Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

aruba

a Hewlett Packard
Enterprise company

Contents

Chapter 6 Local MAC Authentication

Chapter 7 Port-based MAC authentication

Chapter 8 MAC ACLs

Chapter 9 TACACS+ Authentication

Contents   **7**

Contents

Chapter 22 Traffic/Security Filters and Monitors..................................... 497

Chapter 23 Configuring Port and User-Based Access Control (802.1X).518

Contents   **17**

This guide provides information on how to configure access security features and user authentication.

Applicable products

This guide applies to these products:

Aruba 2930F Switch Series (JL253A, JL254A, JL255A, JL256A, JL258A, JL259A, JL260A, JL261A, JL262A, JL263A, JL264A, JL557A, JL558A, JL559A)

Aruba 2930M Switch Series (JL319A, JL320A, JL321A, JL322A, JL323A, JL324A)

Switch prompts used in this guide

Examples in this guide are representative and may not match your particular switch/environment. Examples use simplified prompts as follows:

Prompt	Explanation
`switch#`	`#` indicates manager context (authority).
`switch>`	`>` indicates operator context (authority).
`switch(config)#`	`(config)` indicates the config context.
`switch(vlan-x)#`	`(vlan-x)` indicates the vlan context of config, where *x* represents the VLAN ID. For example: `switch(vlan-128)#`.
`switch(eth-x)#`	`(eth-x)` indicates the interface context of config, where `x` represents the interface. For example: `switch(eth-48)#`.
`switch-Stack#`	`Stack` indicates that stacking is enabled.
`switch-Stack(config)#`	`Stack(config)` indicates the config context while stacking is enabled.
`switch-Stack(stacking)#`	`Stack(stacking)` indicates the stacking context of config while stacking is enabled.
`switch-Stack(vlan-x)#`	`Stack(vlan-x)` indicates the vlan context of config while stacking is enabled, where *x* represents the VLAN ID. For example: `switch-Stack(vlan-128)#`.
`switch-Stack(eth-x/y)#`	`Stack(eth-x/y)` indicates the interface context of config, in the form `(eth-<member-in-stack>/<interface>)`. For example: `switch(eth-1/48)#`

Introduction

This chapter provides an overview of the security features included on your switch. **Access security features** outlines the access security and authentication features, while **Network security features** on page 30 highlights the additional features designed to help secure and protect your network. For detailed information on individual features, see the references provided.

Before you connect your switch to a network, Hewlett Packard Enterprise strongly recommends that you review the section **Getting started with access security** on page 37. It outlines potential threats for unauthorized switch and network access, and provides guidelines on how to prepare the switch for secure network operation.

About this guide

This access security guide describes how to configure security features on your switch.

> **NOTE**
> For an introduction to the standard conventions used in this guide, see "Getting Started" in the basic operation guide for your switch.

For more information

For IPv6-specific security settings and features, see the IPv6 configuration guide for your switch.

For information on which product manual to consult for a specific software feature, see the Software feature index – extended.

For the latest version of all HPE switch documentation, including Release Notes covering recently added features and other software topics, visit the Hewlett Packard Enterprise Networking website at **http://www.hpe.com/support/manuals**.

Access security features

This section provides an overview of the switch access security features, authentication protocols, and methods. For more in-depth information, see the references provided (all chapter and page references are to this access security guide unless a different manual name is indicated).

> **NOTE**
> The Management Interface wizard provides a convenient step-by-step method to prepare the switch for secure network operation. See **Using the Management Interface wizard** on page 38 for details.

Table 1: *Access security and switch authentication features*

Feature	Default setting	Security guidelines	More information and configuration details
Manager password	no password	Configuring a local manager password is a fundamental step in reducing the possibility of unauthorized access through the switch WebAgent and console (CLI and Menu) interfaces. The manager password is set with any of the following methods: • CLI: password manager command, or Management interface wizard • WebAgent: the password options under the Security tab, or Management interface wizard • Menu interface: Console passwords option • SNMP	**Configuring local password security** on page 45 **Using the Management Interface wizard** on page 38 **Using SNMP to view and configure switch authentication features** on page 210
Telnet and web browser access (WebAgent)	enabled	The default remote management protocols enabled on the switch are plain text protocols, which transfer passwords in open or plain text that is easily captured. To reduce the chances of unauthorized users capturing your passwords, secure and encrypted protocols such as SSH and SSL (see below for details) should be used for remote access. This enables you to employ increased access security while still retaining remote client access.	**Using the Management Interface wizard** on page 38 For more on Telnet and the WebAgent, see "Interface Access and System Information" in the management and configuration guide.For RADIUS accounting, see **RADIUS Authentication, Authorization, and Accounting** on page 194

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
		Also, access security on the switch is incomplete without disabling Telnet and the standard web browser access (WebAgent). Unauthorized access can be blocked using these commands: • `no telnet-server`: This command blocks inbound Telnet access. • `no web-management`: This command prevents use of the WebAgent through http (port 80) server access. If you choose not to disable Telnet and the WebAgent, you may want to consider using RADIUS accounting to maintain a record of password-protected access to the switch.	

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Feature	Default setting	Security guidelines	More information and configuration details
SSH	disabled	SSH provides Telnet-like functions through encrypted, authenticated transactions of the following types: • client public-key authentication: uses one or more public keys (from clients) that must be stored on the switch. Only a client with a private key that matches a stored public key can gain access to the switch. • switch SSH and user password authentication: this option is a subset of the client public-key authentication, and is used if the switch has SSH enabled without a login access configured to authenticate the client key. In this case, the switch authenticates itself to clients, and users on SSH clients then authenticate themselves to the switch by providing passwords stored on a RADIUS or TACACS+ server, or locally on the switch. • Secure copy (SC) and secure FTP (SFTP): By opening a secure, encrypted SSH session, you can take advantage of SC and SFTP to provide a secure alternative to TFTP for transferring sensitive switch	**Using the Management Interface wizard** on page 38 **Configuring Secure Shell (SSH)** on page 353

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
		information. For more on SC and SFTP, see the section titled "Using Secure Copy and SFTP" in the "File Transfers" appendix of the management and configuration guide for your switch.	
SSL	disabled	Secure Sockets Layer (SSL) and Transport Layer Security (TLS) provide remote web browser access (WebAgent) to the switch through authenticated transactions and encrypted paths between the switch and management station clients capable of SSL/TLS operation. The authenticated type includes server certificate authentication with user password authentication.	**Using the Management Interface wizard** on page 38 **Configuring Secure Sockets Layer (SSL)** on page 390
SNMP	public, unrestricted	In the default configuration, the switch is open to access by management stations running SNMP management applications capable of viewing and changing the settings and status data in the switch MIB (Management Information Base). Thus controlling SNMP access to the switch and preventing unauthorized SNMP access should be a key element of your network security strategy.	**SNMP security guidelines** on page 40 **Using the Management Interface wizard** on page 38 management and configuration guide, see "Using SNMP Tools to manage the switch".

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
Authorized IP managers	none	This feature uses IP addresses and masks to determine whether to allow management access to the switch across the network through the following: • Telnet and other terminal emulation applications • The WebAgent • SNMP (with a correct community name)	**Using Authorized IP Managers** on page 600
Secure Management VLAN	disabled	This feature creates an isolated network for managing the switches that offer this feature. When a secure management VLAN is enabled, CLI, Menu interface, and WebAgent access is restricted to ports configured as members of the VLAN.	advanced traffic management guide, see "Static Virtual LANs (VLANs)".
ACLs for Management Access Protection	none	ACLs can also be configured to protect management access by blocking inbound IP traffic that has the switch itself as the destination IP address.	**IPv4 Access Control Lists (ACLs)** on page 398

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
TACACS+ Authentication	disabled	This application uses a central server to allow or deny access to TACACS-aware devices in your network. TACACS+ uses username/password sets with associated privilege levels to grant or deny access through either the switch serial (console) port or remotely, with Telnet. If the switch fails to connect to a TACACS+ server for the necessary authentication service, it defaults to its own locally configured passwords for authentication control. TACACS+ allows both login (read-only) and enable (read/write) privilege level access.	**TACACS+ Authentication** on page 160
RADIUS Authentication	disabled	For each authorized client, RADIUS can be used to authenticate operator or manager access privileges on the switch through the serial port (CLI and Menu interface), Telnet, SSH, and Secure FTP/Secure Copy (SFTP/SCP) access methods.	**RADIUS Authentication, Authorization, and Accounting** on page 194

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
802.1X Access Control	none	This feature provides port-based or user-based authentication through a RADIUS server to protect the switch from unauthorized access and to enable the use of RADIUS-based user profiles to control client access to network services. Included in the general features are the following: • User-based access control supporting up to 32 authenticated clients per port. • Port-based access control allowing authentication by a single client to open the port. • Switch operation as a supplicant for point-to-point connections to other 802.1X-compliant switches.	**Configuring Port and User-Based Access Control (802.1X)** on page 518
Web and MAC Authentication	none	These options are designed for application on the edge of a network to provide port-based security measures for protecting private networks and the switch itself from unauthorized access. Because neither method requires clients to run any special supplicant software, both are suitable for legacy systems and temporary access situations where introducing supplicant software is not an attractive option. Both methods rely on using a RADIUS server for authentication. This	**Web and MAC Authentication** on page 99

Feature	Default setting	Security guidelines	More information and configuration details
		simplifies access security management by allowing you to control access from a master database in a single server. It also means that the same credentials can be used for authentication, regardless of which switch or switch port is the current access point into the LAN. Web authentication uses a webpage login to authenticate users for access to the network. MAC authentication grants access to a secure network by authenticating device MAC addresses for access to the network.	

