and secondary TACACS+ server and specify the authentication process.

If you have a backup server configured, the AAA request goes to the backup server if the primary server is not available.

Avaya recommends you use the Avaya Identity Engines Ignition server as your TACACS+ server.

About this task

The TACACS+ server and the switch must have the same:

- Encryption key
- Connection mode (single connection or per-session connection. Per-session connection is the same as multi-connection mode)
- TCP port number

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Add a primary TACACS+ server with an encryption key:

   ```
   tacacs server host {A.B.C.D} key WORD<0-128>
   ```

3. **(Optional)** Configure the parameters for the primary TACACS+ server as required.
   a. **(Optional)** Specify a single connection. The single connection parameter maintains a constant connection between the switch and the TACACS+ daemon:

```
tacacs server host {A.B.C.D} single-connection
```

Note:

The TACACS+ daemon must also support this mode. If you do not configure this, the switch uses the default connection type, which is the per-session connection. Per-session is the same as multi-connection mode.

b. **(Optional)** Specify the TCP port to use when the switch connects to the TACACS+ daemon:

```
tacacs server host {A.B.C.D} port <0-65535>
```

The default port is 49.

c. **(Optional)** Specify the period of time (in seconds) the switch waits for a response from the TACACS+ daemon before it times out and shows an error:

```
tacacs server host {A.B.C.D} timeout <10-30>
```

d. **(Optional)** Designate a fixed source IP address for all outgoing TACACS+ packets and enable this option:

```
tacacs server host {A.B.C.D} source {A.B.C.D}source-ip-interface
enable
```

4. Specify the IP address of the secondary TACACS+ server and specify an encryption key:

```
tacacs server secondary-host {A.B.C.D} key WORD<0-128>
```

5. **(Optional)** Configure the optional parameters on the secondary TACACS+ server as required.

a. **(Optional)** Specify a single connection for the secondary TACACS+ server. The single connection parameter maintains a constant connection between the switch and the TACACS+ daemon:

```
tacacs server secondary-host {A.B.C.D} single-connection
```

Note:

The TACACS+ daemon must also support this mode. If you do not configure this, the switch uses the default connection type, which is the per-session connection. Per-session is the same as multi-connection mode.

b. **(Optional)** Specify the TCP port to use when the switch connects to the TACACS+ daemon:

```
tacacs server secondary-host {A.B.C.D} port <0-65535>
```

c. **(Optional)** Specify the period of time (in seconds) the switch waits for a response from the TACACS+ daemon before it times out and shows an error:

```
tacacs server secondary-host {A.B.C.D} timeout<10-30>
```

d. **(Optional)** Designate a fixed source IP address for all outgoing TACACS+ packets and enable this option:

```
tacacs server secondary-host {A.B.C.D} source {A.B.C.D} source-ip-interface enable
```

6. Display the status of the TACACS+ configuration:

   ```
   show tacacs
   ```

7. **(Optional)** Delete a primary TACACS+ server:

   ```
   no tacacs server host{A.B.C.D} [single-connection][source source-ip-interface enable]
   ```

8. **(Optional)** Delete a backup TACACS+ server:

   ```
   no tacacs server secondary-host{A.B.C.D} [single-connection][source source-ip-interface enable]
   ```

9. **(Optional)** Configure a primary TACACS+ server or secondary TACACS+ server to the default settings:

   ```
   default tacacs server {A.B.C.D} [port][single-connection][source source-ip-interface enable][timeout]
   ```

Example

Configure the primary server with the IP address 192.0.2.1 and the encryption key 1dt41y. Configure the secondary server with the IP address 198.51.100.2 with the same encryption key 1dt41y. Display the configuration to ensure proper configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs server host 192.0.2.1 key 1dt4ly
Switch:1(config)#tacacs server secondary-host 198.51.100.2 key 1dt4ly
Switch:1(config)#show tacacs

Global Status:

   global enable : true

   authentication enabled for : cli

   accounting enabled for : none

   authorization : disabled

   User privilege levels set for command authorization : None

Server:
                  create :

Prio      Status  Key         Port  IP address      Timeout Single Source
SourceEnabled
Primary   Conn    ******      49    192.0.2.1       10      false  0.0.0.0
false
Backup    NotConn ******      49    198.51.100.2    10      false  0.0.0.0
false


Switch:1(config)#no tacacs server host 192.0.2.1
Switch:1(config)#no tacacs server secondary-host 198.51.100.2
```

Variable definitions

Use the data in the following table to use the **tacacs server host** and the **tacacs server secondary-host** commands.

Variable	Value
{A.B.C.D}	Specifies the IP address of the TACACS+ server you want to add. For the current release, only IPv4 addresses are valid.
key *WORD <0-128>*	Configures the authentication and encryption key for all TACACS+ communications between the device and the TACACS+ server. If the key length is zero, that indicates no encryption is used. You must configure the same encryption key for the TACACS+ server and the switch.
port *<0-65535>*	Configures the TCP port, on which the client establishes a connection to the server. A value of 0 indicates the system specified default value is used. The default is 49. You must configure the same TCP port for the TACACS+ server and the switch.
single-connection	Specifies if the TCP connection between the device and the TACACS+ server is a single connection. If you specify the single connection parameter, the connection between the switch and the TACACS+ daemon remains open, which is more efficient because it allows the daemon to handle a higher number of TACACS+ operations. The single-connection is torn down if TACACS+ is disabled due to inactivity. If you do not configure this, the switch uses the default connection type, which is the multi-connection. With the multi-connection, the connection opens and closes each time the switch and TACACS+ daemon communicate. **Note:** You must configure the same connection mode for the TACACS+ server and the switch. To enable single-connection, the TACACS+ daemon has to support this mode as well.
source *{A.B.C.D}*	Designates a fixed source IP address for all outgoing TACACS+ packets, which is useful if the router has many interfaces and you want to make sure all

Table continues…

Variable	Value
	TACACS+ packets from a certain router have the same IP address. If you do not configure an address, the system uses 0.0.0.0 as the default. For the current release, only IPv4 addresses are valid. **Note:** If you configure a valid source IP address that is not 0.0.0.0 without enabling source-ip-interface, the source IP address returns to 0.0.0.0.
source-ip-interface enable	Enables the source address. You must enable this parameter if you configure a valid source IP address. The default is disabled.
timeout *<10-30>*	Configures the maximum time, in seconds, to wait for this TACACS+ server to reply before it times out. The default value is 10 seconds.

Job aid

The following table describes the fields in the output for the **`show tacacs`** command.

Name	Description
`Global Status`	
`global enable`	Displays if the TACACS+ feature is enabled globally.
`authentication enabled for`	Displays which application is authenticated by TACACS+. The possibilities are ACLI, web, or all.
`accounting enabled for`	Displays if accounting is enabled. You can only enable accounting for ACLI. By default, accounting is not enabled.
`authorization`	Displays if authorization is enabled.
`User privilege levels set for command authorization`	Displays the privilege levels set for command authorization. When you configure command authorization for a particular level, all commands that you execute are sent to the TACACS+ server for authorization. The device can only execute the commands the TACACS+ server authorizes. The user privilege levels are: • 0: denied access • 1: read only (ro) access • 2: Layer 1 read and write (l1) access • 3: Layer 2 read and write (l2) access

Table continues…

Name	Description
	• 4: Layer 3 read and write (l3) access • 5: read and write (rw) access • 6: read and write all (rwa) access • 7-14: denied access • 15: read and write all (rwa) access
Server	
Prio	Displays the priority of the TACACS+ server. The switch attempts to use the primary server first, and the secondary server second.
Status	Displays the connection status between the server and the switch – connected or not connected.
Key	Displays as ****** instead of the actual key. The key is secret and is not visible.
Port	Displays the TCP port used to establish the connection to the server. The default port is 49.
IP address	Displays the IP address for the primary and secondary TACACS+ servers.
Timeout	Displays the period of time, in seconds, the switch waits for a response from the TACACS+ daemon before it times out and declares an error. The default is 10 seconds.
Single	Displays if a single open connection is maintained between the switch and TACACS+ daemon, or if the switch opens and closes the TCP connection to the TACACS+ daemon each time they communicate. The default is false, which means the device does not maintain the single open connection.
Source	Displays the fixed source IP address, if you configure one, for all outgoing TACACS+ packets.
SourceEnabled	Displays if the fixed source IP address is enabled for all outgoing TACACS+ packets.

Configuring TACACS+ authentication

Configure what application TACACS+ authenticates: ACLI, web, or all.

TACACS+ authentication provides control of authentication through login and password.

By default, CLI authentication is enabled.

Before you begin

- You must enable TACACS+ globally for TACACS+ authentication to function.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure TACACS+ authentication:

   ```
   tacacs authentication <all/cli/web>
   ```

3. **(Optional)** Disable TACACS+ authentication:

   ```
   no tacacs authentication <all/web>
   ```

4. **(Optional)** Configure TACACS+ authentication to the default settings (default is cli authentication enabled):

   ```
   default tacacs authentication <all/cli/web>
   ```

5. Display the configuration:

   ```
   show tacacs
   ```

Example

Configure TACACS+ to authenticate ACLI and display the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs authentication cli
Switch:1(config)#show tacacs
Global Status:

   global enable : true

   authentication enabled for : cli

   accounting enabled for : none


Server:
                  create :

Prio     Status  Key     Port  IP address  Timeout  SingleSource Source  Enabled
Primary  Conn    ******  49    192.0.2.1       10   false        0.0.0.0 false
Backup   NotConn ******  49    198.51.100.2    10   false        0.0.0.0 false
```

Variable definitions

Use the data in the following table to use the **`tacacs authentication`** command.

Variable	Value
all	Specifies TACACS+ authentication for all applications. By default, CLI authentication is enabled.

Table continues…

Variable	Value
cli	Specifies TACACS+ authentication for command line connections. By default, CLI authentication is enabled.
web	Specifies TACACS+ authentication for web connections. By default, CLI authentication is enabled.

Configuring TACACS+ accounting

Determines for which applications TACACS+ collects accounting information. Use TACACS+ accounting to track the services that users access and the amount of network resources that users consume. If unassigned, TACACS+ does not perform the accounting function.

If enabled, TACACS+ accounting logs the following events:

- User log on and log off
- Log off generated because of activity timeout
- Unauthorized command
- Telnet session closed (not logged off)

If unassigned, TACACS+ does not perform the accounting function. No default value exists.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable TACACS+ accounting:

   ```
   tacacs accounting enable <cli>
   ```

3. **(Optional)** Disable TACACS+ accounting:

   ```
   no tacacs accounting enable <cli>
   ```

Example

Enable TACACS+ accounting for ACLI:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs accounting enable cli
```

Configuring command authorization with TACACS+

Use this procedure to enable TACACS+ authorization for a particular privilege level. Use this option to limit the use of certain commands to certain users.

If command authorization fails, the following log message displays: `Command <command> not authorized for user <username>.`

By default, command authorization is disabled on the switch. The default for the command authorization level is none.

Before you begin

- You must have access to and you must configure a TACACS+ server before the TACACS+ features on your switch are available. You must verify that the switch can reach the TACACS+ server and that you configure TACACS+ properly before you enable command authorization. If a user is TACACS+ authenticated and command authorization is enabled for that level, then if the switch cannot reach the TACACS+ server, the switch does not allow you to issue any command that has privilege level command authorization enabled. If the switch cannot reach the TACACS+ server, you can only issue logout and exit commands.
- To use TACACS+ authorization, you must enable TACACS+ authentication.

About this task

Two kinds of authorization requests exist:

1. Login authorization: Login authorization happens immediately after authentication when the user logs on to the device, authorization provides the user access level. You cannot configure login authorization.
2. Command authorization: When you configure command authorization for a particular level, all commands that you issue are sent to the TACACS+ server for authorization. You need to configure command authorization globally and at individual access levels.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Enable TACACS+ authorization:

   ```
   tacacs authorization enable
   ```

3. Configure TACACS+ privilege level for TACACS+ command authorization:

   ```
   tacacs authorization level <1-6>
   tacacs authorization level all
   tacacs authorization level none
   ```

4. **(Optional)** Disable TACACS+ authorization:

   ```
   tacacs authorization disable
   default tacacs authorization
   ```

Example

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs authorization enable
Switch:1(config)#tacacs authorization level 6
```

Variable definitions

Use the data in the following table to use the **tacacs authorization** command.

Variable	Value
level <*1–6*>	Enables command authorization for a specific privilege level. The default for the command authorization level is none.
level all	Enables command authorization for all privilege levels. The default for the command authorization level is none.
level none	Disables command authorization for all privilege levels. The default for the command authorization level is none.

Changing privilege levels at runtime

Users can change their privilege levels at runtime. The privilege level determines what commands a user can access through TACACS+ server authorization.

A user can only use the **tacacs switch level** command, after TACACS+ authenticates the user. Locally authenticated users, which means users authenticated only by the switch and not by the TACACS+ server, cannot use the **tacacs switch level** command.

Before you begin

- You need to configure separate profiles in the TACACS+ server configuration file for switch level. As part of the profile, you specify a user name, level, and password.

About this task

After you enable TACACS+ authorization, the current privilege-level to command mapping on the switch is no longer relevant because the TACACS+ server has complete responsibility for command authorization. TACACS+ authorization provides access to the system based on username, not based on privilege level.

After you enable TACACS+ command authorization for a particular privilege level, and a user with that privilege level logs on, the user can access commands based on his user name.