Network security features

This section outlines features and defense mechanisms for protecting access through the switch to the network.

Table 2: *Network security – default settings and security guidelines*

Feature	Default setting	Security guidelines	More information and configuration details
Secure File Transfers	not applicable	Secure Copy and SFTP provide a secure alternative to TFTP and auto-TFTP for transferring sensitive information such as configuration files and log information between the switch and other devices.	management and configuration guide, see "File Transfers" and "Using Secure Copy and SFTP".
Traffic/Security Filters	none	These statically configured filters enhance in-band security (and improve control over access to network resources) by forwarding or dropping inbound network traffic according to the configured criteria. Filter options include: • `source-port filters` : Inbound traffic from a designated, physical source-port will be forwarded or dropped on a per-port (destination) basis. • `multicast filters` : Inbound traffic having a specified multicast MAC address will be forwarded to outbound ports or dropped on a per-port (destination) basis. • `protocol filters` : Inbound traffic having the selected frame (protocol) type will be forwarded or dropped on a per-port (destination) basis.	**Traffic/Security Filters and Monitors** on page 497.

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
Access Control Lists (ACLs)	none	ACLs can filter traffic to or from a host, a group of hosts, or entire subnets. Layer 3 IP filtering with Access Control Lists (ACLs) enables you to improve network performance and restrict network use by creating policies for: • **Switch Management Access:** Permits or denies in-band management access. This includes preventing the use of certain TCP or UDP applications (such as Telnet, SSH, WebAgent, and SNMP) for transactions between specific source and destination IP addresses.) • **Application Access Security:** Eliminating unwanted IP, TCP, or UDP traffic by filtering packets where they enter or leave the switch on specific interfaces. **NOTE** On ACL Security Use: ACLs can enhance network security by blocking selected IP traffic, and can serve as one aspect of maintaining network security.	**IPv4 Access Control Lists (ACLs)** on page 398

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
		However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete security solution.	

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
Port Security, MAC Lockdown, and MAC Lockout	none	The features listed below provide device-based access security in the following ways: • **Port security:** Enables configuration of each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch. Some switch models also include eavesdrop prevention in the port security feature. • **MAC lockdown:** This "static addressing" feature is used as an alternative to port security to prevent station movement and MAC address "hijacking" by allowing a given MAC address to use only one assigned port on the switch. MAC lockdown also restricts the client device to a specific VLAN. • **MAC lockout:** This feature enables blocking of a specific MAC address so that the switch drops all traffic to or from the specified address.	**Configuring and Monitoring Port Security** on page 570See also **Precedence of port-based security options** on page 41.

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
Key Management System (KMS)	none	KMS is available in several switch models and is designed to configure and maintain key chains for use with KMS-capable routing protocols that use time-dependent or time-independent keys. (A key chain is a set of keys with a timing mechanism for activating and deactivating individual keys.) KMS provides specific instances of routing protocols with one or more Send or Accept keys that must be active at the time of a request.	**Key Management System** on page 610
ICMP Rate-Limiting	none	This feature helps defeat ICMP denial-of-service attacks by restricting ICMP traffic to percentage levels that permit necessary ICMP functions, but throttle additional traffic that may be due to worms or viruses (reducing their spread and effect).	management and configuration guide see "Port Traffic Controls" and "ICMP Rate-Limiting"

Table Continued

Feature	Default setting	Security guidelines	More information and configuration details
Spanning Tree Protection	none	These features prevent your switch from malicious attacks or configuration errors: • **BPDU Filtering and BPDU Protection:** Protects the network from denial-of-service attacks that use spoofing BPDUs by dropping incoming BPDU frames and/or blocking traffic through a port. • **STP Root Guard:** Protects the STP root bridge from malicious attacks or configuration mistakes.	advanced traffic management guide see "Multiple Instance Spanning-Tree Operation".
DHCP Snooping, Dynamic ARP Protection, and Dynamic IP Lockdown	none	These features provide the following additional protections for your network: • **DHCP Snooping:** Protects your network from common DHCP attacks, such as address spoofing and repeated address requests. • **Dynamic ARP Protection:** Protects your network from ARP cache poisoning. • **Dynamic IP Lockdown:** Prevents IP source address spoofing on a	**Configuring Advanced Threat Protection** on page 460.

Feature	Default setting	Security guidelines	More information and configuration details
		per-port and per-VLAN basis. • **Instrumentation Monitor:** Helps identify a variety of malicious attacks by generating alerts for detected anomalies on the switch.	

Getting started with access security

HPE switches are designed as "plug and play" devices, allowing quick and easy installation in your network. In its default configuration the switch is open to unauthorized access of various types. When preparing the switch for network operation, therefore, Hewlett Packard Enterprise strongly recommends that you enforce a security policy to help ensure that the ease in getting started is not used by unauthorized persons as an opportunity for access and possible malicious actions.

Since security incidents can originate with sources inside as well as outside of an organization, your access security provisions must protect against internal and external threats while preserving the necessary network access for authorized clients and users. It is important to evaluate the level of management access vulnerability existing in your network and take steps to ensure that all reasonable security precautions are in place. This includes both configurable security options and physical access to the switch.

Switch management access is available through the following methods:

- Front panel access to the console serial port, see **Physical security** on page 37
- Inbound Telnet access
- Web-browser access (WebAgent)
- SNMP access

For guidelines on locking down your switch for remote management access, see **Using the Management Interface wizard** on page 38.

Physical security

Physical access to the switch allows the following:

- Use of the console serial port (CLI and Menu interface) for viewing and changing the current configuration and for reading status, statistics, and log messages.

- Use of the switch Clear and Reset buttons for these actions:

  ◦ Clearing (removing) local password protection

  ◦ Rebooting the switch

  ◦ Restoring the switch to the factory default configuration (and erasing any nondefault configuration settings)

Keeping the switch in a locked wiring closet or other secure space helps prevent unauthorized physical access.

As additional precautions, you can do the following:

- Disable or re-enable the password-clearing function of the Clear button.
- Configure the Clear button to reboot the switch after clearing any local usernames and passwords.
- Modify the operation of the Reset+Clear button combination so that the switch reboots, but does not restore the switch factory default settings.
- Disable or re-enable password recovery.

Using the Management Interface wizard

The Management Interface wizard provides a convenient step-by-step method to prepare the switch for secure network operation. It guides you through the process of locking down the following switch operations or protocols:

- Setting local passwords
- Restricting SNMP access
- Enabling/disabling Telnet
- Enabling/disabling SSH
- Enabling/disabling remote Web management (WebAgent)
- Restricting WebAgent access to SSL
- Setting timeouts for SSH/Telnet sessions

The wizard can also be used to view the preconfigured defaults and see the current settings for switch access security. The wizard can be launched either through the CLI or the WebAgent.

> **NOTE**
> The wizard security settings can also be configured using standard commands through the CLI, Menu, or WebAgent.

Configuring security settings using the CLI wizard

To configure the security settings using the CLI wizard, follow the steps below:

Procedure

1. At the command prompt, enter `setup mgmt-interfaces`.

The welcome banner appears and the first setup option is displayed (`operator password`). As you advance through the wizard, each setup option displays the current value in brackets `[]`.

Figure 1: *Management Interface wizard configuration*

```
                    Welcome to the Management Interface Setup Wizard

              This wizard will help you with the initial setup of the var-
              ious management interfaces. The current values are shown in
              brackets[]. Type in a new value, or press <Enter> to keep the
              current value. Press CTRL-C at any time to quit the wizard
              without saving any changes. Press ? for help.

              Operator password             [not configured]:
                                            Confirm password:
              Manager password                    [*******]:
                                            Confirm password:
              Restrict SNMP access to SNMPv3 only       [no]:
              SNMPv2 community name              [notpublic]:
              SNMPv2 Community access level  [unrestricted]:
              Telnet enabled                           [yes]:
              SSH enabled                               [no]:
              Web management enabled                   [yes]:
              Restrict Web access to SSL                [no]:
              Timeout for ssh/telnet sessions           [0]:

              Operator password                          :
              Manager password                           :*******
              Restrict SNMP access to SNMPv3 only        :no
              SNMPv2 community name                      :notpublic
              SNMPv2 Community access level              :unrestricted
              Telnet enabled                             :yes
              SSH enabled                                :no
              Web management enabled                     :yes
              Restrict Web access to SSL                 :no
              Timeout for ssh/telnet sessions            :0

              Do you want to save these changes?     [yes]:
```

Current values are shown in brackets (Password entries must be entered twice and will appear as asterisks.)