Note:

If you want to switch to a privilege level 'X' using **tacacs switch level <1-15>** command, you must create a user "$enabX$" on the TACACS+ server. X is the privilege level to which you want to change.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Change the privilege level for a user at runtime:

   ```
   tacacs switch level <1–15>
   ```

3. Return to the original privilege level:

   ```
   tacacs switch back
   ```

Example

Change the privilege level for a user at runtime. Return to the original privilege level:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#tacacs protocol enable
Switch:1(config)#tacacs switch level 5
Password:******
```

Return to the original privilege level:

```
Switch:1(config)#tacacs switch back
```

Variable definitions

Use the data in the following table to use the **`tacacs switch`** command.

Variable	Value
level <*1–15*>	Specifies the privilege level you want to access. You can change your privilege level at runtime by using this parameter. You are prompted to provide the required password. If you do not specify a level in the command, the administration level is selected by default. **Note:** For switch level, you need to configure separate profiles in the TACACS+ server configuration file. As part of the profile, you specify a username, level, and password. To preconfigure a dummy user for that level on the TACACS+ daemon, the format of the username for the dummy user is **`$enab<n>$`**, where <*n*> is the privilege level to which you want to allow access.
back	Specifies that you want to return to the original privilege level.

TACACS+ configuration using EDM

Configuring TACACS+ globally

Enable TACACS+ globally on the switch. TACACS+ is a security application implemented as a client and server-based protocol that provides centralized validation of users. By default, TACACS+ is disabled.

Before you begin

- You must have access to and you must configure a TACACS+ server before the TACACS+ features on your switch (network access server) are available.

  You must verify that the switch can reach the TACACS+ server and that you configure TACACS+ properly before you enable command authorization.
- If a user is TACACS+ authenticated and command authorization is enabled for that level, then if the switch cannot reach the TACACS+ server, the switch does not allow the user to issue any command that has privilege level command authorization enabled. In such a case, the user can only issue logout and exit commands.
- You must enable TACACS+ globally for TACACS+ authentication to function.
- You must enable TACACS+ authentication for TACACS+ authorization to function.

About this task

Configure what application TACACS+ authenticates. TACACS+ authentication provides control of authentication through login and password dialog, challenge and response. By default, CLI authentication is enabled.

After authentication is complete, the switch starts the authorization process. By default, command authorization is disabled on the switch. The default for the command authorization level is none. If command authorization fails, the following log message displays: `Command <command> not authorized for user <username>.`

Two kinds of authorization requests exist:

1. Login authorization: Login authorization happens immediately after authentication when the user logs on to the device, authorization provides the user access level. You cannot configure login authorization.
2. Command authorization: When you configure command authorization for a particular level, all commands that you issue are sent to the TACACS+ server for authorization. You need to configure command authorization globally and at individual access levels.

Enable TACACS+ accounting function and determine which application TACACS+ accounts. After you enable accounting, the switch reports user activity to the TACACS+ server in the form of accounting records. The default for accounting is none.

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **TACACS+**.

3. Click the **TACACS+ Globals** tab.
4. Select the **GlobalEnable** check box to enable TACACS+ globally.
5. Select the **cli** check box to enable the **Accounting** option.
6. Select the **cli** or **web** check box to enable the **Authentication** option.
7. Click the **AcliCommandAuthorizationEnabled** box to enable TACACS+ authorization.
8. Select the level in the **AcliCommandAuthorizationLevels** box.
9. Click **Apply**.

TACACS+ Globals field descriptions

Use the data in the following table to use the **TACACS+ Globals** tab.

Name	Description
GlobalEnable	Enables or disables the TACACS+ feature globally.
Accounting	Determines for which applications TACACS+ collects accounting information. Use TACACS+ accounting to track the services that users access and the amount of network resources that users consume. If unassigned, TACACS+ does not perform the accounting function. The default is none. If enabled, TACACS+ accounting logs the following events: • User log on and log off • Log off generated because of activity timeout • Unauthorized command • Telnet session closed (not logged off)
Authentication	Configures what application TACACS+ authenticates. The options include: • cli • web TACACS + authentication provides control of authentication through login and password dialog, challenge and response. By default, CLI authentication is enabled.
LastUserName	Displays the last user for which the system attempted authentication.
LastAddressType	Displays the type of address to access the TACACS + server.
LastAddress	Displays the last address to access the TACACS+ server.

Table continues…

Name	Description
AcliCommandAuthorizationEnabled	Enables TACACS+ authorization for a particular privilege level. Use this option to limit the use of certain commands to certain users. To use TACACS + authorization, you must also use TACACS+ authentication. The VSP switch allows the user to access the switch according to the access level. The default is disabled.
AcliCommandAuthorizationLevels	Enables command authorization for a specific privilege level. The default for the command authorization level is none.

Adding a TACACS+ server

Add a TACACS+ server, configure the TACACS+ server, and specify the authentication process.

If you have a secondary server configured, the AAA request goes to the backup server if the primary server is not available.

Avaya recommends you use the Avaya Identity Engines Ignition server as your TACACS+ server.

Before you begin

You must have access to and you must configure a TACACS+ server before the TACACS+ features on your switch are available.

About this task

The TACACS+ server and the switch must have the same:

- Encryption key
- Connection mode (single connection or per-session connection. Per-session is the same as multi-connection mode.)
- TCP port number

Procedure

1. In the navigation pane, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **TACACS+**.
3. Click the **TACACS+ Servers** tab.
4. Click **Insert**.
5. In the **AddressType** box, select **ipv4**.
6. In the **Address** field, type the IP address of the TACACS+ server.

7. **(Optional)** In the **PortNumber** field, type the TCP port on which the client establishes a connection to the TACACS+ server.
8. **(Optional)** In the **ConnectionType** box, select either **singleConnection** or **perSessionConnection** to specify the TCP connection type between the switch and TACACS+ server.
9. **(Optional)** In the **Timeout** field, type the period of time (in seconds) the switch waits for a response from the TACACS+ server.
10. In the **Key** field, enter the key that the switch and the TACACS+ server share.
11. **(Optional)** Select **SourceIpInterfaceEnabled**, if you want to enable the switch to designate a fixed source IP address for all outgoing TACACS+ packets.
12. In the **SourceIPInterfaceType** box, select **ipv4**.
13. **(Optional)** In the **SourceIpInterface** field, type a fixed source IP address if you want to designate a fixed source IP address for all outgoing TACACS+ packets.
14. In the **Priority** box, select either **primary** or **backup** to determine the order the switch uses the TACACS+ servers.
15. Click **Insert**.
16. If you want to delete an existing TACACS+ configuration perform the following procedure. In the navigation pane, expand the following folders: **Configuration** > **Security** > **Control Path**.
17. Click **TACACS+**.
18. In the **TACACS+** tab, click **TACACS+ Servers** tab.
19. Identify the configuration to delete by clicking anywhere in the row.
20. Click **Delete**.

TACACS+ Servers field descriptions

Use the data in the following table to use the **TACACS+ Servers** tab.

Name	Description
AddressType	Specifies the type of IP address to use on the TACACS+ server. For the current release, you must set the value to IPv4.
Address	Specifies the IP address of the TACACS+ server.
PortNumber	Configures the TCP port on which the client establishes a connection to the server. The default is 49. A value of 0 indicates that the system specified default value is used. You must configure the same TCP port for the TACACS+ server and the switch.

Table continues…

Name	Description
ConnectionType	Specifies if the TCP connection between the device and the TACACS+ server is a single connection. If you specify the single connection parameter, the connection between the switch and the TACACS+ daemon remains open, which is more efficient because it allows the daemon to handle a higher number of TACACS+ operations. The single-connection session is torn down if TACACS+ is disabled due to inactivity. If you do not configure this parameter, the switch uses the default connection type, which is the multi-connection. With the multi-connection, the connection opens and closes each time the switch and TACACS+ daemon communicate. **Note:** You must configure the same connection mode for the TACACS+ server and the switch. To enable single-connection, the TACACS+ daemon has to support this mode as well.
ConnectionStatus	Specifies if the TCP connection between the device and TACACS+ server is connected or not connected.
Timeout	Configures the maximum time, in seconds, to wait for this TACACS+ server to reply before it times out. The default value is 10 seconds.
Key	Configures the authentication and encryption key for all TACACS+ communications between the device and the TACACS+ server. If the key length is zero, that indicates no encryption is used. You must configure the same encryption key for the TACACS+ server and the switch.
SourceIpInterfaceEnabled	Enables the source address specification. If **SourceIpInterfaceEnabled** is true (the check box is selected), and you change **SourceIpInterfaceEnabled** to false (the check box is cleared), the **SourceIpInterface** is reset to 0.0.0.0. The default is disabled. You must enable this parameter if you configure a valid source IP address
SourceIpInterfaceType	Specifies the type of IP address to use on the interface that connects to the TACACS+ server.

Table continues…

Name	Description
	Note: For the current software release, you must set the value to IPv4.
SourceIpInterface	Designates a fixed source IP address for all outgoing TACACS+ packets, which is useful if the router has many interfaces and you want to make sure all TACACS+ packets from a certain router have the same IP address. If you do not configure an address, the system uses 0.0.0.0 as the default. For the current release, only IPv4 addresses are valid. **Note:** If you configure a valid source IP address that is not 0.0.0.0 without enabling source-ip-interface, the source IP address returns to 0.0.0.0.
Priority	Determines the order in which the switch uses the TACACS+ servers, where 1 is the highest priority. The priority values are primary and backup. If more than one server shares the same priority, the device uses the servers in the order they exist in the table.

Modifying a TACACS+ configuration

Modify an existing TACACS+ configuration to customize the server.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.
2. Click **TACACS+**.
3. Click **TACACS+ Servers** tab.
4. Double-click in the fields that you want to modify.

   In some of the fields, the text becomes bold, which indicates that you can edit them. In other fields, a list appears.
5. In the fields that you can edit, type the desired values.
6. In the fields with lists, select the desired option.
7. Click **Apply**.

TACACS+ configuration examples

This section provides configuration examples to configure the Avaya Virtual Services Platform switch and Avaya Identity Engines Ignition Server to use TACACS+.

Figure 12: VSP switch connects to the Identity Engines Ignition Server

TACACS+ configuration on the VSP switch

The following section shows the steps required to configure TACACS+ on the switch.

The example displays how to:

- Configure a key to be used by the TACACS+ server and the switch. In the example, the key is configured to the word `secret`.
- Configure an IP address for the TACACS+ server. In the example the IP address for the primary server is 192.0.2.8, which is accessible by the Management Router VRF.
- Configure the TACACS+ server to authenticate ACLI sessions.
- Enable TACACS+.

VSP switch

```
TACACS CONFIGURATION

tacacs server  host 192.0.2.8 key ******
tacacs protocol enable
tacacs accounting enable cli
tacacs authorization enable
tacacs authorization level 6
```

Verify your configuration

The **`show tacacs`** output must show as `global enable: true` to confirm TACACS is enabled.

The output for the **`show tacacs`** command must display the IP addresses for the TACACS+ Identity Engines Ignition Server. The IP addresses must be accessible to the Management Router VRF on the switch.

If you want to use the TACACS+ server to authenticate sessions in ACLI, the output must display as `authentication enabled for: cli`. If you want to authenticate EDM sessions, the output must display as `authentication enabled for: web`.

Ensure the other parameters match what you have configured.

```
Global Status:

   global enable : true

   authentication enabled for : cli

   accounting enabled for : cli

   authorization : enabled

   User privilege levels set for command authorization : rwa

Server:
                  create :