Type in a new value to change a setting, or press <Enter> to keep the current value.

Summary of current settings (displayed after last wizard option has been set)

To save these settings, press [**Enter**].
To cancel any changes, type [**n**] (for no), then press [**Enter**].

2. When you enter the wizard, you have the following options:

   a. To update a setting, type in a new value, or press **Enter** to keep the current value.

   b. To quit the wizard without saving any changes, press **CTRL-C** at any time.

   c. To access online Help for any option, press **?**.

   After you have gone through each setup option, the wizard displays the summary configuration together with a prompt to save the changes.

3. When the message appears asking if you want to save these changes, you have the following options:

   a. To save your changes, press **Enter**.

   b. To cancel any changes without saving, enter **n** and then press **Enter**.
   After pressing **Enter**, the wizard exits to the command line prompt.

CLI Wizard: Operating notes and restrictions

• Once a password has been configured on the switch, you cannot remove it using the CLI wizard. passwords can be removed by executing the `no password` command directly from the CLI.

• When you restrict SNMP access to SNMPv3 only, the options SNMPv2 community name and access level will not appear.

• The wizard displays the first available SNMPv2 community and allows the user to modify the first community access parameters.

• The wizard creates a new SNMP community only when no communities have been configured on the switch.

WebAgent: Management Interface wizard

To use the Management Interface wizard from the WebAgent, follow the steps below:

Procedure

1. In the navigation tree, select **Security**.

2. Click **Security Wizard**. The Welcome window appears.

3. Select a setup type:

   a. **Typical**—provides a multiple page, step-by-step method to configure security settings, with on-screen instructions for each option.

   b. **Advanced**—provides a single summary screen in which to configure all security settings at once.

See the WebAgent Online Help for detailed information about using the management interface wizard.

SNMP security guidelines

In the default configuration, the switch is open to access by management stations running SNMP, management applications capable of viewing and changing the settings and status data in the switch MIB. Thus, controlling SNMP access to the switch and preventing unauthorized SNMP access should be a key element of your network security strategy.

General SNMP access to the switch

The switch supports SNMP versions 1, 2c, and 3, including SNMP community and trap configuration. The default configuration supports versions 1 and 2c compatibility, which uses plain text and does not provide security options.

Hewlett Packard Enterprise recommends you enable SNMP version 3 for improved security. SNMPv3 includes the ability to configure restricted access and to block all non-version 3 messages (which blocks version 1 and 2c unprotected operation).

SNMPv3 security options include:

• Configuring device communities as a means for excluding management access by unauthorized stations

• Configuring for access authentication and privacy

• Reporting events to the switch CLI and to SNMP trap receivers

- Restricting non-SNMPv3 agents to either read-only access or no access

- Coexisting with SNMPv1 and v2c if necessary.

SNMP access to the authentication configuration MIB

A management station running an SNMP networked device management application can access the management information base (MIB) for read access to the switch status and read/write access to the switch authentication configuration (`hpSwitchAuth`). This means that the switch default configuration now allows SNMP access to security settings in `hpSwitchAuth`.

> **CAUTION**
>
> If SNMP access to the `hpSwitchAuth` MIB is considered a security risk in your network, then you should implement the security precautions described in this procedure.

> **NOTE**
>
> Downloading and booting enables SNMP access to the authentication configuration MIB (the default action). If SNMPv3 and other security safeguards are not in place, the switch authentication configuration MIB is exposed to unprotected SNMP access and you must use the command shown below to disable this access.

For details on this feature, see **Using SNMP to view and configure switch authentication features** on page 210.

See "Configuring for Network Management Applications" in the management and configuration guide for your switch.

Procedure

- If SNMP access to the authentication configuration (`hpSwitchAuth`) MIB described above is not desirable for your network, then immediately use the following command to disable this feature:

  ```
  snmp-server mib hpswitchauthmib excluded
  ```

- If you choose to leave the authentication configuration MIB accessible, then you must do the following to help ensure that unauthorized workstations cannot use SNMP tools to access the MIB:

  1. Configure SNMP version 3 management and access security on the switch.

  2. Disable SNMP version 2c on the switch.

Precedence of security options

This section explains how port-based security options, and client-based attributes used for authentication, get prioritized on the switch.

Precedence of port-based security options

Where the switch is running multiple security options, it implements network traffic security based on the OSI (Open Systems Interconnection model) precedence of the individual options, from the lowest to the highest. The following list shows the order in which the switch implements configured security features on traffic moving through a given port.

1. Disabled/Enabled physical port

2. MAC lockout (applies to all ports on the switch.)

3. MAC lockdown

4. Port security

5. Authorized IP managers

6. Application features at higher levels in the OSI model, such as SSH.

The above list does not address the mutually exclusive relationship that exists among some security features.

Precedence of client-based authentication: Dynamic Configuration Arbiter (DCA)

The Dynamic Configuration Arbiter (DCA) is implemented to determine the client-specific parameters that are assigned in an authentication session.

A client-specific authentication configuration is bound to the MAC address of a client device and may include the following parameters:

- Untagged client VLAN ID

- Tagged VLAN IDs

- Per-port CoS (802.1p) priority

- Per-port rate-limiting on inbound traffic

- Client-based ACLs

DCA allows client-specific parameters configured in any of the following ways to be applied and removed as needed in a specified hierarchy of precedence. When multiple values for an individual configuration parameter exist, the value applied to a client session is determined in the following order (from highest to lowest priority) in which a value configured with a higher priority overrides a value configured with a lower priority:

1. 802.1X authentication parameters (RADIUS-assigned)

2. Web or MAC authentication parameters (RADIUS-assigned)

3. Local, statically configured parameters

Although RADIUS-assigned settings are never applied to ports for unauthenticated clients, the DCA allows configuring and assigning client-specific port configurations to unauthenticated clients, provided that a client MAC address is known in the switch in the forwarding database. DCA arbitrates the assignment of attributes on both authenticated and unauthenticated ports.

DCA does not support the arbitration and assignment of client-specific attributes on trunk ports.

Arbitrating client-specific attributes

In previous releases, client-specific authentication parameters for 802.1X Web, and MAC authentication are assigned to a port using different criteria. A RADIUS-assigned parameter is always given highest priority and overrides statically configured local passwords. 802.1X authentication parameters override Web or MAC authentication parameters.

DCA stores client-specific authentication parameters and prioritizes them according to the following hierarchy of precedence:

1. RADIUS-assigned

    a. 802.1X authentication

    b. Web or MAC authentication

2. Statically (local) configured

Client-specific configurations are applied on a per-parameter basis on a port. In a client-specific profile, if DCA detects that a parameter has configured values from two or more levels in the hierarchy of precedence described above, DCA decides which parameters to add or remove, or whether to fail the authentication attempt due to an inability to apply the parameters.

In addition, DCA supports conflict resolution for QoS (port-based CoS priority) and rate-limiting (ingress) by determining whether to configure either strict or nonstrict resolution on a switchwide basis.

For information on how to configure RADIUS-assigned and locally configured authentication settings, see:

- RADIUS-assigned 802.1X authentication: **Configuring Port and User-Based Access Control (802.1X)** on page 518

- RADIUS-assigned Web or MAC authentication: **Web and MAC Authentication** on page 99

- RADIUS-assigned CoS, rate-limiting, and ACLs: "Configuring RADIUS Server Support for Switch Services"

- Statically (local) configured: **Configuring Username and Password Security** on page 44

Overview

Console access includes both the menu interface and the CLI. There are two levels of console access: manager and operator. For security, you can set a **password pair** (username and password) on each of these levels.

> Usernames are optional. Also, in the menu interface, you can configure passwords, but not usernames. To configure usernames, use the CLI or the WebAgent.
>
> Usernames and passwords for manager and operator access can also be configured using SNMP. See **Using SNMP to view and configure switch authentication features** on page 210.
>
> Usernames and passwords for manager and operator access can also be configured using the Management Interface Wizard. See **Using the Management Interface wizard** on page 38.

Level	Actions Permitted
Manager:	Access to all console interface areas.**This is the default level.** That is, if a manager password has not been set prior to starting the current console session, then anyone having access to the console can access any area of the console interface.
Operator:	Access to the Status and Counters menu, the Event Log, and the CLI, but no Configuration capabilities. On the operator level, the configuration menus, Download OS, and Reboot Switch options in the Main Menu are not available.

Configuring password security

To set up password security:

Procedure

1. Set a Manager password pair (and an operator password pair, if applicable for your system).

2. Exit from the current console session. A Manager password pair will now be needed for full access to the console.

If you do steps 1 and 2, above, then the next time a console session is started for either the menu interface or the CLI, a prompt appears for a password. Assuming you have protected both the manager and operator levels, the level of access to the console interface is determined by which password is entered in response to the prompt.

If you set a manager password, you may also want to configure an inactivity timer. Doing this causes the console session to end after the specified period of inactivity, thus giving you added security against unauthorized console access.

> If the console inactivity-timer expires, any outbound Telnet or SSH sessions open on the switch are terminated.

You can use either of the following to set the inactivity timer:

- **Menu Interface:** System Information screen, Select option 2 — Switch Configuration.

- **CLI:** Use the command (and options) as follows:

```
console inactivity-timer <0|1|5|10|15|20|30|60|120>
```

> ⚠ **CAUTION**
>
> If the switch has no a manager or operator password, anyone having access to the switch through either Telnet, the serial port, or the WebAgent can access the switch with full manager privileges. Also, if you configure only an operator password, entering the operator password enables full manager privileges.

> 📝 **NOTE**
>
> The manager and operator passwords and (optional) usernames control access to the menu interface, CLI, and WebAgent.
>
> If you configure only a manager password (with no operator password), and in a later session the manager password is not entered correctly in response to a prompt from the switch, then the switch does not allow management access for that session.
>
> If the switch has a password for both the manager and operator levels, and neither is entered correctly in response to the switch password prompt, then the switch does not allow management access for that session.
>
> Passwords are case-sensitive.

Configuring local password security

Setting passwords (Menu)

1. From the **Main Menu** select:**3. Console passwords**

   Figure 2: *Set password screen*

2. To set a new password:

a. Select **Set manager password** or **Set operator password**. You will then be prompted with `Enter new password`.

b. Type a password of up to 64 ASCII characters with no spaces, and press **Enter**. (Remember that passwords are case-sensitive.)

c. When prompted with `Enter new password again`, retype the new password and press **Enter**.

After you configure a password, if you subsequently start a new console session, you will be prompted to enter the password. Remember that usernames are optional. If you use the CLI or WebAgent to configure an optional username, the switch will prompt you for the username, and then the password.

Deleting password protection

This procedure deletes **all** usernames (if configured) and passwords (manager and operator).

Option one

Procedure

1. If you have physical access to the switch, press and hold the **Clear** button (on the front of the switch) for a minimum of one second to clear all password protection.

2. Enter new passwords.

Option two

If you do not have physical access to the switch, you will need manager-level access:

• Enter the console at the manager level.

• Select the **Set manager password** option.

• Select to **Delete password Protection**. Selection will prompt the following:

`Continue Deletion of password protection? No/Yes`

  ◦ Press the Space bar to select **Yes**, then press **Enter**.

  ◦ Press **Enter** to clear the **Delete password Protection** message.

Recovering from a lost manager password

If you cannot start a console session at the manager level because of a lost manager password, clear the password by following these steps:

Procedure

1. Get physical access to the switch.

2. Press and hold the **Clear** button on the switch for a minimum of one second.

This deletes all passwords and usernames (manager and operator) used by the console and the WebAgent.

Setting passwords and usernames (CLI)

NOTE You can now configure manager and operator passwords in one step.

Syntax:

```
[no] password <manager|operator|all|port-access> [user-name ASCII-STR] [<plaintext|sha1> ASCII-STR]
```

Sets or clears a local username/password for a given access level.

The command sets or changes existing password(s). If no password is provided in the command, you are prompted to enter the new password twice.

The `[no]` form of the command removes specific local password protection.

NOTE `port-access` is available only if `include-credentials` is enabled.

`<manager|operator|port-access|all>`: Level of access

manager

Configures access to the switch with manager-level privileges.

operator

Configures access to the switch with operator-level privileges.

port-access

Configures access to the switch through 802.1X authentication with operator-level privileges.

user-name <*name*>

The optional text string of the user name associated with the password. Username up to 64 characters.

<plaintext|sha1>

Format for the password entry, and the password itself (up to 64 characters). Specifies the type of algorithm (if any) used to hash the password. Valid values are `plaintext` or `sha-1` The default type is `plaintext`, which is also the only type accepted for the `port-access` parameter.

Configuring manager and operator passwords

- Password entries appear as asterisks.

- You must type the password entry twice.

```
HP Switch(config)# password manager
New password: *******
Please retype new password: *******
HP Switch(config)# password operator
New password: ********
Please retype new password: ********
```

Removing password protection

Removing password protection means to eliminate password security. This command prompts you to verify that you want to remove one or both passwords, then clears the indicated password(s). (This command also clears

the username associated with a password you are removing.) For example, to remove the operator password (and username, if assigned) from the switch, you would do the following:

Syntax:

```
[no] password
```

Executing this command removes password protection from the operator level so anyone able to access the switch console can gain operator access without entering a username or password.

Syntax:

```
[no] password all
```

This command removes both operator and manager password protection.

Removing a password and associated username from a switch

```
HP Switch(config)# no password
Password protection will be deleted, do you want to continue [y/n]? y
HP Switch(config)#
```

Username and password length

The limit on username and password length is 64 characters for the following authentication methods:

- Front-end—WEB User Interface, SSH, and Telnet

- Back-end—RADIUS, TACACS+, and Local

General rules for usernames and passwords

Usernames and passwords are case-sensitive. ASCII characters in the range of 33-126 are valid, including:

- A through Z uppercase characters

- a through z lower case characters

- 0 through 9 numeric characters

- Special characters ' ~ ! @ # $ % ^ & * ( ) - _ = + [ ] { } \ | ; : ' " , < > / ?.

> **NOTE**
>
> The SPACE character is allowed to form a username or password pass-phrase. The username must be in quotes, for example "The little brown fox". A space is not allowed as part of a username without the quotes. A password that includes a space or spaces should not have quotes.

Restrictions for the setmib command

Usernames and passwords can be set using the CLI command `setmib`. They cannot be set using SNMP.

- Quotes are permitted for enclosing other characters, for example, a username or password of abcd can be enclosed in quotes "abcd" without the quotes becoming part of the username or password itself. Quotes can also be inserted between other characters of a username or password, for example, ab"cd. A pair of quotes enclosing characters followed by any additional characters is invalid, for example, "abc"d.

- Spaces are allowed in usernames and passwords. The username or password must be enclosed in quotes, for example, "one two three". A blank space or spaces between quotes is allowed, for example, " ".

Additional restrictions

Some authentication servers prevent the usage of special symbols such as the backslash (\) and quotes (" "). The switch allows the use of these symbols in configurable credentials, but using them can limit access for some users who can use different client software. See the vendor's documentation for specific information about these restrictions.

Passwords implications when upgrading or downgrading software versions

When you update software from a version that does not support long passwords to a version that does support long passwords, the existing usernames and passwords continue to be there and no further action is required.

Before downgrading to a software version that does not include this feature, use one of the following procedures:

Procedure

- Reset the username and/or password to be no more than 16 characters in length, without using any special characters, from the CLI command `password`.

  1. Execute a CLI `write memory` command (required if the `include-credentials` feature has ever been enabled.)

```
HP Switch(config)# password manager
New password: ********
Please retype new password: *******
HP Switch(config)# write mem
```

- Execute the CLI command `[no] password all`. This clears all the passwords.

  1. Execute a CLI `write memory` command (required if the `include-credentials` feature has ever been enabled.)

```
HP Switch(config)# no password all
Password protections will be deleted, do you want to
continue [y/n]? y
HP Switch(config)# write mem
```

- Clear the password by using the **Clear** button on the switch.

  1. Execute a CLI `write memory` command (required if the `include-credentials` feature has ever been enabled.)

Unable to use previous password

If you cannot access the switch after a software version downgrade, clear the password by using the **Clear** button on the switch to regain access. Then boot into a software version that supports long passwords, and perform one of the procedures in **Passwords implications when upgrading or downgrading software versions** .

Setting passwords and usernames (WebAgent)

In the WebAgent you can enter passwords and (optional) usernames. See the WebAgent online help for detailed information.

Configure SHA-256 format passwords

Syntax

```
password non-plaintext-sha256

[no] password non-plaintext-sha256
```

Description

Configure the password in SHA-256 format.

Limitations

- After `password non-plaintext-sha256` is executed, the password cannot be converted back to plaintext; you must reconfigure the password.

- This feature is not applicable for passwords used in protocol handshaking (for example, SNMPv3, OSPF, and BFD).

- The SHA-256 password format is not supported when the password complexity feature is enabled.

- If the passwords in the configuration are in SHA-256 format, downgrading to a version where this feature is not supported results in the deletion of the passwords. HPE recommends that you disable this feature and reconfigure the password before downgrading.