Prio      Status  Key         Port  IP address      Timeout Single Source
SourceEnabled
Primary   Conn    ******      49    192.0.2.8       10      false  0.0.0.0
false
```

Identity Engine Ignition Server TACACS+ configuration example

The following section shows the steps required to configure TACACS+ on Avaya Identity Engines Ignition Server, Release 8.0. Use the preceding information to configure the VSP switch.

A TACACS+ server responds to and audits network access requests. In an Avaya installation, the Identity Engines Ignition Server is the TACACS+ server.

The example displays how to do the following:

- Enable TACACS+
- Configure a user
- Create a command set
- Configure the authentication protocol policy
- Create the authorization policy
- Configure TACACS+ authenticators

For more information on the Avaya Ignition Server, see *Avaya Identity Engines Ignition Server Administration*, NN47280–600.

Before you begin

- Configure the Ignition Server appliance and set up its network settings. For more information, see *Avaya Identity Engines Ignition Server Getting Started*, NN47280–300.
- Install the Ignition Dashboard on your Windows OS.
- Configure each authenticator (VSP switch) to recognize the Ignition Server appliance as its TACACS+ server.
- Configure your switch to send packets to the Ignition Server appliance with the appropriate IP address and port.
- Ensure licenses are up-to-date.

Procedure

1. If the Ignition Server Dashboard is not connected to your Ignition Server, select **Administration: Login** to connect.
   a. The default login credentials for **User Name** and **Password** are `admin`/`admin`. Avaya recommends you change the default values.
   b. In the **Connect to** field enter the IP address of the Ignition Server for TACACS+. In this example, the IP address for the TACACS+ server is 192.0.2.8.
2. Enable TACACS+.
   a. In the Ignition Server Dashboard, select **Site 0**.
   b. In the Sites window, select the **Services** tab.
   c. Under the Services tab, select the **TACACS+** tab.
   d. Click the **Edit** button in the TACACS+ tab.
   e. In the **Edit TACACS+ Configuration** dialog box, select the **Protocol is enabled** box.
   f. In the **Bound Interface** field, select **Admin Port**.
   g. In the **Port** field, enter `49`.
   h. Select **Accept Requests from Any Authenticator**.

      Select this option if you want to create a global TACACS+ authenticator that sets policy for all authenticators that do not match a specific TACACS+-enabled authentication in your Ignition server configuration.
   i. In the **Access Policy** field, select **default-tacacs-admin**.

      Use this configuration in the case of a global TACACS+ authenticator. Choose your global TACACS+ policy that you want applied if the device finds no better matching authenticator.
   j. In **TACACS+ Shared Secret** field, enter the secret that the VSP switch and TACACS+ Ignition server share. In this example, the shared secret is `secret`.
   k. Click **OK**.
3. Configure a user recognized by the TACACS + server.
   a. In the Ignition Server Dashboard, expand the following in the Configuration tree: **Site Configuration** > **Directories** > **Internal Store** > **Internal Users**.
   b. Click **New**.
   c. Fill in the appropriate fields.

      As an example:

      User Name: jsmith

      First Name: John

      Last Name: Smith

Password: test

Confirm password: test

4. If your TACACS+ policy uses per-command authorization, create a command set.
   a. In the Ignition Server Dashboard, expand the following in the Configuration tree: **Site Configuration** > **Access Policies** > **TACACS+**.
   b. Click **Define Command Sets**.
   c. Click **New**.
   d. In the New Device Command Set window, type a **Name** and **Description** for the command set; for instance, level5.

      In this window you build your command set by adding commands to the list. You can build the command list manually or you can import a list. For more information on importing a command list, see *Avaya Identity Engines Ignition Server Administration*, NN47280–600.
   e. To manually add the commands, click **Add** in the New/Edit Device Command Set window.
   f. Click the **Simple Command Using Keywords and Arguments** box.
   g. In the **Command** field, type the command, and optionally its arguments.
   h. To allow the command to be used with any argument, select the **Allow** box.
   i. To allow only the specific command and arguments you have types, tick the **Deny** box.
   j. Click **OK** to add the command to the list.
   k. Continue to add the commands that you want.
5. If your TACACS+ policy uses privilege-level authorization, create the TACACS+ access policy to allow the TACACS+ Ignition Server to communicate with the VSP switch.
   a. In the Ignition Server Dashboard, expand the following in the Configuration tree: **Site Configuration** > **Access Policies** > **TACACS+**.
   b. Select **default-tacacs-admin**.
   c. Click on the **Authorization Policy** tab and select the name of the policy you want to edit.
   d. Click **Edit** and the **Edit Authorization Policy** window appears.
   e. In the **Rules** section, select the rule you want to edit. In this case select level5, to which you have already added commands.

      The **Rules** list at the left lets you browse and sort the rules in your policy. Use the up and down arrow buttons at the right to set the rule sequence, and click a rule name in the list to edit that rule. The Selected Rule Details section lets you edit the rule you have selected.
   f. In the Selected Rule Details section, under **Rule Name**, for this example, it reads `level5`.

g. Select **Rule Enabled**.

h. With level5 selected in the Rules list, go to the buttons to the right of the **Constraint** list and click **New**.

i. In the Action section, select **Allow**.

j. Select the **Command Sets** tab, in the Action section. Allow Commands in Set should read level-5, in this example, and under All Command Sets all the commands that are accessible under level5 should be listed.

k. Click **OK**.

For this example to function properly, the summary window must display:

IF User: user-id = level5 THEN Allow

Permit commands in Command Set: level-5

6. Configure the Ignition Server to connect to authenticators, which is the VSP switch:

a. In the Ignition Server Dashboard, expand the following folders: **Site Configuration** > **Authenticators** > **default** and the Authenticator Summary window appears.

b. Click **New**, and the Authenticator Details window appears.

c. For this example, type `VSPswitch` under name.

d. To the right select **Enable Authenticator**.

e. Type the IP address for the VSP switch, which is the authenticator. Use the primary CPU address or the management virtual address.

f. In the **Vendor** field, select **Nortel**.

g. In the **Device template** field, select **ers-switches-nortel**.

h. Select the **TACACS+ Settings** tab.

i. Select **Enable TACACS+ Access**.

j. In the **TACACS+ Shared Secret** field, type the key value you entered into VSP 4000. In this example, the key is the word `secret`.

To connect using TACACS+, you must use the shared secret for each device. In your switch documentation, the shared secret can also be referred to as a specific key string or an encryption string.

k. Under **Access Policy**, select **default-tacacs-user**.

l. Click **OK**.

Chapter 7: Simple Network Management Protocol (SNMP)

You can use the Simple Network Management Protocol (SNMP) to remotely collect management data and configure devices.

An SNMP agent is a software process that monitors the UDP port 161 for SNMP messages. Each SNMP message sent to the agent contains a list of management objects to retrieve or modify.

Related links

SNMPv3 on page 136
SNMP community strings on page 142
SNMPv3 support for VRF on page 143
SNMP configuration using ACLI on page 144
SNMP configuration using Enterprise Device Manager on page 158

SNMPv3

The SNMP version 3 (v3) is the third version of the Internet Standard Management Framework and is derived from and builds upon both the original Internet Standard Management Framework SNMP version 1 (v1) and the second Internet Standard Management Framework SNMP version 2 (v2).

The SNMPv3 is not a stand-alone replacement for SNMPv1 or SNMPv2. The SNMPv3 defines security capabilities you must use in conjunction with SNMPv2 (preferred) or SNMPv1. The following figure shows how SNMPv3 specifies a user-based security model (USM) that uses a payload of either an SNMPv1 or an SNMPv2 Protocol Data Unit (PDU).

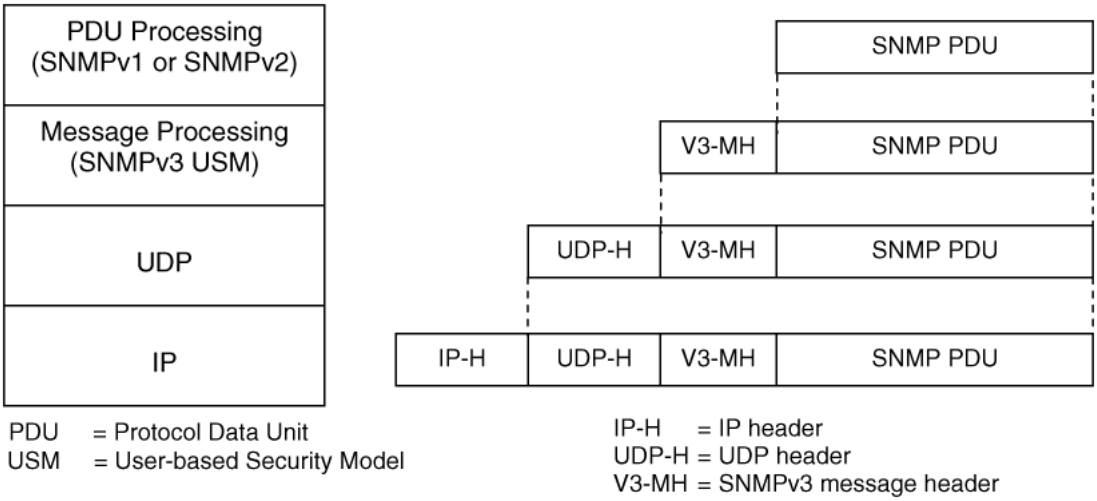

Figure 13: SNMPv3 USM

SNMPv3 is an SNMP framework that supplements SNMPv2 by supporting the following:

- new SNMP message formats
- security for messages
- access control
- remote configuration of SNMP parameters

The recipient of a message can use authentication within the USM to verify the message sender and to detect if the message is altered. According to RFC2574, if you use authentication, the USM checks the entire message for integrity.

An SNMP entity is an implementation of this architecture. Each SNMP entity consists of an SNMP engine and one or more associated applications.

SNMP engine

An SNMP engine provides services for sending and receiving messages, authenticating and encrypting messages, and controlling access to managed objects. A one-to-one association exists between an SNMP engine and the SNMP entity, which contains the SNMP engine.

EngineID

Within an administrative domain, an EngineID is the unique identifier of an SNMP engine. Because a one-to-one association between SNMP engines and SNMP entities exists, the ID also uniquely and unambiguously identifies the SNMP entity within that administrative domain. The system generates an EngineID during the startup process. The SNMP engine contains a

- Dispatcher on page 138
- Message processing subsystem on page 138
- Security subsystem on page 138
- Access control subsystem on page 139

Dispatcher

The dispatcher is part of an SNMP engine. You can use the dispatcher for concurrent support of multiple versions of SNMP messages in the SNMP engine through the following ways:

- To send and receive SNMP messages to and from the network
- To determine the SNMP message version and interact with the corresponding message processing model
- To provide an abstract interface to SNMP applications for delivery of a PDU to an application
- To provide an abstract interface for SNMP applications to send a PDU to a remote SNMP entity

Message processing subsystem

The message processing subsystem prepares messages for sending and extracts data from received messages. The subsystem can contain multiple message processing models.

Security subsystem

The security subsystem provides the following features:

- authentication
- privacy
- security

Authentication

You can use authentication within the SNMPv3 to verify the message sender and whether the message is altered. If you use authentication, the integrity of the message is verified. The supported SNMPv3 authentication protocols are HMAC-MD5 and HMAC-SHA-96.

Privacy

SNMPv3 is an encryption protocol for privacy. Only the data portion of a message is encrypted; the header and the security parameters are not. The privacy protocol that SNMPv3 supports is CBC-DES Symmetric Encryption Protocol.

Security

The SNMPv3 security protects against the following:

- modification of information—protects against altering information in transit
- masquerade—protects against an unauthorized entity assuming the identity of an authorized entity
- message Stream Modification—protects against delaying or replaying messages
- disclosure—protects against eavesdropping
- discovery procedure—finds the EngineID of an SNMP entity for a certain transport address or transport endpoint address.
- time synchronization procedure—facilitates authenticated communication between entities

The SNMPv3 does not protect against the following:

- denial-of-service—prevention of exchanges between manager and agent
- traffic analysis—general pattern of traffic between managers and agents

Access control subsystem

SNMPv3 provides a group option for access policies.

The access policy feature in Virtual Services Platform 9000 determines the access level for the users connecting to the device with different services like File Transfer Protocol (FTP), Trivial FTP (TFTP), Telnet, and rlogin. The system access policy feature is based on the user access levels and network address. This feature covers services, such as TFTP, HTTP, SSH, rlogin, and SNMP. However, with the SNMPv3 engine, the community names do not map to an access level. The View-based Access Control Model (VACM) determines the access privileges.

Use the configuration feature to specify groups for the SNMP access policy. You can use the access policy services to cover SNMP. Because the access restriction is based on groups defined through the VACM, the synchronization is made using the SNMPv3 VACM configuration. The administrator uses this feature to create SNMP users (USM community) and associate them to groups. You can configure the access policy for each group and network.

The following are feature specifications for the group options:

• After you enable SNMP service, this policy covers all users associated with the groups configured under the access policy. The access privileges are based on access allow or deny. If you select allow, the VACM configuration determines the management information base (MIB)-views for access.

• The SNMP service is disabled by default for all access policies.

• The access level configured under **`access-policy policy <id>`** does not affect SNMP service. The VACM configuration determines the SNMP access rights.

User-based security model

In a USM system, the security model uses a defined set of user identities for any authorized user on a particular SNMP engine. A user with authority on one SNMP engine must also have authorization on all SNMP engines, with which the original SNMP engine communicates.

The USM provides the following levels of communication:

- NoAuthNoPriv—communication without authentication and privacy
- AuthNoPriv—communication with authentication and without privacy
- AuthPriv—communication with authentication and privacy

The following figure shows the relationship between USM and VACM.

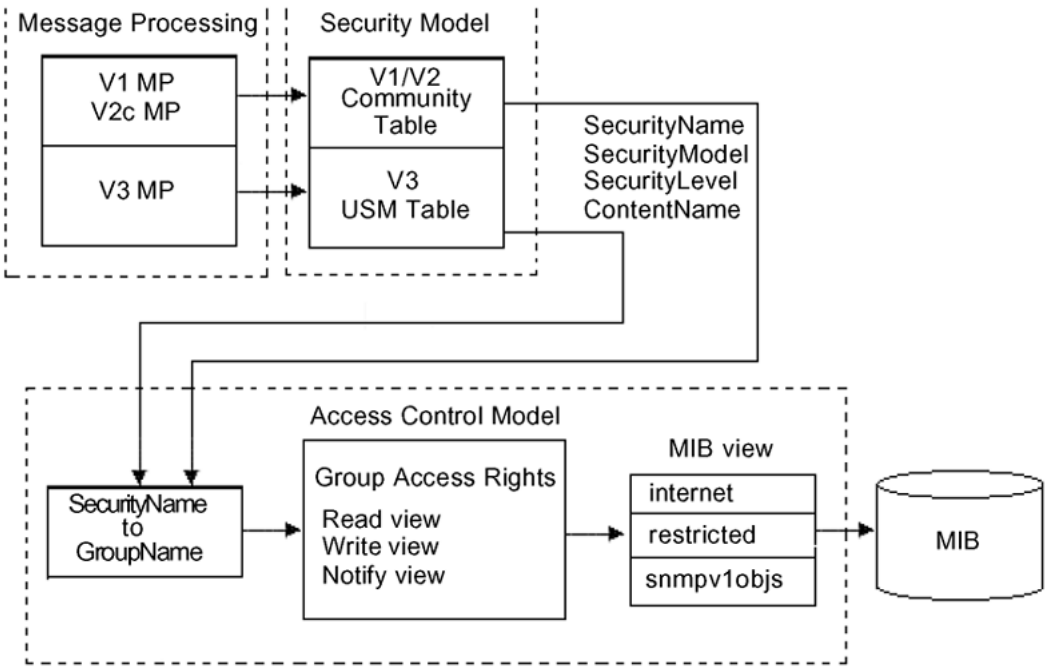

Figure 14: USM association with VACM

View-based Access Control

View-based Access Control Model (VACM) provides group access, group security levels, and context based on a predefined subset of MIB objects. These MIB objects define a set of managed objects and instances.

VACM is the standard access control mechanism for SNMPv3, and it provides the following:

- authorization service to control access to MIB objects at the PDU level
- alternative access control subsystems

The access is based on principal, security level, MIB context, object instance, and type of access requested (read or write). You can use the VACM MIB to define the policy and control remote management.

SNMPv3 encryption

A user-based security module for SNMPv3 is defined as a security subsystem within an SNMP engine. Currently Virtual Services Platform 9000 USM uses HMAC-MD5-96 and HMAC-SHA-96 as the authentication protocols, and CBC-DES as the privacy protocol. Use USM to use other protocols instead of, or concurrently with, these protocols. CFB128-AES-128, an AES-based Symmetric Encryption Protocol, is an alternative privacy protocol for the USM.

The AES standard is the encryption standard (FIPS-197) intended to be used by the U.S. Government organizations to protect sensitive information. The AES standard is also becoming a global standard for commercial software and hardware that uses encryption or other security features.

Important:

Due to export restrictions, the SNMPv3 encryption capability is separate from the main image. For more information about downloading and enabling the SNMPv3 encryption image, see Downloading the software on page 145 and Loading the SNMPv3 encryption modules on page 146. SNMPv3 does not function properly without the use of this image.

The AES-based symmetric encryption protocol

This symmetric encryption protocol provides support for data confidentiality. The system encrypts the designated portion of the SNMP message and includes it as part of the transmitted message.

The USM specifies that the scoped PDU is the portion of the message that requires encryption. An SNMP engine that can legitimately originate messages on behalf of the appropriate user shares a secret value, in combination with a timeliness value and a 64-bit integer, used to create the (localized) encryption/decryption key and the initialization vector.

The AES encryption key and Initialization Vector

The AES encryption key uses the first 128 bits of the localized key. The 128-bit Initialization Vector (IV) is the combination of the authoritative SNMP engine 32-bit snmpEngineBoot, the SNMP engine 32-bit snmpEngineTime, and a local 64-bit integer. The system initializes the 64-bit integer to a pseudo-random value at startup time.

Data encryption

Virtual Services Platform 9000 handles data encryption in the following manner:

1. The system treats data as a sequence of octets.
2. The system divides the plaintext into 128-bit blocks.

   The first input block is the IV, and the forward cipher operation is applied to the IV to produce the first output block.
3. The system produces the first cipher text block by executing an exclusive-OR function on the first plaintext block with the first output block.
4. The system uses the cipher text block as the input block for the subsequent forward cipher operation.
5. The system repeats the forward cipher operation with the successive input blocks until it produces a cipher text segment from every plaintext segment.
6. The system produces the last cipher text block by executing an exclusive-OR function on the last plaintext segment of r bits (r is less than or equal to 128) with the segment of the r most significant bits of the last output block.

Data decryption

Virtual Services Platform 9000 handles data decryption in the following manner:

1. In CFB decryption, the IV is the first input block, the system uses the first cipher text for the second input block, the second cipher text for the third input block, and this continues until the system runs out of blocks to decrypt.
2. The system applies the forward cipher function to each input block to produce the output blocks.
3. The system passes the output blocks through an exclusive-OR function with the corresponding cipher text blocks to recover the plaintext blocks.

4. The system sends the last cipher text block (whose size r is less than or equal to 128) through an exclusive-OR function with the segment of the r most significant bits of the last output block to recover the last plaintext block of r bits.

Trap notifications

You configure traps by creating SNMPv3 trap notifications, creating a target address, to which you want to send the notifications, and specifying target parameters. For more information about how to configure trap notifications, see *Troubleshooting Avaya Virtual Services Platform 9000,* NN46250-700.

Related links

Simple Network Management Protocol (SNMP) on page 136

SNMP community strings

For security reasons for SNMPv1 and SNMPv2, the SNMP agent validates each request from an SNMP manager before responding to the request by verifying that the manager belongs to a valid SNMP community. An SNMP community is a logical relationship between an SNMP agent and one or more SNMP managers (the manager software implements the protocols used to exchange data with SNMP agents). You define communities locally at the agent level.

The agent establishes one community for each combination of authentication and access control characteristics that you choose. You assign each community a unique name (community string), and all members of a community have the same access privileges, either read-only or read-write:

- Read-only: members can view configuration and performance information.
- Read-write: members can view configuration and performance information, and change the configuration.

By defining a community, an agent limits access to its MIB to a selected set of management stations. By using more than one community, the agent can provide different levels of MIB access to different management stations.

SNMP community strings are used when a user logs on to the device over SNMP, for example, using an SNMP-based management software. You set the SNMP community strings using ACLI . If you have read/write/all access authority, you can modify the SNMP community strings for access to the device through Enterprise Device Manager (EDM).

When you create a new SNMP-server community, you must specify a group name if you use a new security name.

The security name creates the link between the SNMP-server community and the groups, which allows the SNMP agent to provide different levels of MIB access to different management stations. If you do not create the link among the SNMP-server community, security name, and group, the SNMP-server community does not function.

You can only create the link among the group, security name, and SNMP-server community name (community string), when you initially create the SNMP-server community. After you create the

SNMP-server community, if you want to link a group (new or existing) to a security name, you must first delete the SNMP-server community name (community string), and then recreate it.

Avaya provides community strings for SNMPv1 and SNMPv2. If you want to use SNMPv3 only, you must disable SNMPv1 and SNMPv2 access by deleting the default community string entries and create the SNMPv3 user and group. The groups specify the level of authentication and the access privileges. When you create a group name, the system does not allow you to leave a blank space as the group name parameter in ACLI.

For more information, see SNMPv3 on page 136.

The following table lists the default community strings for SNMPv1 and SNMPv2.

VRF	Default community string	Access
GlobalRouter VRF	public	Read access
	private	Write access
ManagementRouter VRF	public:512	Read access
	private:512	Write access

Community strings are encrypted using the blowfish algorithm. Community strings do not appear on the device and are not stored in the configuration file.

Security alert:

Security risk

For security reasons, Avaya recommends that you set the community strings to values other than the factory defaults.

Virtual Services Platform 9000 handles community string encryption in the following manner:

- When the device starts up, community strings are restored from the hidden file.
- When the SNMP community strings are modified, the modifications are updated to the hidden file.

Hsecure with SNMP

If you enable hsecure, the system disables SNMPv1, SNMPv2 and SNMPv3. If you want to use SNMP, you must use the command **`no boot config flag block-snmp`** to re-enable SNMP.

Related links

Simple Network Management Protocol (SNMP) on page 136

SNMPv3 support for VRF

Use Virtual Router Forwarding (VRF) to offer networking capabilities and traffic isolation to customers that operate over the same node (Virtual Services Platform 9000). Each virtual router emulates the behavior of a dedicated hardware router and is treated by the network as a separate physical router. You can use VRF Lite to perform the functions of many routers using a single router

running VRF Lite. This substantially reduces the cost associated with providing routing and traffic isolation for multiple clients.

Related links

Simple Network Management Protocol (SNMP) on page 136

SNMP configuration using ACLI

Configure the SNMP engine to provide services to send and receive messages, authenticate and encrypt messages, and control access to managed objects. A one-to-one association exists between an SNMP engine and the SNMP entity.

- Before you can use SNMPv3 with Data Encryption Standard (DES) or Advanced Encryption Standard (AES) to access the device, you must load the appropriate SNMPv3 encryption module. For more information, see Loading the SNMPv3 encryption modules on page 146.
- To perform the procedures in this section, you must log on to the Global Configuration mode in ACLI. For more information about how to use ACLI, see *Using ACLI and EDM on Avaya Virtual Services Platform 9000,* NN46250-103.

This task flow shows you the sequence of procedures you perform to configure basic elements of SNMP when using ACLI.

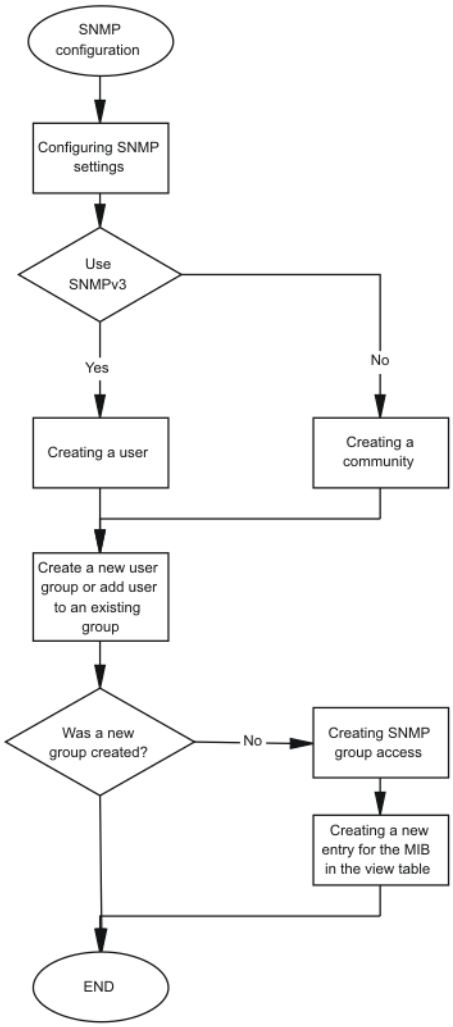

Figure 15: SNMP configuration procedures

Downloading the software

Download new software to upgrade the Avaya Virtual Services Platform 9000. Software downloads can include encryption modules and software images.

Download patches and readme files from the Avaya support site at www.avaya.com/support.

Before you begin

- You must have access to the new software from the Avaya support site: www.avaya.com/support. You need a valid user or site ID and password.

About this task

Download the Advanced Encryption Standard (AES) and Data Encryption Standard (DES) software before you enable the encryption algorithms and use SNMPv3. The AES and DES encryption modules exist in a single file and you can enable them when the file is stored on flash.

Download the SSH encryption software before you enable the 3DES encryption module and use SSH.

Due to export restrictions, the encryption capability is separate from the main software image. SNMPv3 and the SSH server do not function properly without the use of this image.

For more information about file names for the current release, see *Release Notes for Avaya Virtual Services Platform 9000,* NN46250-401.

Important:

You must load the security encryption modules on the device before you can use the protocol.

Procedure

1. From an Internet browser, browse to www.avaya.com/support.
2. Click **Support by Product**.
3. Click **Downloads**.
4. In the product search field, type **Virtual Services Platform 9000**.
5. In the **Choose Release** field, click a release number.
6. Click the download title to view the selected information.
7. Click the file you want to download.
8. Login to download the required software file.
9. Use an FTP client in binary mode to transfer the file to either the Virtual Services Platform 9000 or an external USB device.

Loading the SNMPv3 encryption modules

Before you begin

- Download the file containing the SNMPv3 encryption software. For more information about downloading the SNMPv3 encryption software, see Downloading the software on page 145.

  Important:

  Due to export restrictions, the SNMPv3 encryption capability is separate from the main image. You must copy the SNMPv3 encryption software to Avaya Virtual Services Platform 9000 before you can load the SNMPv3 encryption modules. SNMPv3 does not function properly without this image.

- You must log on to the Global Configuration mode in ACLI.

About this task

Before you can use SNMPv3 with Data Encryption Standard (DES) or Advanced Encryption Standard (AES) to access the device, you must load the appropriate SNMPv3 encryption module.

Procedure

Load the encryption module file on the device:

```
load-encryption-module <DES|AES>
```

Important:

You must load the AES and DES encryption routines by issuing two separate load-encryption-module commands. If you issue the load-encryption-module command for AES, the image is loaded into memory and only the AES routines are enabled; the DES routines are not enabled. To enable the DES routines, you must issue a separate load-encryption-module command for DES.

Example

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

Load the Advanced Encryption Standard security encryption image:

```
VSP-9012:1(config)#load-encryption-module AES
```

Variable definitions

Use the data in the following table to use the **load-encryption-module** command.

Table 25: Variable definitions

Variable	Value
{DES|AES}	Loads the AES or DES SNMPv3 encryption module.

Configuring SNMP settings

Before you begin

- You must log on to the Global Configuration mode in ACLI.

About this task

Configure Simple Network Management Protocol (SNMP) to define or modify the SNMP settings, and specify how secure you want SNMP communications.

Procedure

1. Enable the generation of authentication traps:

   ```
   snmp-server authentication-trap enable
   ```

2. Create a community to form a relationship between an SNMP agent and manager:

```
snmp-server community WORD<1-32>
```

3. Configure the contact information for the system:

   ```
   snmp-server contact WORD<0-255>
   ```

4. Configure the SNMP and IP sender flag to the same value:

   ```
   snmp-server force-iphdr-sender enable
   ```

5. Send the configured source address (sender IP) as the sender network in the notification message:

   ```
   snmp-server force-trap-sender enable
   ```

6. Create an SNMPv1 server host:

   ```
   snmp-server host WORD<1-256> [port <1-65535>] v1 WORD<1-32> [filter
   WORD<1-32>]
   ```

7. Create an SNMPv2 server host:

   ```
   snmp-server host WORD<1-256> [port <1-65535>] v2c WORD<1-32> [inform
   [timeout <1-2147483647>][retries <0-255>][mms <0-2147483647>]]
   [filter WORD<1-32>]
   ```

8. Create an SNMPv3 server host:

   ```
   snmp-server host WORD<1-256> [port <1-65535>] v3 {noAuthNoPriv|
   authNoPriv|authPriv WORD<1-32> [inform [timeout <1-2147483647>]
   [retries <0-255>]] [filter WORD<1-32>]
   ```

9. Configure the system location:

   ```
   snmp-server location WORD<0-255>
   ```

10. Configure the system name:

    ```
    snmp-server name WORD<0-255>
    ```

11. Create a new entry in the notify filter table:

    ```
    snmp-server notify-filter WORD<1-32> WORD<1-32>
    ```

12. Configure the SNMP trap receiver and source IP addresses:

    ```
    snmp-server sender-ip {A.B.C.D} {A.B.C.D}
    ```

Example

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

Enable the generation of authentication traps:

```
VSP-9012:1(config)#snmp-server authentication-trap enable
VSP-9012:1(config)#snmp-server contact xxxx@avaya.com
VSP-9012:1(config)#snmp-server force-iphdr-sender enable
```

Configure hosts to receive SNMP notifications

```
VSP-9012:1(config)#snmp-server host 45.16.149.128 port 1 v1 SNMPv1 filter
SNMPfilterv1
```

Variable definitions

Use the data in the following table to use the **snmp-server** command.

Table 26: Variable definitions

Variable	Value
authentication-trap enable	Enables the generation of authentication traps.
community WORD<1-32>	Specifies a community string to create a community to use in forming a relationship between an SNMP agent and one or more SNMP managers. You require SNMP community strings to access the system using an SNMP-based management software. Use the no option to delete the community string: no snmp-server community *WORD<1–32>*
contact WORD<0-255>	Changes the sysContact information for Virtual Services Platform 9000. WORD<0-255> is an ASCII string from 0–255 characters (for example a phone extension or e-mail address).
force-iphdr-sender enable	Automatically configures the SNMP and IP sender to the same value. The default is no force-iphdr-sender enable.
force-trap-sender enable	Sends the configured source address (sender IP) as the sender network in the notification message.
host WORD<1-256> [port <1-65535>] {v1 WORD<1-32>|v2c WORD<1-32> [inform [timeout <1-2147483647>][retries <0-255>] [mms <0-2147483647>]]|v3 {noAuthPriv|authNoPriv|authPriv} WORD<1-32> [inform [timeout <1-2147483647>][retries <0-255>]]} [filter WORD<1-32>]	Configures hosts to receive SNMP notifications. • host WORD<1-256>— Specifies the IPv4 or IPv6 host address. • port <1-65535>—Specifies the port number that will be the destination port at the UDP level in the trap packet. • v1 WORD<1-32>—Specifies the SNMP v1 security name. • v2c WORD<1-32>—Specifies the SNMPv2 security name. • inform—Configures the attribute to send an inform message. • timeout <1-2147483647>—Specifies the time to wait for a reply before resending the inform message. The time is specified in centiseconds. • retries <0-255>—Specifies the number of packets to be sent if no reply is received. • mms <1-2147483647>—Specifies the maximum message size. • v3 specifies SNMPv3 • noAuthPriv|authNoPriv|authPriv —Specifies the security level. • *WORD<1-32>*—Specifies the user name.

Table continues…

Variable	Value
	• filter—Specifies a filter profile name.
location WORD<0-255>	Configures the sysLocation information for the system. <WORD 0-255> is an ASCII string from 0–255 characters.
login-success-trap	Enables the generation of login-success traps.
name WORD<0-255>	Configures the sysName information for the system. <WORD 0-255> is an ASCII string from 0–255 characters.
notify-filter WORD<1-32> WORD<1-32>	Creates a new entry in the notify filter table. The first WORD<1-32> specifies the filter profile name, and the second WORD<1-32> specifies the subtree OID.
sender-ip {A.B.C.D} {A.B.C.D}	The first {A.B.C.D} configures the SNMP trap receiver and source IP addresses. Specify the IP address of the destination SNMP server receives the SNMP trap notification in the first IP address. The second {A.B.C.D} specifies the source IP address of the SNMP trap notification packet that is transmitted in the second IP address. If you set this to 0.0.0.0, the system uses the IP address of the local interface that is closest (from an IP routing table perspective) to the destination SNMP server.

Creating a user

Before you begin

- You must log on to Global Configuration mode in ACLI.
- Before you can use SNMPv3 with Data Encryption Standard (DES) or Advanced Encryption Standard (AES) to access the device, you must load the appropriate SNMPv3 encryption module. For more information, see Loading the SNMPv3 encryption modules on page 146.

About this task

Create a new user in the USM table to authorize a user on a particular SNMP engine.

Procedure

1. Create a user on a remote system:

   ```
   snmp-server user WORD<1-32> [engine-id WORD<1-32>] [{md5|sha} WORD<1-32>] [{aes|des} WORD<1-32>]
   ```

2. Create a user on the local system:

   ```
   snmp-server user WORD<1-32> [read-view WORD<1-32>] [write-view WORD<1-32>] [notify-view WORD<1-32>] [[{md5|sha} WORD<1-32>] [read-view WORD<1-32>] [write-view WORD<1-32>] [notify-view WORD<1-32>] [{aes|des|3des} WORD<1-32> [read-view WORD<1-32>] [write-view WORD<1-32>] [notify-view WORD<1-32>]]
   ```

3. Add the user to a group:

```
snmp-server user WORD<1-32> group WORD<1-32> [{md5|sha} WORD<1-32>]
[{aes|des} WORD<1-32>]
```

Example

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

Create a local user test1 with MD5:

```
VSP-9012:1(config)# snmp-server user test1 md5 auth-password
```

Variable definitions

Use the data in the following table to use the **snmp-server user** command.

Table 27: Variable definitions

Variable	Value
{aes|des} WORD<1-32>	Specifies a privacy protocol. If no value is entered, no authentication capability exists. The choices are aes, des, or 3des. *WORD<1-32>* assigns a privacy password. If no value is entered, no privacy capability exists. The range is 1–32 characters. **Important:** You must set authentication before you can set the privacy option.
engine-id *WORD<1-32>*	Assigns an SNMPv3 engine ID. The range is 10–64 characters. Use the no operator to remove this configuration.
group WORD<1-32>	Specifies the group access name.
{md5|sha} WORD<1-32>	Specifies an authentication protocol. If no value is entered, no authentication capability exists. The protocol choices are: MD5 and SHA. *WORD<1-32>* specifies an authentication password. If no value is entered, no authentication capability exists. The range is 1–32 characters.
notify-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view. The switch uses elements from this view in the trap messages sent for the user.
read-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view. The switch makes elements from this view available for reading for the user.

Table continues…

Variable	Value
write-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view. The switch makes elements from this view available to be modified by the user.
user WORD<1-32>	Creates the new entry with this security name. The name is used as an index to the table. The range is 1–32 characters. Use the no operator to remove this configuration.

Creating a new user group

Before you begin

- You must log on to Global Configuration mode in ACLI.

About this task

Create a new user group to logically group users who require the same level of access. Create new access for a group in the View-based Access Control Model (VACM) table to provide access to managed objects.

Note:

Avaya created several default groups (public and private) that you can use. To see the list of default groups and their associated security names (secnames), enter `show snmp-server group`. If you use one of these groups, there is no need to create a new group.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Create a new user group:

   ```
   snmp-server group WORD <1-32> WORD<1-32> {auth-no-priv|auth-priv|no-auth-no-priv} [notify-view WORD<1-32>] [read-view WORD<1-32>] [write-view WORD<1-32>]
   ```

Example

This example uses the following variable names:

- The new group name is *lan6grp*.
- The context of the group is *""*, which represents the Global Router (VRF 0).
- The security level is *no-auth-no-priv*.
- The access view name is *v1v2only* for all three views: `notify-view`, `read-view`, and `write-view`.

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

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

Create a new user group:

```
VSP-9012:1(config)# snmp-server group lan6grp "" no-auth-no-priv notify-
view 20 v1v2only read-view v1v2only write-view v1v2only
```

Variable definitions

Use the data in the following table use the **`snmp-server group`** command.

Table 28: Variable definitions

Variable	Value
auth-no-priv	Assigns the minimum level of security required to gain the access rights allowed by this conceptual row. If the auth-no-priv parameter is included, it creates one entry for SNMPv3 access.
auth-priv	Assigns the minimum level of security required to gain the access rights allowed by this conceptual row. If the auth-priv parameter is included, it creates one entry for SNMPv3 access.
group WORD<1-32> WORD<1-32>	The first WORD<1–32> specifies the group name for data access. The range is 1–32 characters. Use the no operator to remove this configuration. The second WORD<1–32> specifies the context name. The range is 1–32 characters. If you use a particular group name value but with different context names, you create multiple entries for different contexts for the same group. You can omit the context name and use the default. If the context name value ends in the wildcard character (*), the resulting entries match a context name that begins with that context. For example, a context name value of foo* matches contexts starting with foo, such as foo6 and foofofum. Use the no operator to remove this configuration.
no-auth-no-priv	Assigns the minimum level of security required to gain the access rights allowed by this conceptual row. If the no-auth-no-priv parameter is included, it creates 3 entries, one for SNMPv1 access, one for SNMPv2c access, and one for SNMPv3 access.
notify-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The switch uses elements from this in the trap messages sent for the communities associated with this group.
read-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The switch makes elements from this available for reading for the communities associated with this group.
write-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The switch makes elements from this view vailable to be modified by the communities associated with this group

Creating a new entry for the MIB in the view table

Before you begin

- You must log on to Global Configuration mode in ACLI.

About this task

Create a new entry in the MIB view table. The default Layer 2 MIB view cannot modify SNMP settings. However, a new MIB view created with Layer 2 permission can modify SNMP settings.

Procedure

Create a new entry:

```
snmp-server view WORD<1-32> WORD<1-32>
```

Example

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

Create MIB views:

```
VSP-9012:1(config)# snmp-server view 2 1.3.8.7.1.4
```

Variable definitions

Use the data in the following table to use the **`snmp-server view`** command.

Table 29: Variable definitions

Variable	Value
The first *WORD<1-32>*	The first WORD <1-32> is used to name the view. The range is 1–32 characters.
The second *WORD<1-32>*	The second WORD <1-32> is used to specify the OID associated with this view. If the subtree OID uses a '+' prefix, or no prefix, this indicates include. The '-' prefix indicates exclude.

Creating a community

About this task

Create a community to use in forming a relationship between an SNMP agent and one or more SNMP managers. You require SNMP community strings to access the system using an SNMP-based management software.

When you create a new SNMP-server community, you must specify a group name if you use a new security name.

The security name creates the link between the SNMP-server community and the groups, which allows the SNMP agent to provide different levels of MIB access to different management stations. If

you do not create the link among the SNMP-server community, security name, and group, the SNMP-server community does not function.

You can only create the link among the group, security name, and SNMP-server community name (community string), when you initially create the SNMP-server community. After you create the SNMP-server community, if you want to link a group (new or existing) to a security name, you must first delete the SNMP-server community name (community string), and then recreate it.

Procedure

1. Enter Global Configuration mode:

   ```
   enable

   configure terminal
   ```

2. Create a community:

   ```
   snmp-server community WORD<1-32> [group WORD<1-32>] [index
   WORD<1-32>] [secname WORD<1-32>]
   ```

   Important:

   - The *group* parameter is only required if you created a new user group. If you use any of the default groups, the *secname* automatically links the community to its associated group so there is no need specify the group in this command.
   - If you do create a new group, use the **`snmp-server community command`** to create an SNMP community with a new security name and link it to the new group you created. There is no separate command to create a security name (secname). You use the **`snmp-server community`** command. The security name is the key to link the community name to a group.
   - You cannot use the `@` character or the string `::` when you create community strings.

Example

In the following example, the community name is *anewcommunity*, the index is *third*, and the secname is *readview*. There is no group specified because this is a default public/read only group.

```
VSP-9012:1>enable
VSP-9012:1#configure terminal
VSP-9012:1(config)#snmp-server community anewcommunity index third secname readview
```

Variable definitions

Use the data in the following table to use the **`snmp-server community`** command.

Table 30: Variable definitions

Variable	Value
community WORD<1-32>	Specifies a community string. The range is 1–32 characters.

Table continues…

Variable	Value
group WORD<1-32>	Specifies the group name. The range is 1–32 characters. When you create a group name, the system does not allow you to leave a blank space as the group name parameter.
index WORD<1-32>	Specifies the unique index value of a row in this table. The range is 1–32 characters.
secname WORD<1-32>	Maps the community string to the security name in the VACM Group Member Table. The range is 1-32 characters.

Adding a user to a group

Before you begin

- You must log on to Global Configuration mode in ACLI.

About this task

Add a user to a group to logically group users who require the same level of access.

Procedure

Create a new user group:

```
snmp-server user WORD<1-32> group WORD<1-32> [{md5 WORD<1-32>|sha
WORD<1-32>) [{aes WORD<1-32>|des WORD<1-32>}]]
```

Example

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

Add a user to a group to logically group users who require the same level of access:

```
VSP-9012:1(config)# snmp-server user test1 group Grouptest1 md5 winter
aes summer
```

Variable definitions

Use the data in the following table to use the **`snmp-server user`** command.

Table 31: Variable definitions

Variable	Value
{aes|des} WORD<1-32>	Specifies a privacy protocol. If no value is entered, no authentication capability exists. The choices are aes, des, or 3des. *WORD<1-32>* assigns a privacy password. If no value is entered, no privacy capability exists. The range is 1–32 characters.

Table continues…

Variable	Value
	Important: You must set authentication before you can set the privacy option.
engine-id *WORD<1-32>*	Assigns an SNMPv3 engine ID. The range is 10–64 characters. Use the no operator to remove this configuration.
group WORD<1-32>	Specifies the group access name.
{md5|sha} WORD<1-32>	Specifies an authentication protocol. If no value is entered, no authentication capability exists. The protocol choices are: MD5 and SHA. *WORD<1-32>* specifies an authentication password. If no value is entered, no authentication capability exists. The range is 1–32 characters.
notify-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view. The switch uses elements from this view in the trap messages sent for the user.
read-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view. The switch makes elements from this view available for reading for the user.
write-view WORD<1-32>	Specifies the view name in the range of 0–32 characters. The first instance is a noAuth view. The second instance is an auth view and the last instance is an authPriv view. The switch makes elements from this view available to be modified by the user.
user WORD<1-32>	Creates the new entry with this security name. The name is used as an index to the table. The range is 1–32 characters. Use the no operator to remove this configuration.

Blocking SNMP

Before you begin

- You must log on to Global Configuration mode in ACLI.

About this task

Disable SNMP by using the SNMP block flag. By default, SNMP access is enabled.

Procedure

Disable SNMP:

```
boot config flags block-snmp
```

Example

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

Disable SNMP:

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

Variable definitions

Use the data in the following table to use the **boot config flags** command.

Table 32: Variable definitions

Variable	Value
block-snmp	Configures the block SNMP flag as active. Use the no operator to remove this configuration. The default is off. To set this option to the default value, use the default operator with the command.

Displaying SNMP system information

About this task

Display SNMP system information to view trap and authentication profiles. For a comprehensive set of SNMP-related `show` commands, see *ACLI Commands Reference for Avaya Virtual Services Platform 9000,* NN46250-104.

Procedure

Display SNMP system information:

```
show snmp-server
```

Example

```
VSP-9012:1(config)#show snmp-server

                trap-sender :
         force-trap-sender : FALSE
       force-iphdr-sender : FALSE
                  contact : http://support.avaya.com/
                 location : 211 Mt. Airy Road,Basking Ridge,NJ 07920
                     name : VSP-9012
       AuthenticationTrap : false
         LoginSuccessTrap : false
```

SNMP configuration using Enterprise Device Manager

Configure SNMP to provide services to send and receive messages, authenticate and encrypt messages, and control access to managed objects with Enterprise Device Manager (EDM).

The following task flow shows you the sequence of procedures you perform to configure basic elements of SNMP using EDM.

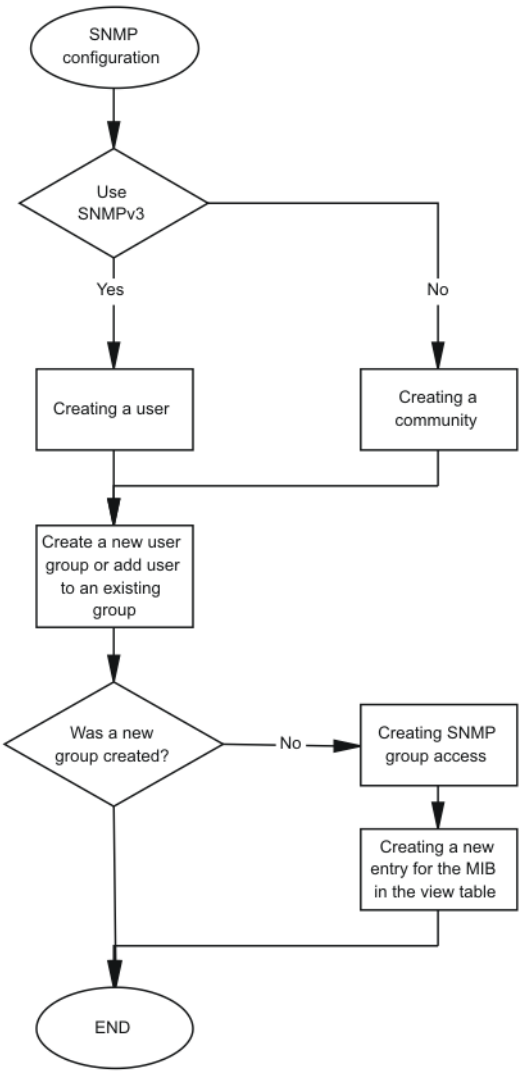

Figure 16: SNMP configuration using Enterprise Device Manager procedures

Creating a user

Create a new user in the USM table to authorize a user on a particular SNMP engine.

About this task

Note:

In EDM, to create new SNMPv3 users you must use the **CloneFromUser** option. However, you cannot clone the default user, named initial. As a result, you must first use ACLI to configure at least one user, and then you can use EDM to create subsequent users with the **CloneFromUser** option.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **SnmpV3**.
2. Click **USM Table**.
3. Click **Insert**.
4. In the **EngineID** box, use the default Engine ID provided or type an administratively-unique identifier to an SNMP engine.
5. In the **User Name** box, type a name.
6. From the **CloneFromUser** list, select a security name, from which the new entry copies authentication data and private data, if required.
7. From the **Auth Protocol** list, select an authentication protocol.
8. In the **Cloned User's Auth Password** box, type the authentication password of the cloned user.
9. In the **New User's Auth Password** box, type an authentication password for the new user.
10. From the **Priv Protocol** list, select a privacy protocol.
11. In the **Cloned User's Priv Password** box, type the privacy password of the cloned user.
12. In the **New User's Priv Password** box, type a privacy password for the new user.
13. Click **Insert**.

Security alert:

Security risk

To ensure security, change the GroupAccess table default view after you set up a new user in the USM table. This prevents unauthorized people from accessing the system using the default user logon. Also, change the Community table defaults, because the community name is used as a community string in SNMPv1/v2 PDU.

USM Table field descriptions

Use the data in the following table to use the **USM Table** tab and the **Insert USM Table** dialog box. Some fields appear only on the Insert USM Table dialog box.

Name	Description
EngineID	Specifies an administratively-unique identifier to an SNMP engine.

Table continues…

Name	Description
UserName	Creates the new entry with this security name. The name is used as an index to the table. The range is 1–32 characters.
SecurityName	Identifies the name on whose behalf SNMP messages are generated.
Clone From User	Specifies the security name, from which the new entry must copy privacy and authentication parameters. The range is 1–32 characters. This option appears only in the **Insert USM Table** dialog box.
Auth Protocol (Optional)	Assigns an authentication protocol (or no authentication) from a list. If you select an authentication protocol, you must enter an old AuthPass and a new AuthPass.
Cloned User's Auth Password	Specifies the current authentication password of the cloned user. This option appears only in the **Insert USM Table** dialog box.
New User's Auth Password	Specifies the authentication password of the new user. This option appears only in the **Insert USM Table** dialog box.
Priv Protocol (Optional)	Assigns a privacy protocol (or no privacy) from a list. If you select a privacy protocol, you must enter an old PrivPass and a new PrivPass.
Cloned User's Priv Password	Specifies the current privacy password of the cloned user. This option appears only in the **Insert USM Table** dialog box.
New User's Priv Password	Specifies the privacy password of the new user. This option appears only in the **Insert USM Table** dialog box.

Creating a new group membership

About this task

Create a new group membership to logically group users who require the same level of access.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **SnmpV3**.
2. Click **VACM Table**.
3. Click the **Group Membership** tab.
4. Click **Insert**.
5. From the **SecurityModel** options, select a security model.
6. In the **SecurityName** box, type a security name.
7. In the **GroupName** box, type a group name.
8. Click **Insert**.

Group Membership field descriptions

Use the data in the following table to use the **Group Membership** tab.

Name	Description
SecurityModel	Specifies the security model to use with this group membership.
SecurityName	Specifies the security name assigned to this entry in the View-based Access Control Model (VACM) table. The range is 1–32 characters.
GroupName	Specifies the name assigned to this group in the VACM table. The range is 1–32 characters.

Creating access for a group

About this task

Create access for a group in the View-based Access Control Model (VACM) table to provide access to managed objects.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **SnmpV3**.
2. Click **VACM Table**.
3. Click the **Group Access Right** tab.
4. Click **Insert**.
5. In the **GroupName** box, type a VACM group name.
6. In the **ContextPrefix** box, select a VRF instance.
7. From the **SecurityModel** options, select a model.
8. From the **SecurityLevel** options, select a security level.
9. In the **ContextMatch** option, select a value to match the context name. This value is **exact** by default.
10. In the **ReadViewName** box, type the name of the MIB view that forms the basis of authorization when reading objects.
11. In the **WriteViewName** box, type the name of the MIB view that forms the basis of authorization when writing objects.
12. In the **NotifyViewName** box, type MIB view that forms the basis of authorization for notifications.
13. Click **Insert**.

Group Access Right field descriptions

Use the data in the following table to use the **Group Access Right** tab.

Name	Description
GroupName	Specifies the name of the new group in the VACM table. The range is 1–32 characters.
ContextPrefix	Specifies if the contextName must match the value of the instance of this object exactly or partially. The range is an SnmpAdminString, 1–32 characters.
SecurityModel	Specifies the authentication checking to communicate to the switch. The security models are: • SNMPv1 • SNMPv2 • USM
SecurityLevel	Specifies the minimum level of security required to gain the access rights allowed. The security levels are: • noAuthNoPriv • authNoPriv • authpriv
ContextMatch	Specifies if the prefix and the context name must match. If the value is exact, all rows where the contextName exactly matches vacmAccessContextPrefix are selected. If you do not select exact, all rows where the contextName with starting octets that exactly match vacmAccessContextPrefix are selected.
ReadViewName	Identifies the MIB view of the SNMP context, to which this conceptual row authorizes read access. The default is the empty string.
WriteViewName	Identifies the MIB view of the SNMP context, to which this conceptual row authorizes write access. The default is the empty string.
NotifyViewName	Identifies the MIB view of the SNMP context, to which this conceptual row authorizes access for notifications. The default is the empty string.

Creating access policies for SNMP groups

About this task

Create an access policy to determine the access level for the users who connect to Avaya Virtual Services Platform 9000 with different services like File Transfer Protocol (FTP), Trivial FTP (TFTP), Telnet, and rlogin.

You only need to create access policies for SNMP groups if you have the access policy feature enabled. For more information about access policies, see *Administering Avaya Virtual Services Platform 9000,* NN46250-600.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Security** > **Control Path**.

2. Click **Access Policies**.
3. Click the **Access Policies-SNMP Groups** tab.
4. Click **Insert**.
5. Beside the **ID** box, click the ellipsis (**...**) button.
6. Select a policy ID from the ID list, and then click **Ok**.
7. In the **Name** box, type a name.
8. From the **Model** options, select a security model.
9. Click **Insert**.

Access Policies — SNMP Groups field descriptions

Use the data in the following table to use the **Access Polices-SNMP Groups** tab.

Name	Description
Id	Specifies the ID of the group policy.
Name	Specifies the name assigned to the group policy. The range is 1–32 characters.
Model	Specifies the security model {SNMPv1|SNMPv2c|USM}.

Assigning MIB view access for an object

About this task

Create a new entry in the MIB View table.

You cannot modify SNMP settings with the default Layer 2 MIB view. However, you can modify SNMP settings with a new MIB view created with Layer 2 permissions.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **SnmpV3**.
2. Click **VACM Table**.
3. In the VACM Table tab, click the **MIB View** tab.
4. Click **Insert**.
5. In the **ViewName** box, type a view name.
6. In the **Subtree** box, type a subtree.
7. In the **Mask** box, type a mask.
8. From the **Type** options, select whether access to the MIB object is granted.
9. Click **Insert**.

MIB View field descriptions

Use the data in the following table to use the **MIB View** tab.

Name	Description
ViewName	Creates a new entry with this group name. The range is 1–32 characters.
Subtree	Specifies a valid object identifier that defines the set of MIB objects accessible by this SNMP entity, for example, 1.3.6.1.1.5.
Mask (optional)	Specifies a bit mask with vacmViewTreeFamilySubtree to determine whether an OID falls under a view subtree.
Type	Determines whether access to a MIB object is granted (included) or denied (excluded). The default is included.

Creating a community

Create a community to use in forming a relationship between an SNMP agent and one or more SNMP managers. You require SNMP community strings for access to Avaya Virtual Services Platform 9000 using an SNMP-based management software.

About this task

When you create a new SNMP-server community, you must specify a group name if you use a new security name.

The security name creates the link between the SNMP-server community and the groups, which allows the SNMP agent to provide different levels of MIB access to different management stations. If you do not create the link among the SNMP-server community, security name, and group, the SNMP-server community does not function.

You can only create the link among the group, security name, and SNMP-server community name (community string), when you initially create the SNMP-server community. After you create the SNMP-server community, if you want to link a group (new or existing) to a security name, you must first delete the SNMP-server community name (community string), and then recreate it.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **SnmpV3**.
2. Click **Community Table**.
3. Click **Insert**.
4. In the **Index** box, type an index.
5. In the **Name** box, type a name that is a community string.
6. In the **SecurityName** box, type a security name.
7. In the **ContextName** box, type the context name.
8. Click **Insert**.

Community Table field descriptions

Use the data in the following table to use the **Community Table** tab.

Name	Description
Index	Specifies the unique index value of a row in this table. The range is 1–32 characters.
Name	Specifies the community string, for which a row in this table represents a configuration.
SecurityName	Specifies the security name in the VACM group member table, to which the community string is mapped. The range is 1–32 characters.
ContextEngineID	Indicates the location of the context, in which management information is accessed when using the community string specified in **Name**.
ContextName	Specifies the context, in which management information is accessed when you use the specified community string.

Viewing all contexts for an SNMP entity

About this task

View contexts to see the contents of the context table in the View-based Access Control Model (VACM). This table provides information to SNMP command generator applications so that they can properly configure the VACM access table to control access to all contexts at the SNMP entity.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **SnmpV3**.
2. Click **VACM Table**.
3. In the **VACM Table** tab, click the **Contexts** tab.

Contexts field descriptions

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

Variable	Value
ContextName	Shows the name identifying a particular context at a particular SNMP entity. The empty contextName (zero length) represents the default context.

Chapter 8: MACsec

The following sections describe Media Access Control Security (MACsec) and its configuration.

MACsec fundamentals

MAC Security (MACsec) is based on the IEEE 802.1ae standard that allows authorized systems in a network to transmit data confidentially and to protect against data transmitted or modified by unauthorized devices.

You can use MACsec for core and enterprise edge switches to secure site-to-site connectivity between data centers, provide data security on links that run over public ground, or outside the physically secure boundaries of a site. You can use MACsec on access switches to secure host to switch connectivity, and host to switch connectivity in an environment where both trusted and untrusted hosts co-exist.

In addition to host level authentication, MACsec capable LANs provide data origin authentication, data confidentiality, and data integrity between authenticated hosts or systems. MACsec protects data from external hacking while the data passes through the public network to reach a receiver host.

MACsec enabled hosts encrypt and decrypt every frame exchanged between them using a MACsec key. The source MACsec host encrypts data frames and destination MACsec host decrypts the frames, ensuring delivery of the frame in its original condition to the recipient host. This ensures secure data communication.

You can configure MACsec encryption over any type of point-to-point Ethernet or emulated Ethernet connection, which includes:

- Dark fiber
- Conventional wavelength-division multiplexing/dense wavelength-division multiplexing (CWDM/DWDM) service
- Multiprotocol label switching (MPLS) point-to-point (ELINE)
- Provider Backbone Bridge Traffic Engineering (PBB-TE)

You can configure MACsec on a physical port or on a trunk group level, which includes: Split MultiLink Trunking (SMLT), distributed MultiLink Trunking (DMLT), or Link aggregate group (LAG).

You configure a pre-shared key on either end of the MACsec link. The pre-shared key is an interface parameter, not a switch-wide parameter.

Note:

MACsec encrypts all packets. If you configure MACsec on one or more MultiLink Trunking (MLT) port members on one side, you must configure MACsec on the same port members on the other side. If you do not do this, the port can physically be up, but any overlying protocols can be down. You do not have to provision MACsec on all MLT port members, but if you configure MACsec on an MLT port member on one side, you must also provision MACsec on the corresponding MLT port on the other side.

One way to detect a mismatch of MACsec configuration is to use Virtual Link Aggregation Protocol (VLACP) on the links.

MACsec provides security at the data link layer or the physical layer. It provides enhancements at the MAC service sub layer for its operation and services to the upper layer.

MACsec is an interface level feature and is disabled by default.

MACsec security modes

The static Connectivity Association Key (CAK) security mode is the only supported MACsec security mode on the VSP platform, and is also the most common mode to enable MACsec.

When you use the static connectivity association key (CAK) security mode to enable MACsec, you configure a community association on both ends of the link. A pre-shared key establishes the MACsec relationship between the switches on each end of the ethernet link. The two pre-shared security association keys (SAKs) include a connectivity association key name (CKN) and its own connectivity association key (CAK). The MACsec CKN and CAK are configured in a connectivity association and the CAK must match on both ends of the link to initially enable MACsec.

To ensure link security, the system periodically refreshes keys based on traffic volume and link speed.

To enable MACsec at the port level, you must first associate the port to the connectivity association. You complete the configuration within the connectivity association, but outside of the secure channel.

When you use the static CAK security mode, the system automatically creates two secure channels, one for inbound traffic and another for outbound traffic. You cannot configure any parameters in the automatically-created secure channels.

The CAK security mode ensures security by frequently refreshing to a new random security key, and by only sharing the security key between the two devices on the MACsec-secured point-to-point link.

MACsec provides options to encrypt user payload, or send in a clear confidential offset, to start the encryption from selectable bytes of 0, 30, and 50 after the SecTag header.

Note:

MACsec replay protect is not supported on VSP 9000 in the current release, and you cannot configure MACsec replay protect on the switch.

Do not enable MACsec replay protect on VOSS switches (VSP 8400, VSP 8200, VSP 7200, or VSP 4000), because replay protect on a VOSS switch may cause the switch to drop packets.

You can choose to configure the following optional features:

- Data encryption — If you disable encryption, MACsec forwards traffic in clear text. You can view that data that is not encrypted in the Ethernet frame that travels across the link. Even if you disable encryption the MACsec header applies to the frame and integrity checks make sure that traffic has not been tampered with.
- Confidentiality offset — If encryption is enabled, and an offset is not configured, all traffic in the connectivity is encrypted. The confidentiality offset provides a way to start encryption after a few bytes following the Ethernet header. The confidentiality offset facilitates traffic flow inspection and classification on intermediate devices by not encrypting the Network Layer header for IPv4 or IPv6. For instance, if you configure the offset to 30, the IPv4 header and the TCP/UDP header are not encrypted. If you configure the offset to 50, the IPv6 header and the TCP/UDP header is not encrypted.

MACsec keys

MACsec provides industry-standard security through secure point-to-point Ethernet links. The point-to-point links are secured after matching security keys.

Security keys are of two types:

- connectivity association key (CAK)—A configured *pre-shared key*, if you enable MACsec using the static connectivity association key (CAK) security mode.

  Important:

  Avaya currently supports the configuration of a pre-shared key to enable MACsec using the static connectivity association key (CAK) security mode.

  The CAK must be identical across both ends of MACsec links.

- secure association key (SAK)—A configured *static secure association key*, if you use the static secure association key (SAK) security mode. SAKs are short-lived keys derived from the CAK or pre-configured for a particular secure channel (SC). MACsec uses a timer to refresh these keys so that the key, as well the session, is secure.

MACsec uses derived keys to encrypt or decrypt data at each end of the MACsec links.

Integrity Check Verification (ICV)

MACsec ensures data integrity using Integrity Check Verification (ICV). MACsec introduces an 8 or 16 byte SecTag after the Ethernet header, and an 8 or 16 byte calculated ICV after the Encrypted Payload. MACsec computes the ICV for the entire frame, starting from the Ethernet header, SecTag until the Checksum. The receiving side recalculates the ICV after data decryption and verifies if the received ICV and computed ICV match. If the ICVs do not match, it indicates that data is modified, and MACsec drops the frame.

Connectivity associations (CA) and secure channels (SC)

You configure MACsec in connectivity associations. You can enable MACsec after you attach a connectivity association to an interface. To use the static CAK security mode to enable MACsec, you must create, and configure connectivity associations on both ends of the link.

You can configure a maximum of 512 connectivity associations on a switch. You can only associate a port to one MACsec connectivity association at a time.

A connectivity association (CA) is a logical representation of a MACsec domain within a network. Each connectivity association is associated with a connectivity association key (CAK). MACsec links are associated with a CA to establish end-to-end MACsec communication. Every MACsec enabled interface is a member of one connectivity association. Switch ports are members of a connectivity association, and can only be a member of one connectivity association.

A secure channel (SC) is a unidirectional channel that connects two endpoints of MACsec. A secure channel is a long-term relationship that persists through the sequence of secure associations.

A secure association (SA) is a short-lived relationship within an SC. MACsec identifies each security association by AN, and supported Secure association key (SAK), which is derived from the CAK. The secure association key is used on both ends of MACsec links to encrypt and decrypt the frames. SAKs are frequently refreshed for security reasons. Periodically changing SAs allows the use of fresh keys without terminating the SC relationship.

You configure connectivity associations. Secure channels and secure associations are internally created in the hardware.

MACsec components

MACsec has three major components:

- **Security entity (SecY)**

  SecY is the entity that operates the MACsec protocol within the system. You configure a secure community association (CA) to meet the requirements of MACsec for connectivity between stations that attach to an individual LAN. Unidirectional secure channels (SC) support each CA. Each SC supports secure transmission of frames through the use of symmetric key cryptography from one of the systems to all the others in the CA.

  Each SecY transmits frames conveying secure MACsec service requests on a single SC, and receives frames conveying secure service indications on separate SCs, one for each of the other SecYs that participate in the secure CA.

  A connectivity association (CA) is a logical representation of a MACsec domain within a network. Each connectivity association is associated with a connectivity association key (CAK). MACsec links are associated with a CA to establish end-to-end MACsec communication. Every MACsec enabled interface is a member of one connectivity association. Switch ports are

members of a connectivity association, and can only be a member of one connectivity association.

A secure channel (SC) is a unidirectional channel that connects two endpoints of MACsec. A secure channel is a long-term relationship that persists through the sequence of secure associations. An SC is a unidirectional point to multipoint communication, and can persist through Secure Association Key (SAK) changes. A sequence of Secure Associations (SAs) support each SC and allow for the periodic use of fresh keys without terminating the relationship. A single secret key or a set of keys support each SA, where the cryptographic operations used to protect one frame require more than one key. An SCI identifies each SC. An SCI is comprised of a unique 48-bit universally administered MAC address, identifying the system to which the transmitting SecY belongs, concatenated with a 16-bit port number, identifying the SecY within that system.

The SCI concatenated with a two-bit AN identifies each SA. The Secure Association Identifier (SAI) created allows the receiving SecY to identify the SA, and the SAK used to decrypt and authenticate the received frame. The AN, and hence the SAI, are only unique for the SAs that can be used or recorded by participating SecYs at any instant.

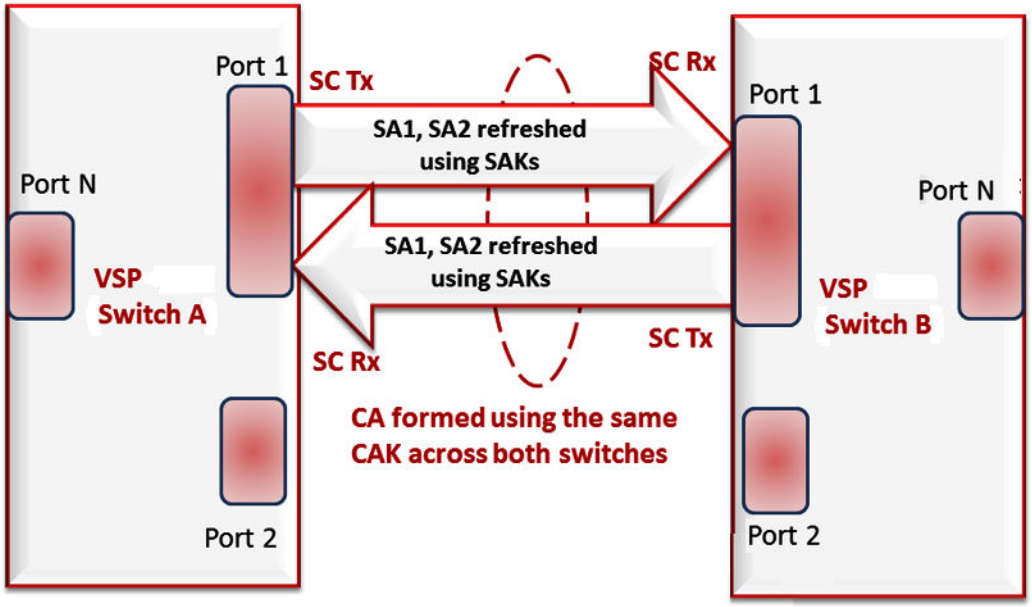

Figure 17: MACsec relationship

- **Key agreement entity (KaY)**

  The KaY in MACsec is responsible for CAK and SAK computations, distributions and maintenance of those keys. CAK is a global key which is persistent until the CA exists. When you configure the CAK, ensure that it is identical across MACsec links. SAK are short-lived keys derived from the CAK, or pre-configured for a particular SC. MACsec uses a timer to refresh these keys so that the key, as well the session, is secure.

  A separate 802.1x-2010 standard is available to automate the above key exchanges and maintenance. The keys are pre-configured.

- **Integrity check verification (ICV) or Cryptographic entity**

The Cryptographic entity provides integrity check protection and validation for frames transmitted or received through the SecY layer. The ICV is calculated for the frame SA/DA, SecTag, User Payload, and CRC. The calculated ICV is appended at the end-of-frame, recalculated at the receiver side of MACsec link and validated to see if they are equal. This is called Integrity Check Verification (ICV). The frames that pass the integrity check are further processed, while the system drops the frames that fail the integrity check.

MACsec configuration provides options to encrypt user payload or send in the clear. The option to start the encryption from N bytes after the Ethernet header also exists.

In the following figure, CA connects switches A, B, and C by their respective SC and SAK. Station D cannot participate in the secure communication between A, B, or C as station D does not know the SAK.

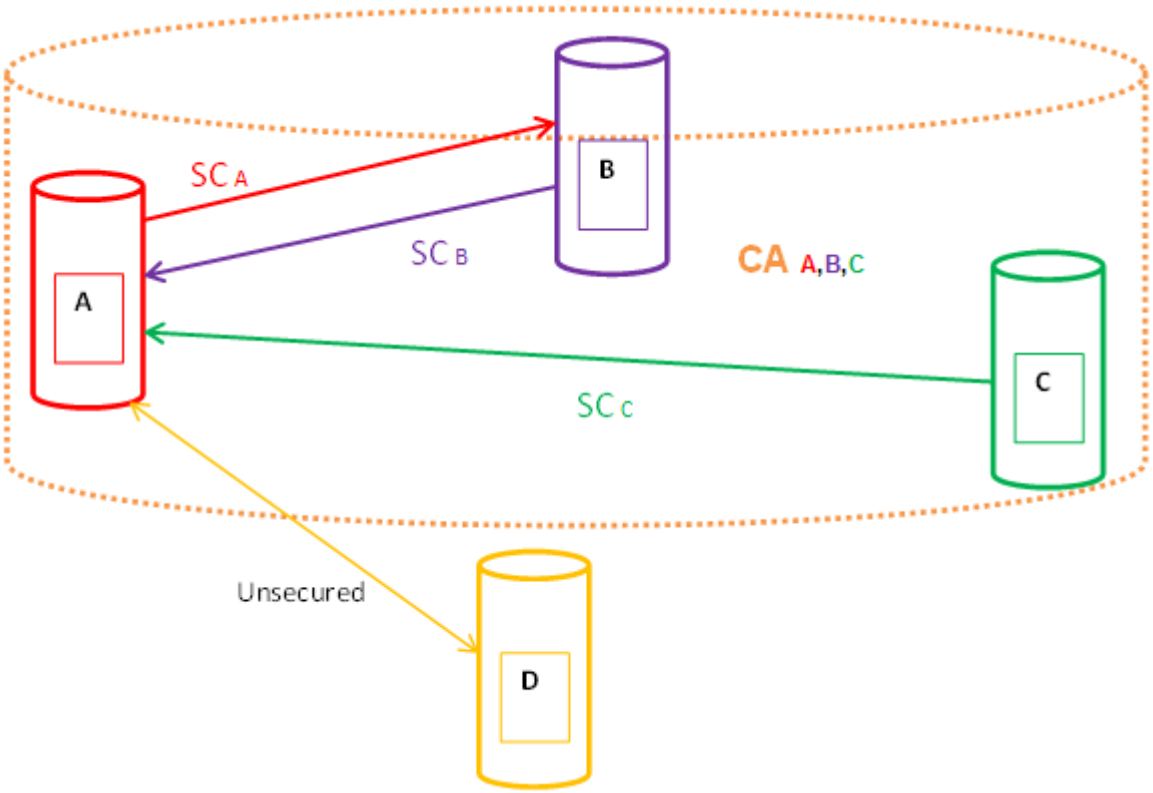

MACsec operation

As shown in the following figure, a host that connects to Switch A sends an Ethernet frame to a host that connects to Switch B. Switch A encrypts the frame, excluding the Ethernet header and optionally the 802.1Q header. Switch A also appends MACsec information like SecTag and ICV to the encrypted payload and transmits the frame using normal frame transmission. This process ensures data confidentiality.

On receiving the frame, Switch B decrypts the frame. Switch B recalculates the ICV using a MACsec key and the SecTag present in the frame. If the ICV present in the received frame matches the recalculated ICV, the switch processes the frame. If the two ICVs do not match, the switch discards the frame. This process ensures data origin authenticity and data integrity. The encryption and decryption algorithms follow the AES-GCM-128 standard.

The MACsec key between switches A and B are statically pre-configured.

Note:

MACsec will be operational between two VSP switches across Point-to-Point Connectivity only when the VSP switches are either directly connected or across a network cloud that provides P2P connectivity between the two VSP switches.

For example, in the following figure you can enable MACsec between two VSP switches across a network cloud where P2P connectivity between the VSPs is provided via services such as P2P, MPLS, L2VPN (ELINE), or connectivity across Dark Fiber. However, it is important to note that MACsec will not be operational between two VSP switches across a network cloud if the intermediate routers/switches need to inspect the VLAN tag or IP header for service classification. This is because MACsec encrypts the entire data frame including the VLAN header and as such the intermediate switches/routers will not have visibility into the same to perform service classification.

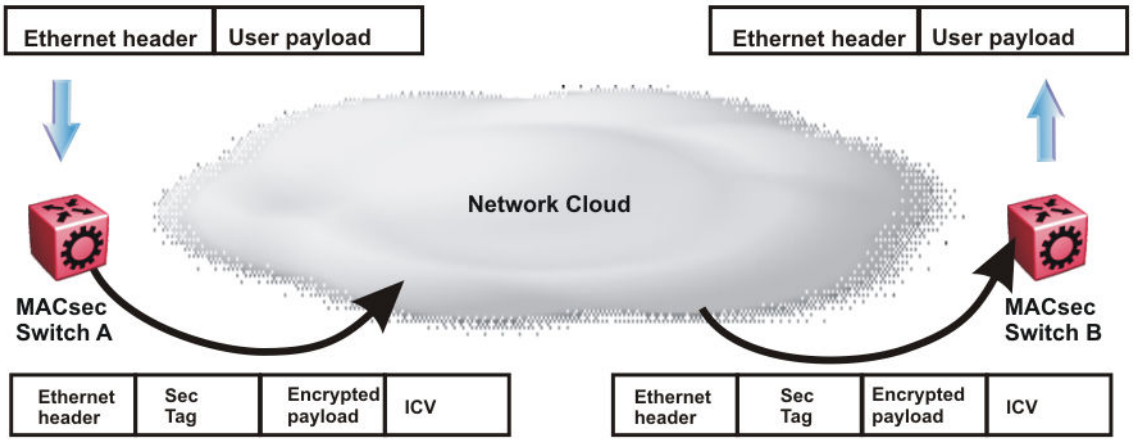

Figure 18: MACsec operation

Hardware requirement

Ports on the 9048XS-2 I/O module support the MACsec feature.

MACsec performance

To monitor MACsec performance, view the performance statistics. For information on the supported statistics, see *Monitoring Performance on Avaya Virtual Services Platform 9000,* NN46250-701.

MACsec configuration using ACLI

Configuring a connectivity association

Use the following procedure to configure a connectivity association (CA) in static CAK security mode using ACLI.

Procedure

1. Enter Global Configuration mode:

   ```
   enable
   configure terminal
   ```

2. Configure a connectivity association (CA):

   ```
   macsec connectivity-association WORD<5-15> connectivity-association-key WORD<10-32>
   ```

3. Enter GigabitEthernet Interface Configuration mode:

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

4. Associate a port with CA:

   ```
   macsec connectivity-association WORD<5-15>
   ```

5. Enable MACsec on the port:

   ```
   macsec enable
   ```

Example

Configure a connectivity association and enable MACsec on a port:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#macsec connectivity-association caname1 connectivity-association-key
1029384756abcdef
Switch:1(config)#interface gigabitethernet 4/17
Switch:1(config-if)#macsec connectivity-association caname1
```