- If the `password non-plaintext-sha256` feature is enabled, you are not allowed to enter the password in SHA-1 format.

The following three tables show the output from the `show running-config` command for each password storage format.

Table 3: *Passwords configured using the plaintext option*

include credentials enabled	encrypt-credentials enabled	non-plaintext-sha256 enabled	`show running-config` output (manager/operator/local-user)
No	No	No	password manager password operator aaa authentication local-user <username> group <groupname>
No	No	Yes	Manager and operator credentials are not displayed. aaa authentication local-user <username> group <groupname>
No	Yes	No	password manager password operator aaa authentication local-user <username> group <groupname>
No	Yes	Yes	Manager and operator credentials are not displayed. aaa authentication local-user <username> group <groupname>

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Yes	No	No	password manager user-name \<username\> \<SHA-1 password\> password manager user-name \<username\> \<SHA-1 password\> aaa authentication local-user \<username\> group \<groupname\> password sha1 \<SHA-1 password\>
Yes	No	Yes	password manager user-name \<username\>sha256 \<SHA-256 password\> password manager user-name \<username\>sha256 \<SHA-256 password\> aaa authentication local-user \<username\> group \<groupname\> password \<SHA-256 password\>
Yes	Yes	No	encrypted-password manager user-name \<username\> \<encrypted SHA-1 password\> encrypted-password manager user-name \<username\> \<encrypted SHA-1 password\> aaa authentication local-user \<username\> group \<groupname\> password sha1 \<SHA-1 password\>
Yes	Yes	Yes	encrypted-password manager user-name \<username\> \<encrypted SHA-256 password\> encrypted-password manager user-name \<username\> \<encrypted SHA-256 password\> aaa authentication local-user \<username\> group \<groupname\> password sha 256 \<SHA-256 password\>

Table 4: *Passwords configured using the sha1 option*

include credentials enabled	encrypt-credentials enabled	non-plaintext-sha256 enabled	`show running-config` output (manager/operator/local-user)
Yes	No	No	password manager user-name \<username\> sha-1 \<SHA-1 password\> password operator user-name \<username\> sha-1 \<SHA-1 password\> aaa authentication local-user \<username\> group \<groupname\> password sha1 \<SHA-1 password\>
Yes	No	Yes	Passwords cannot be configured using the sha1 option when non-plaintext sha256 is enabled.

Table Continued

Yes	Yes	No	encrypted-password manager user-name <username> <encrypted SHA-1 password> encrypted-password manager user-name <username> <encrypted SHA-1 password> aaa authentication local-user <username> group <groupname> password sha1 <SHA-1 password>
Yes	Yes	Yes	Passwords cannot be configured using the sha1 option when non-plaintext sha256 is enabled.

Table 5: *Passwords configured using the sha256 option*

include credentials enabled	encrypt-credentials enabled	non-plaintext-sha256 enabled	`show running-config` output (manager/operator/local-user)
Yes	No	No	Manager and operator credentials are not displayed because SHA-1 passwords are not available. aaa authentication local-user <username> group <groupname>
Yes	No	Yes	password manager user-name <username> sha256 <SHA-256 password> password manager user-name <username> sha256 <SHA-256 password> aaa authentication local-user <username> group <groupname> password sha 256 <SHA-256 password>
Yes	Yes	No	Manager and operator credentials are not displayed because SHA-1 passwords are not available. aaa authentication local-user <username> group <groupname>
Yes	Yes	Yes	encrypted-password manager user-name <username> <encrypted SHA-256 password> encrypted-password manager user-name <username> <encrypted SHA-256 password> aaa authentication local-user <username> group <groupname> password sha 256 <SHA-256 password>

Saving security credentials in a config file

You can store and view the following security settings in the running-config file associated with the current software image by entering the `include-credentials` command (formerly this information was stored only in internal flash memory):

• Local manager and operator passwords and (optional) usernames that control access to a management session on the switch through the CLI, menu interface, or WebAgent.

• SNMP security credentials used by network management stations to access a switch, including authentication and privacy passwords.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

- Port-access passwords and usernames used as 802.1X authentication credentials for access to the switch.

- TACACS+ encryption keys used to encrypt packets and secure authentication sessions with TACACS+ servers keys.

- RADIUS shared secret (encryption) keys used to encrypt packets and secure authentication sessions with RADIUS servers.

- Secure Shell (SSH) public keys used to authenticate SSH clients that try to connect to the switch.

Benefits of saving security credentials

The benefits of including and saving security credentials in a configuration file are:

- After making changes to security parameters in the running configuration, you can experiment with the new configuration and, if necessary, view the new security settings during the session. After verifying the configuration, you can then save it permanently by writing the settings to the startup-config file.

- By permanently saving a switch security credentials in a configuration file, you can upload the file to a TFTP server or Xmodem host, and later download the file to the switches on which you want to use the same security settings without having to manually configure the settings (except for SNMPv3 user parameters) on each switch.

- By storing different security settings in different files, you can test different security configurations when you first download a new software version that supports multiple configuration files, by changing the configuration file used when you reboot the switch.

For more information about how to experiment with, upload, download, and use configuration files with different software versions, see:

- "Switch Memory and Configuration" in the management and configuration guide.

- **Configuring password security** on page 44.

Enabling the storage and display of security credentials

To enable the security settings, enter the `include-credentials` command.

Syntax:

```
[no] include-credentials [radius-tacacs-only|store-in-config]
```

Enables the inclusion and display of the currently configured manager and operator usernames and passwords, RADIUS shared secret keys, SNMP and 802.1X authenticator (port-access) security credentials, and SSH client public keys in the running configuration. (Earlier software releases store these security configuration settings only in internal flash memory and do not allow you to include and view them in the running-config file.)

To view the currently configured security settings in the running configuration, enter one of the following commands:

- `show running-config`

  : Displays the configuration settings in the current running-config file.

- `write terminal`

  : Displays the configuration settings in the current running-config file.

For more information, see "Switch Memory and Configuration" in the basic operation guide.

To view the current status of include-credentials on the switch, enter `show include-credentials`. See **Displaying the status of include-credentials** on page 62.

The `[no]` form of the command disables only the display and copying of these security parameters from the running configuration, while the security settings remain active in the running configuration.

Default: The security credentials described in **Security settings that can be saved** on page 54 are not stored in the running configuration.

`radius-tacacs-only`

When executed with the `radius-tacacs-only` option, only the RADIUS and TACACS security keys are included in the configuration when saving files remotely.

The `radius-tacacs-only` option can be disabled with either command

```
[no]include-credentials
```

```
[no]include-credentials radius-tacacs-only
```

`store-in-config:`

Stores passwords and SSH authorized keys in the configuration files. This happens automatically when `include-credentials` is enabled.

`[no]include-credentials store-in-config`

The `[no]include-credentials store-in-config` command disables `include-credentials` and removes credentials stored in the configuration files. The switch reverts to storing only a single set of passwords and SSH keys, regardless of which configuration file is booted.

Security settings that can be saved

The security settings that can be saved to a configuration file are:

- Local manager and operator passwords and usernames

- SNMP security credentials, including SNMPv1 community names and SNMPv3 usernames, authentication, and privacy settings

- 802.1X port-access passwords and usernames

- TACACS+ encryption keys

- RADIUS shared secret (encryption) keys

- Public keys of SSH-enabled management stations that are used by the switch to authenticate SSH clients that try to connect to the switch

Executing include-credentials or include-credentials store-in-config

When `include-credentials` or `include-credentials store-in-config` is executed for the first time (for example, on a new switch) or when you have successfully executed the `no include-credentials store-in-config` command, the passwords and SSH keys are not currently stored in the configuration file (not activated). The following example shows the caution message displayed.

Caution message

```
HP Switch(config)# include-credentials
```

```
                          **** CAUTION ****

You have invoked the command 'include-credentials'. This action will make changes
to the password and SSH public-key storage.

It will affect *all* stored configurations, which might need to be updated.
Those credentials will no longer be readable by older software revisions.
It also may break some of your existing user scripts. Continue?[y/n] y

Erasing configurations with 'include-credentials' enabled will erase stored
passwords and security credentials. The system will reboot with the factory
default configuration.

Proceed?[y/n]
```

This caution message can also appear if you have successfully executed the `no include-credentials store-in-config` command.

The no include-credentials store-in-config option

The `no include-credentials` command disables include-credentials. Credentials continue to be stored in the active and inactive configurations, but are not displayed in the config file.

When `no include-credentials` is used with the store-in-config option, `include-credentials` is disabled and the credentials stored in the config files are removed. The switch is restored to its default state and only stores one set of operator/manager passwords and SSH keys.

If you choose to execute the `no include-credentials store-in-config` command, you are also presented with the option of setting new switch passwords.

You are queried about retaining the current SSH authorized keys on the switch. If you enter "y", the currently active authorized key files are renamed to the pre-include-credentials names, for example:

`/file/mgr_auth_keys.2 -> /file/mgr_auth_keys`

`/file/authorized_keys.2 -> /file/authorized_keys`

All remaining authorized keys files with an extension are deleted.

The no include-credentials store-in-config messages and options

```
HP Switch(config)# no include-credentials store-in-config

This will remove any switch passwords and inactive SSH authorized keys from all
configuration files. This will also restore the functionality to store only a
single set of passwords and authorized keys on the switch.
Do you want to continue (y/n)? y

The SSH authorized keys associated with the active configuration will be deleted.
Would you like to retain these as the switch global SSH authorized keys (y/n)? y

Do you want to set new switch passwords (y/n)? y

operator username: admin
operator password: ********
Confirm password: ********
manager username: GeorgeV
manager password: ********
Confirm password: ********

HP Switch(config)#
```

Local manager and operator passwords

The information saved to the running-config file when the include-credentials command is entered includes:

```
password manager [user-name <name>] <hash-type> <pass-hash>
```

```
password operator [user-name <name>] <hash-type> <pass-hash>
```

where

`<name>` is an alphanumeric string for the user name assigned to the manager or operator.

`<hash-type>` indicates the type of hash algorithm used: SHA-1 or plain text.

`<pass-hash>` is the SHA-1 authentication protocol's hash of the password or clear ASCII text.

For example, a manager username and password can be stored in a `runningconfig` file as follows:

```
password manager user-name George SHA1
2fd4e1c67a2d28fced849ee1bb76e7391b93eb12
```

Use the `write memory` command to save the password configurations in the `startup-config` file. The passwords take effect when the switch boots with the software version associated with that configuration file.

> ⚠️ **CAUTION**
>
> If a startup configuration file includes other security credentials, but does not contain a manager or operator password, the switch will not have password protection and can be accessed through Telnet or the serial port of the switch with full manager privileges.

Password command options

The `password` command has the following options:

Syntax:

```
[no] password <manager|operator|port-access> [user-name <name>] <plaintext> <password>>
```

Set or clear a local username/password for a given access level.

`manager`:

Configures access to the switch with manager-level privileges.

`operator`:

Configures access to the switch with operator-level privileges.

`port-access`:

Configures access to the switch through 802.1X authentication with operator-level privileges.

`user-name <name>` :

The optional text string of the user name associated with the password.

`<password>` :

The clear ASCII text string of the password.

For more information about configuring local manager and operator passwords, see **Configuring Username and Password Security** on page 44.

SNMP security credentials

SNMPv1 community names and write-access settings, and SNMPv3 usernames continue to be saved in the running configuration file even when you enter the `include-credentials` command.

In addition, the following SNMPv3 security parameters are also saved:

```
snmpv3 user <Name> [auth [<md5 | sha>] <Auth-Pass>]
                   [priv [<des | aes>] <Priv-Pass>]
```

Parameters

<Name>

   The name of an SNMPv3 management station.

<md5 | sha>

   Selects the authentication method.

<Auth-Pass>

   Specifies the authentication password.

<des | aes>

   Selects the privacy encryption method.

<Priv-Pass>

   Specifies the privacy password used by a privacy protocol to encrypt SNMPv3 messages between the switch and the station.

The following example shows the additional security credentials for SNMPv3 users that can be saved in a `running-config` file.

Security credentials saved in the running-config

```
snmpv3 user boris \
auth md5 "9e4cfef901f21cf9d21079debeca453" \
priv "82ca4dc99e782db1a1e914f5d8f16824"

snmpv3 user alan \
auth sha "8db06202b8f293e9bc0c00ac98cf91099708ecdf" \
priv "5bc4313e9fd7c2953aaea9406764fe8bb629a538"
```

Although you can enter an SNMPv3 authentication or privacy password in either clear ASCII text or the SHA-1 hash of the password, the password is displayed and saved in a configuration file only in hashed format, as shown in the preceding example.

See "Configuring for Network Management Applications" in the management and configuration guide for your switch for more information about the configuration of SNMP security parameters.

TACACS+ encryption key authentication

You can use TACACS+ servers to authenticate users who request access to a switch through Telnet (remote) or console (local) sessions. TACACS+ uses an authentication hierarchy consisting of:

Remote passwords assigned in a TACACS+ server

Local manager and operator passwords configured on the switch.

When you configure TACACS+, the switch first tries to contact a designated TACACS+ server for authentication services. If the switch fails to connect to any TACACS+ server, it defaults to its own locally assigned passwords for authentication control if it has been configured to do so.

For improved security, you can configure a global or server-specific encryption key that encrypts data in TACACS + packets transmitted between a switch and a RADIUS server during authentication sessions. The key configured on the switch must match the encryption key configured in each TACACS+ server application. (The encryption key is sometimes referred to as "shared secret" or "secret" key.)

TACACS+ shared secret (encryption) keys can be saved in a configuration file by entering this command:

```
HP Switch(config)# tacacs-server key <keystring>
```

The option `<keystring>` is the encryption key (in clear text) used for secure communication with all or a specific TACACS+ server.

RADIUS shared-secret key authentication

You can use RADIUS servers as the primary authentication method for users who request access to a switch through Telnet, SSH, console, or port access (802.1X). The shared secret key is a text string used to encrypt data in RADIUS packets transmitted between a switch and a RADIUS server during authentication sessions. Both the switch and the server have a copy of the key; the key is never transmitted across the network.

RADIUS shared secret (encryption) keys can be saved in a configuration file by entering this command:

```
HP Switch(config)# radius-server key <keystring>
```

The option `<keystring>` is the encryption key (in clear text) used for secure communication with all or a specific RADIUS server.

The include-credentials radius-tacacs-only option

This option allows you to execute `include-credentials` for **only** RADIUS and TACACS. The `radius-tacacs-only` option does not cause the switch to store authentication passwords and SSH keys in the configuration file.

Syntax:

```
[no] include-credentials [radius-tacacs-only|store-in-config]
```

Enables the inclusion of passwords and security credentials in each configuration file when the file is saved onto a remote server or workstation. When `[no]include-credentials` is executed, include-credentials is disabled. Credentials continue to be stored in the active and inactive configuration files but are not displayed.

`radius-tacacs-only`:

When executed with the `radius-tacacs-only` option, only the RADIUS and TACACS security keys are included in the configuration when saving files remotely.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

The `radius-tacacs-only` option can be disabled with either command:

```
[no] include-credentials
```

```
[no] include-credentials radius-tacacs-only
```

`store-in-config`:

Stores passwords and SSH authorized keys in the configuration files. This happens automatically when `include-credentials` is enabled.

The `[no] include-credentials store-in-config` command disables the `include-credentials` command and removes credentials stored in the configuration files. The switch reverts to storing only a single set of passwords and SSH keys, regardless of which configuration file is booted.

When `include-credentials radius-tacacs-only` is executed, a warning message displays.

Caution message displayed for the radius-tacacs-only option

```
HP Switch(config)# include-credentials radius-tacacs-only
                          **** CAUTION ****
This will insert possibly sensitive information in switch configuration files,
and as a part of some CLI commands output. It is strongly recommended that you
use SFTP rather than TFTP for transfer of the configuration over the network,
and that you use the web configuration interface only with SSL enabled.

Erasing configurations with 'include-credentials' enabled will erase stored
passwords and security credentials. The system will reboot with the factory
default configuration.
```

SSH Re-Keying for SSH Server and SSH Client

To comply with RFC 4251, session rekeying ensures that either the SSH server or the SSH client initiates a rekey. This results in a new set of encryption and integrity keys to be exchanged between them. Once the rekey is complete, new keys are used for further communication, which ensures that the same key is not used for a long duration and the security of the session is maintained.

Syntax

```
ip ssh rekey {time <MINUTES> | volume <SIZE>}
```

```
no ip ssh rekey
```

Description

Enable SSH key re-exchange.

The no form of the command disables SSH rekeying. If time or volume is specified, the no form reverts to the default value.

Command context
`config`

Parameters

`time <MINUTES>`

Specifies the number of minutes for rekey initiation. Range: 10 to 60 minutes. Default: 60 minutes.

```
volume <SIZE>
```

Specifies the volume size in KB for rekey initiation. Range: 100 KB to 1048576 KB. Default: 1048576 KB.

Examples

Initiate rekeying every 45 minutes:

```
switch(config)# ip ssh rekey time 45
```

Reset the configured time to the default value (60 minutes):

```
switch(config)# no ip ssh rekey time
```

Initiate rekeying after every 2000 KB of data is transferred:

```
switch(config)# ip ssh rekey volume 2000
```

Reset the configured volume to the default value (1048576 KB):

```
switch(config)# no ip ssh rekey volume
```

SSH client public-key authentication

Secure Shell version 2 (SSHv2) is used by HPE switches to provide remote access to SSH-enabled management stations. Although SSH provides Telnet-like functions, unlike Telnet, SSH provides encrypted, two-way authenticated transactions. SSH client public-key authentication is one of the types of authentication used.