Variable definitions

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

Variable	Value
connectivity-association *WORD<5–15>*	Specifies a connectivity-association name. It is a 5 to 15 character alphanumeric string.
connectivity-association-key *WORD<10–32>*	Specifies the value of the connectivity-association key (CAK). Avaya recommends you use a 32 character hexadecimal string.

Use the data in the following table to use the **interface gigabitethernet** command.

Variable	Value
{slot/port[-slot/port][,...]}	Specifies the port that you want to associate with the connectivity association (CA). Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port).

Updating the connectivity association key (CAK)

Use the following procedure to update the connectivity association key (CAK).

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

2. Disable MACsec on the port:

   ```
   no macsec enable
   ```

3. Update the connectivity association key (CAK):

   ```
   macsec connectivity-association WORD<5-15> connectivity-association
   key WORD<10-32>
   ```

4. Enable MACsec on the port:

   ```
   macsec enable
   ```

Example

Update the connectivity association key (CAK):

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabit 4/17
Switch:1(config-if)#no macsec enable
Switch:1(config-if)#macsec connectivity-association caname1 connectivity-association-key
1029384756abcdef
Switch:1(config-if)#macsec enable
```

Variable definitions

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

Variable	Value
connectivity-association *WORD<5–15>*	Specifies a connectivity-association name. It is a 5 to 15 character alphanumeric string.
connectivity-association-key *WORD<10–32>*	Specifies the value of the connectivity-association key (CAK). Avaya recommends you use a 32 character hexadecimal string.

Use the data in the following table to use the **interface gigabitethernet** command.

Variable	Value
{slot/port[-slot/port][,...]}	Specifies the port that you want to associate with the connectivity association (CA). Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port).

Configuring MACsec encryption on a port

Use the following procedure to enable or disable encryption on a MACsec capable port. The default is disabled.

About this task

If you disable encryption, MACsec forwards traffic in clear text. You can view that data that is not encrypted in the Ethernet frame that travels across the link. Even if you disable encryption the MACsec header applies to the frame and integrity checks make sure that traffic has not been tampered with.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

   ```
   enable
   configure terminal
   ```

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

2. Enable MACsec encryption on the port:

```
macsec encryption enable
```

3. Disable MACsec encryption on the port:

```
no macsec encryption enable
```

Example

Configure MACsec encryption on a port:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabit 4/17
Switch:1(config-if)#macsec encryption enable
```

Variable definitions

Use the data in the following table to use the **`interface gigabitethernet`** command.

Variable	Value
{slot/port[-slot/port][,...]}	Specifies the port that you want to associate with the connectivity association (CA). Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port).

Configuring the confidentiality offset on a port

Use the following procedure to configure the confidentiality offset on a port. The default is disabled.

About this task

The confidentiality offset provides a way to start encryption after a few bytes following the Ethernet header. The confidentiality offset facilitates traffic flow inspection and classification on intermediate devices by not encrypting the Network Layer header for IPv4 or IPv6. For instance, if you configure the offset to 30, the IPv4 header and the TCP/UDP header are not encrypted. If you configure the offset to 50, the IPv6 header and the TCP/UDP header is not encrypted.

Procedure

1. Enter GigabitEthernet Interface Configuration mode:

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

2. Configure confidentiality offset on the port:

```
macsec confidentiality-offset <30-50>
```

3. Disable the confidentiality offset on the port:

   ```
   no macsec confidentiality-offset
   ```

Example

Configuring the confidentiality offsett on the port:

```
Switch:1>enable
Switch:1#configure terminal
Switch:1(config)#interface gigabit 4/17
Switch:1(config-if)#macsec confidentiality-offset 30
```

Variable definitions

Use the data in the following table to use the **`macsec confidentiality-offset`** command.

Variable	Value
<30–50>	Specifies the bytes after the Ethernet header from which data encryption begins. Valid values are 30 and 50.

Use the data in the following table to use the **`interface gigabitethernet`** command.

Variable	Value
{slot/port[-slot/port][,...]}	Specifies the port that you want to associate with the connectivity association (CA). Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port).

Viewing the MACsec connectivity association details

Perform this procedure to view the MACsec connectivity association (CA) details.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View the MACsec connectivity association (CA) details:

   ```
   show macsec connectivity-association WORD<5-15>
   ```

   Note:

   This command displays the MACsec connectivity association (CA) details, including the MD5 hashed value of the CA key.

Example

View the MACsec connectivity association details:

```
Switch:1>enable
Switch:1#show macsec connectivity-association ca333

================================================================================
                     MACSEC Connectivity Associations Info
================================================================================
  Connectivity                Connectivity                      Port
Association Name             Association Key Hash             Members
--------------------------------------------------------------------------------
 ca333                 1304a8fcc51296e7229683ff6882424a       4/17
```

Variable definitions

Use the data in the following table to use the **`show macsec connectivity-association`** command.

Variable	Value
WORD<5–15>	Specifies a connectivity-association name in a 5 to 15 character alphanumeric string.

Viewing MACsec status

Perform this procedure to view MACsec status.

About this task

This command displays the status for the following:

- MACsec status
- MACsec encryption status
- MACsec replay protect status and window
- The associated Connectivity Association (CA) name

Note:

If you do not specify a port number, the information on all MACsec capable interfaces is displayed.

Procedure

1. Enter Privileged EXEC mode:

   ```
   enable
   ```

2. View the MACsec status:

   ```
   show macsec status {slot/port[-slot/port][,...]}
   ```

3. Display all MACsec related information:

   ```
   show MACsec
   ```

Example

View the MACsec status:

```
Switch:1>enable
Switch:1#show macsec status

================================================================================
                                MACSEC Port Status
================================================================================
        MACSEC   Encryption  Replay       Replay        Encryption     CA
PortId  Status     Status    Protect   Protect W'dow     Offset       Name
--------------------------------------------------------------------------------
4/13    disabled   disabled   disabled       --        none         Nil
4/17    enabled    enabled    disabled       50      ipv4Offset(30) ca333

Switch:1#show macsec status 4/17

================================================================================
                                MACSEC Port Status
================================================================================
        MACSEC   Encryption  Replay       Replay        Encryption     CA
PortId  Status     Status    Protect   Protect W'dow     Offset       Name
--------------------------------------------------------------------------------
4/17    enabled    enabled    disabled       50      ipv4Offset(30) ca333
```

Display all MACsec information:

```
Switch:1>show macsec

================================================================================
                     MACSEC Connectivity Associations Info
================================================================================
  Connectivity               Connectivity                    Port
Association Name          Association Key Hash               Members
--------------------------------------------------------------------------------
SMLT105              82a439c7b005be7a5a05087d41df25d4       6/33
Building1           11dfc6a854879f910ba3e396812ebe05

All 2 out of 2 Total Num of Macsec connectivity associates displayed


================================================================================
                                MACSEC Port Status
================================================================================
         MACSEC     Encryption  Replay      Replay       Encryption       CA

PortId   Status     Status     Protect     Protect W'dow   Offset        Name

--------------------------------------------------------------------------------
6/1      disabled   disabled   disabled      --             none          Nil

6/2      disabled   disabled   disabled      --             none          Nil

6/3      disabled   disabled   disabled      --             none          Nil

6/4      disabled   disabled   disabled      --             none          Nil

6/5      disabled   disabled   disabled      --             none          Nil

6/6      disabled   disabled   disabled      --             none          Nil

6/7      disabled   disabled   disabled      --             none          Nil
```

```
6/8       disabled   disabled   disabled      --             none          Nil
6/9       disabled   disabled   disabled      --             none          Nil
6/10      disabled   disabled   disabled      --             none          Nil
6/11      disabled   disabled   disabled      --             none          Nil
6/12      disabled   disabled   disabled      --             none          Nil

--More-- (q = quit)
```

Variable definitions

Use the data in the following table to use the **`show macsec status`** command.

Variable	Value
{slot/port[-slot/port][,...]}	Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port).

MACsec configuration using EDM

Configuring connectivity associations

Use the following procedure to configure connectivity associations (CA) using EDM.

Note:

For VSP 9000, MACsec is only supported on the 9048XS-2 module.

Procedure

1. In the navigation tree, expand the following folders: **Configuration** > **Edit**.
2. Click **Chassis**.
3. In the Chassis window, click the **MAC Security** tab.
4. Click **Insert**.
   a. In the **AssociationName** field, type the connectivity-association name.
   b. In the **AssociationKey** field, type the value of the connectivity-association key.

      Note:

      The connectivity-association key appears as an MD5-hashed text in the MAC security table.
   c. Click **Insert** to save the configuration.

5. Click **Apply**.

Configuring CA field descriptions

Use the data in the following table to use the **MAC Security** tab.

Name	Description
AssociationName	Specifies a name for each connectivity association configured on the device.
AssociationKey	Specifies a pre-shared, connectivity association key associated with each connectivity association configured on the device.
AssociationPortMembers	Specifies the set of ports for which this connectivity association is associated.

Associating a port with a connectivity association

Use the following procedure to associate a port with a connectivity association (CA) using EDM.

For VSP 9000, MACsec is only support on the 9048XS-2 module.

Procedure

1. In the Device Physical View, click on the port that you want to associate with the connectivity association.
2. In the navigation tree, expand the following folders: **Configuration** > **Edit** > **Port**.
3. Click **General**.
4. In the Port General window, click the **MAC Security** tab.
5. In the **CAName** field, type the connectivity-association name.
6. In the **OffsetValue** field, select the value of confidentiality offset to be achieved.
7. Select the **EncryptionEnable** checkbox to enable encryption for the frames transmitted on the port.
8. Select the **Macsec Enable** checkbox to enable MACsec on the port.
9. Click **Apply** to save the configuration.

Associating a port with CA field descriptions

Use the data in the following table to configure the **MAC security** tab.

Name	Description
CAName	Specifies the name of the connectivity association attached to the port or interface.

Table continues…

Name	Description
OffsetValue	Offsets MACsec encryption in an IPv4 TCP/UDP header or IPv6 TCP/UDP header. The confidentiality offset provides a way to start encryption after a few bytes following the Ethernet header. The confidentiality offset facilitates traffic flow inspection and classification on intermediate devices by not encrypting the Network Layer header for IPv4 or IPv6. For instance, if you configure the offset to 30, the IPv4 header and the TCP/UDP header are not encrypted. If you configure the offset to 50, the IPv6 header and the TCP/UDP header is not encrypted.
EncryptionEnable	Specifies the encryption status per port. Use this field to enable or disable encryption for each MACsec capable port.
Macsec Enable	Enables or disables MACsec on the port.

Glossary

American Standard Code for Information Interchange (ASCII)	A code to represent characters in computers. ASCII uses uppercase and lowercase alphabetic letters, numeric digits, and special symbols.
authentication server	A RADIUS server that provides authorization services to the authenticator, which is software that authorizes or rejects a supplicant attached to the other end of the LAN segment.
Authentication, Authorization, and Accounting (AAA)	Authentication, Authorization, and Accounting (AAA) is a framework used to control access to a network, limit network services to certain users, and track what users do. Authentication determines who a user is before allowing the user to access the network and network services. Authorization allows you to determine what you allow a user to do. Accounting records what a user is doing or has done.
authenticator	Software on Virtual Services Platform 9000 that authorizes or rejects a supplicant, such as a PC, attached to the other end of a LAN segment.
AV pairs	AV pairs are strings of text in the form “attribute-value” that are sent between a network access server (NAS) and a TACACS+ daemon as part of the TACACS+ protocol.
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.
Challenge Handshake Authentication Protocol (CHAP)	An access protocol that exchanges a random value between the server and the client and is encrypted with a challenge password.
Control Processor Unit High Availability (CPU-HA)	CPU-HA activates two CP modules simultaneously. The CP modules exchange topology data so, if a failure occurs, either CP module can take precedence in less than one second with the most recent topology data.
controlled port	In relation to EAPoL, any port on the device with EAPoL enabled.

daemon/server	A daemon is a program that services network requests for authentication and authorization, verifies identities, grants or denies authorizations, and logs accounting records.
Data Encryption Standard (DES)access control entry (ACE)	A cryptographic algorithm that protects unclassified computer data. The National Institute of Standards and Technology publishes the DES in the Federal Information Processing Standard Publication 46-1.
Extensible Authentication Protocol over LAN (EAPoL)	A port-based network access control protocol. EAPoL provides security in that it prevents users from accessing network resources before they are authenticated.
Global routing engine (GRE)	The base router or routing instance 0 in the Virtual Routing and Forwarding (VRF).
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.
Internet Engineering Task Force (IETF)	A standards organization for IP data networks.
Layer 2	Layer 2 is the Data Link Layer of the OSI model. Examples of Layer 2 protocols are Ethernet and Frame Relay.
Layer 3	Layer 3 is the Network Layer of the OSI model. An example of a Layer 3 protocol is Internet Protocol (IP).
Local Area Network (LAN)	A data communications system that lies within a limited spatial area, uses a specific user group and topology, and can connect to a public switched telecommunications network (but is not one).
management information base (MIB)	The MIB defines system operations and parameters used for the Simple Network Management Protocol (SNMP).
mask	A bit string that the device uses along with an IP address to indicate the number of leading bits in the address that correspond with the network part.
Media Access Control (MAC)	MAC arbitrates access to and from a shared medium.
Message Digest 5 (MD5)	A one-way hash function that creates a message digest for digital signatures.
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.
network access server (NAS)	A network access server (NAS) is a single point of access to a remote device. The NAS acts as a gateway to guard the remote device. A client connects to the NAS and then the NAS connects to another device to verify the credentials of the client. Once verified the NAS allows or disallows access to the device. Network access servers are almost exclusively used with Authentication, Authorization, and Accounting (AAA) servers.
next hop	The next hop to which a packet can be sent to advance the packet to the destination.
Packet Capture Tool (PCAP)	A data packet capture tool that captures ingress and egress (on Ethernet modules only) packets on selected ports. You can analyze captured packets for troubleshooting purposes.
Point-to-Point Protocol (PPP)	Point-to-Point Protocol is a basic protocol at the data link layer that provides its own authentication protocols, with no authorization stage. PPP is often used to form a direct connection between two networking nodes.
port	A physical interface that transmits and receives data.
Port Access Entity (PAE)	Software that controls each port on the switch. The PAE, which resides on the device, supports authenticator functionality. The PAE works with the Extensible Authentication Protocol over LAN (EAPoL).
Protocol Data Units (PDUs)	A unit of data that is specified in a protocol of a specific layer and that consists of protocol-control information of the specific layer and possibly user data of that layer.
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.
Read Write All (RWA)	An access class that lets users access all menu items and editable fields.
Remote Access Dial-In User Services (RADIUS)	Remote Access Dial-In User Services (RADIUS) can secure networks against unauthorized access, and allow communication servers and clients to authenticate the identity of users through a central database. You can use RADIUS to secure access to the device (console/Telnet/SSH), and RADIUS accounting to track the management sessions for ACLI only. RADIUS authentication allows the remote server to authenticate logons. RADIUS accounting logs all of the activity of each remote user in a session on the centralized RADIUS accounting server. RADIUS uses UDP.

remote login (rlogin)	An application that provides a terminal interface between hosts (usually UNIX) that use the TCP/IP network protocol. Unlike Telnet, rlogin assumes the remote host is, or behaves like, a UNIX host.
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.
Secure Copy (SCP)	Secure Copy securely transfers files between the switch and a remote station.
Simple Network Management Protocol (SNMP)	SNMP administratively monitors network performance through agents and management stations.
supplicant	A device, such as a PC, that applies for access to the network.
Terminal Access Controller Access Control System plus	Terminal Access Controller Access Control System plus (TACACS+) is a security protocol that provides centralized validation of users who attempt to gain access to a router or network access server. TACACS+ uses Transmission Control Protocol (TCP) for its transport to ensure reliable delivery and encrypts the entire body of the packet. TACACS+ provides separate authentication, authorization, and accounting services. TACACS+ is not compatible with previous versions of TACACS.
User Datagram Protocol (UDP)	In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. TCP/IP host systems use UDP for application-to-application programs.
user-based policies (UBP)	Establishes and enforces roles and conditions on an individual user basis for access ports in the network.
view-based access control model (VACM)	Provides context, group access, and group security levels based on a predefined subset of management information base (MIB) objects.
virtual 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.