Client public-key authentication uses one or more public keys (from clients) that must be stored on the switch. Only a client with a private key that matches a public key stored on the switch can gain access at the manager or operator level. For more information about how to configure and use SSH public keys to authenticate SSH clients that try to connect to the switch, see **Configuring Secure Shell (SSH)** on page 353.

The SSH security credential that is stored in the running configuration file is configured with the `ip ssh public-key` command used to authenticate SSH clients for manager or operator access, along with the hashed content of each SSH client public key.

Syntax:

```
ip ssh public-key <manager|operator> keystring
```

Set a key for public-key authentication.

manager:

Allows manager-level access using SSH public-key authentication.

operator:

Allows operator-level access using SSH public-key authentication.

`keystring:`

> A legal SSHv2 (RSA or DSA) public key. The text string for the public key must be a single quoted token. If the keystring contains double-quotes, it can be quoted with single quotes (`'keystring'`). The following restrictions for a keystring apply.

> A keystring cannot contain both single and double quotes.

> A keystring cannot have extra characters, such as a blank space or a new line. However, to improve readability, you can add a backlash at the end of each line.

NOTE

The `ip ssh public-key` command allows you to configure only one SSH client public key at a time. The `ip ssh public-key` command behavior includes an implicit append that never overwrites existing public-key configurations on a running switch.

If you download a software configuration file that contains SSH client publickey configurations, the downloaded public keys overwrite any existing keys, as happens with any other configured values.

To display the SSH public-key configurations (72 characters per line) stored in a configuration file, enter the `show config` or `show running-config` command. The following example shows the SSH public keys configured for manager access, along with the hashed content of each SSH client public key, that are stored in a configuration file.

SSH public keys

```
...
include-credentials
ip ssh public-key manager "ssh-dss \
AAAAB3NzaC1kc3MAAACBAPwJHSJmTRtpZ9BUNC+ZrsxhMuZEXQhaDME1vc/ \
EvYnTKxQ31bWvr/bT7W58NX/YJ1ZKTV2GZ2QJCicUUZVWjNFJCsa0v03XS4 \
BhkXjtHhz6gD701otgizUOO6/Xzf4/J9XkJHkOCnbHIqtB1sbRYBTxj3NzA \
K1ymvIaU09X5TDAAAAFQCPwKxnbwFfTPasXnxfvDuLSxaC7wAAAIASBwxUP \
pv2scqPPXQghgaTkdPwGGtdFW/+K4xRskAnIaxuG0qLbnekohi+ND4TkKZd \
EeidgDh7qHusBhOFXM2g73RpE2rNqQnSf/QV95kdNwWIbxuusBAzvfaJptd \
gca6cYR4xS4TuBcaKiorYj60kk144E1fkDWieQx8zABQAAAIEAu7/1kVOdS \
G0vE0eJD23TLXvu94plXhRKCUAvyv2UyK+piG+Q1el1w9zsMaxPA1XJzSY/ \
imEp4p6WXEMcl0lpXMRnkhnuMMpaPMaQUT8NJTNu6hqf/LdQ2kqZjUuIyV9 \
LWyLg5ybS1kFLeOt0oo2Jbpy+U2e4jh2Bb77sX3G5C0= spock@sfc.gov" \
ip ssh public-key manager 'ssh-rsa \
AAAAB3NzaC1yc2EAAAADAQABAAAAgQDyO9RDD52JZP8k2F2YZXubgwRAN0R \
JRs1Eov6y1RK3XkmgVatzl+mspiEmPS4wNK7bX/IoXNdGrGkoE8tPkxlZOZ \
oqGCf5Zs50P1nkxXvAidFs55AWqOf4MhfCqvtQCe1nt6LFh4ZMig+YewgQG \
M6H1geCSLUbXXSCipdPHysakw== "TectiaClientKey [1024-bit rsa, \
nobody@testmachine, Mon Aug 15 2005 14:47:34]"'
ip ssh public-key manager "ssh-rsa \
AAAAB3NzaC1yc2EAAABIwAAAIEA1Kk9sVQ9LJOR6XO/hCMPxbiMNOK8C/ay \
+SQ10qGw+K9m3w3TmCfjh0ud9hivgbFT4F99AgnQkvm2eVsgoTtLRnfF7uw \
NmpzqOqpHjD9YzItUgSK1uPuFwXMCHKUGKa+G46A+EWxDAIypwVIZ697QmM \
qPFj1zdI4sIo5bDett2d0= joe@hp.com"
...
```

If a switch configuration contains multiple SSH client public keys, each public key is saved as a separate entry in the configuration file. You can configure up to 10 SSH client public keys on a switch.

X.509v3 certificate authentication for SSH

This feature supports user-authentication in SSH using X.509v3-based certificates.

Syntax

```
aaa authentication ssh {enable | login | client} <primary-method> [<backup-method>]
```

Description

Configure the authentication mechanism used to control SSH access to the switch. The X.509 certificate authentication for the SSH server works only when both `enable` and `login` options are configured to use `certificate` as the primary authentication method.

Options

- `enable`: Configure access to the privileged mode commands.

- `login`: Configure login access to the switch.

- `client`: Configure SSH client authentication for the switch.

Example

Use the X.509 certificate for SSH client authentication. To disable this feature, use `none` as the primary authentication method.

```
switch# aaa authentication ssh client certificate none
```

Displaying the status of include-credentials

The `show include-credentials` command provides the current status of include-credentials on the switch.

Syntax:

```
show include-credentials
```

Displays information about the passwords and SSH keys stored in the configuration.

Stored in configuration — yes:

The passwords and SSH keys are stored in the configuration. Include-credentials was executed.

Stored in configuration — no:

There is only one set of operator/manager passwords and one set of SSH keys for the switch.

Enabled in active configuration:

`Include-credentials` is either enabled or disabled.

RADIUS/TACACS only:

Displayed when the option is configured.

Output for show include credentials command

```
HP Switch(config)# show include-credentials

Stored in Configuration          : Yes
Enabled in Active Configuration  : N/A
RADIUS/TACACS Only               : Yes
```

Storage states when using include-credentials

The following table shows the states of several access types when the factory default settings are in effect or when `include-credentials` is enabled or not enabled.

Table 6: *Switch storage states*

Type	Factory default	Include-credentials enabled	Include-credentials disabled but active	No include-credentials executed (factory default)
Manager/Operator passwords & port access	• Single set for switch • Not displayed in config • Not copied with config to file server	• One set per stored config • Displayed in config • Copied with the config to file server	• One set per stored config • Not displayed in config • Not copied with config to file server	• Single set for switch • Not displayed in config • Not copied with config to file server
SSH Public Key	• Single set for switch • Not displayed in config • Not copied with config to file server	• One set per stored config • Displayed in config • Copied with the config to file server	• One set per stored config • Not displayed in config • Not copied with config to file server	• Single set for switch • Not displayed in config • Not copied with config to file server
SNMPv3 auth and priv	• One set per stored config • Not displayed in config • Not copied with config to file server	• One set per stored config • Displayed in config • Copied with the config to file server	• One set per stored config • Not displayed in config • Not copied with config to file server	• One set per stored config • Not displayed in config • Not copied with config to file server
SNTP auth	• One set per stored config • Not displayed in config • Not copied with config to file server	• One set per stored config • Displayed in config • Copied with the config to file server	• One set per stored config • Not displayed in config • Not copied with config to file server	• One set per stored config • Not displayed in config • Not copied with config to file server

Table Continued

Type	Factory default	Include-credentials enabled	Include-credentials disabled but active	No include-credentials executed (factory default)
RADIUS & TACACS keystrings	• One set per stored config • Displayed in config • Not copied with config to file server	• One set per stored config • Displayed in config • Not copied with config to file server	• One set per stored config • Not displayed in config • Not copied with config to file server	• One set per stored config • Displayed in config • Not copied with config to file server

> **NOTE**
> • When the `no include-credentials store-in-config` command is executed, the switch is restored to its default state and only stores one set of operator/manager passwords and SSH keys.
>
> • In the Factory Default state, the comments 'password operator' and/or 'password manager' are displayed to indicate passwords are configured. No indication is displayed for the state Include-Credentials disabled but Active.

Operating notes

⚠️ **CAUTION**

When you first enter the `include-credentials` command to save the additional security credentials to the running configuration, these settings are moved from internal storage on the switch to the running-config file.You are prompted by a warning message to perform a `write memory` operation to save the security credentials to the startup configuration. The message reminds you that if you do not save the current values of these security settings from the running configuration, they will be lost the next time you boot the switch and will revert to the values stored in the startup configuration.

When you boot a switch with a startup configuration file that contains the `include-credentials` command, any security credentials that are stored in internal flash memory are ignored and erased. The switch will load only the security settings in the startup configuration file.

Security settings are no longer automatically saved internally in flash memory and loaded with the startup configuration when a switch boots up. The configuration of all security credentials requires that you use the `write memory` command to save them in the startup configuration in order for them to not be lost when you log off. A warning message reminds you to permanently save a security setting.

• After you enter the `include-credentials` command, the currently configured manager and operator usernames and passwords, RADIUS shared secret keys, SNMP and 802.1X authenticator (port-access) security credentials, and SSH client public keys are saved in the running configuration.Use the `no include-credentials` command to disable the display and copying of these security parameters from the running configuration using the `show running-config` and `copy running-config` commands without disabling the configured security settings on the switch. After you enter the `include-credentials` command, you

can toggle between the non-display and display of security credentials in `show` and `copy` command output by alternately entering the `no include-credentials` and `include-credentials` commands.

- After you permanently save security configurations to the current startup-config file using the `write memory` command, you can view and manage security settings with the following commands:

    ◦ `show config`

      : Displays the configuration settings in the current startup-config file.

    ◦ `copy config <source-filename> config <target-filename>`

      : Makes a local copy of an existing startup-config file by copying the contents of the startup-config file in one memory slot to a new startup-config file in another, empty memory slot.

    ◦ `copy config tftp`

      : Uploads a configuration file from the switch to a TFTP server.

    ◦ `copy tftp config`

      : Downloads a configuration file from a TFTP server to the switch.

    ◦ `copy config xmodem`

      : Uploads a configuration file from the switch to an Xmodem host.

    ◦ `copy xmodem config`

      : Downloads a configuration file from an Xmodem host to the switch.

    For more information, see "Transferring startup-config files to or from a remote server" in the management and configuration guide.

- The switch can store up to three configuration files. Each configuration file contains its own security credentials and these security configurations can differ. It is the responsibility of the system administrator to ensure that the appropriate security credentials are contained in the configuration file that is loaded with each software image and that all security credentials in the file are supported.

- If you have already enabled the storage of security credentials (including local manager and operator passwords) by entering the `include credentials` command, the `reset-on-clear` option is disabled. When you press the **Clear** button on the front panel, the manager and operator usernames and passwords are deleted from the running configuration. However, the switch does not reboot after the local passwords are erased. (The `reset-on-clear` option normally reboots the switch when you press the **Clear** button.)For more in formation, see **Configuring front panel security** on page 71.

Restrictions on enabling security credentials

The following restrictions apply when you enable security credentials to be stored in the running configuration with the `include-credentials` command:

- The private keys of an SSH host cannot be stored in the running configuration. Only the public keys used to authenticate SSH clients can be stored. An SSH host's private key is only stored internally, for example, on the switch or on an SSH client device.

- SNMPv3 security credentials saved to a configuration file on a switch cannot be used after downloading the file on a different switch. The SNMPv3 security replaceables in the file are only supported when loaded on the same switch for which they were configured. This is because when SNMPv3 security credentials are saved to a configuration file, they are saved with the engine ID of the switch as shown here:

`snmpv3 engine-id 00:00:00:0b:00:00:08:00:09:01:10:01`

If you download a configuration file with saved SNMPv3 security credentials on a switch, when the switch loads the file with the current software version the SNMPv3 engine ID value in the downloaded file must match the engine ID of the switch in order for the SNMPv3 users to be configured with the authentication and privacy passwords in the file. (To display the engine ID of a switch, enter the `show snmpv3 engine-id` command. To configure authentication and privacy passwords for SNMPv3 users, enter the `snmpv3 user` command.)If the engine ID in the saved SNMPv3 security settings in a downloaded configuration file does not match the engine ID of the switch:

○ The SNMPv3 users are configured, but without the authentication and privacy passwords. You must manually configure these passwords on the switch before the users can have SNMPv3 access with the privileges you want.

○ Only the `snmpv3 user`**< user_name>** credentials from the SNMPv3 settings in a downloaded configuration file are loaded on the switch, for example:

```
snmpv3 user boris
snmpv3 user alan
```

You can store 802.1X authenticator (port access) credentials in a configuration file. However, 802.1X supplicant credentials cannot be stored.

Encrypting credentials in the configuration file

A security risk is present when credentials used for authentication to remote devices such as RADIUS or TACACS + servers are displayed in the configuration file in plain text. The `encrypt-credentials` command allows the storing, displaying, and transferring of credentials in encrypted form.

When the encrypt-credentials feature is enabled, the affected credentials will be encrypted using aes-256-cbc encryption. By default, a fixed, hard-coded 256-bit key that is common to all networking devices is used. This allows transfer of configurations with all relevant credentials and provides much more security than plaintext passwords in the configuration.

Additionally, you can set a separate, 256-bit pre-shared key, however, you must now set the pre-shared key on the destination device before transferring the configuration. The pre-shared key on the destination device must be identical to the pre-shared key on the source device or the affected security credentials will not be usable. This key is only accessible using the CLI, and is not visible in any file transfers.

> **NOTE** It is expected that plaintext passwords will continue to be used for configuring the switch. The encrypted credentials option is available primarily for the backup and restore of configurations.

Only the aes-256-cbc encryption type is available.

Enabling encrypt-credentials

To enable `encrypt-credentials`, enter this command.

Syntax:

```
[no] encrypt-credentials [pre-shared-key <plaintext|hex>]
```

When `encrypt-credentials` is enabled without any parameters, it enables the encryption of relevant security parameters in the configuration.

The `[no]` form of the command disables the `encrypt-credentials` feature. If specified with `pre-shared-key` option, clears the `preshared-key` used to encrypt credentials.

`pre-shared-key`: When specified, sets the pre-shared-key that is used for all AES encryption. If no key is set, an switch default AES key is used.

Default: switch default AES key

`plaintext`: Set the key using plaintext.

`hex`: Set the key as a 64 hexadecimal character string (32 bytes). You must enter 64 hexadecimal digits to set this key.

When `encrypt-credentials` is enabled without any parameters, a caution message displays advising you about the effect of the feature with prior software versions, and actions that are recommended. All versions of the command force a configuration `save` after encrypting or re-encrypting sensitive data in the configuration.

Enabling encrypt-credentials with caution message

```
HP Switch(config)# encrypt-credentials

                     **** CAUTION ****

This will encrypt all passwords and authentication keys.

The encrypted credentials will not be understood by older software versions.
The resulting config file cannot be used by older software versions.
It may also break some of your existing user scripts.

Before proceeding, please save a copy of your current config file, and associate
the current config file with the older software version saved in flash memory.
See "Best Practices for Software Updates" in the Release Notes.

A config file with 'encrypt-credentials' may prevent previous software versions
from booting. It may be necessary to reset the switch to factory defaults. To
prevent this, remove the encrypt-credentials command or use an older config file.

Save config and continue [y/n]? y
```

Creating a pre-shared-key in plaintext

```
HP Switch(config)# encrypt-credentials pre-shared-key plaintext SecretKey1

Save config and continue [y/n]? y
```

Creating a pre-shared key in hex

```
HP Switch(config)# encrypt-credentials pre-shared-key hex
1234567891234567891234567891234567891234567891234567891
Save config and continue [y/n]? y
```

Displaying the state of encrypt-credentials

To display whether `encrypt-credentials` is enabled or disabled, enter the `show encrypt-credentials` command. This command is available only from the manager context.

Status of encrypt-credentials when the pre-shared key has not been set

```
HP Switch(config)# show encrypt-credentials
```

```
Encryption    : Disabled
Pre-shared Key: None
```

Status of encrypt-credentials when the pre-shared key has been set

```
HP Switch(config)# show encrypt-credentials

Encryption    : Disabled
Pre-shared Key:
055d7b3b6bc22d18d29533ba2b549b3991bc23b7cbfc8e5769bdcc9ec748af27
```

Affected commands

Several commands will have encryption available for configuration.

Table 7: *Affected commands*

Existing command	New equivalent option
`HP Switch(config)# radius-server key secret1`	`HP Switch(config)# radius-server encrypted-key U2FsdGVkX18XWadTeFN +bxHxKA/q+s5cV1NiYvx+TuA=`
`HP Switch(config)# radius-server host 10.0.0.1 key secret1`	`HP Switch(config)# radius-server host 10.0.0.1 encrypted-key U2FsdGVkX18XWadTeFN+bxHxKA q+s5cV1NiYvx +TuA=`
`HP Switch(config)# tacacs-server key secret1`	`HP Switch(config)# tacacs-server encrypted-key U2FsdGVkX18XWadTeFN +bxHxKA/q+s5cV1NiYvx+TuA=`
`HP Switch(config)# tacacs-server host 10.0.0.1 key secret1`	`HP Switch(config)# tacacs-server host 10.0.0.1 encrypted-key U2FsdGVkX18XWadTeFN+bxHxKA/ q+s5cV1NiYvx +TuA=`
`HP Switch(config)# key-chain example key 1 key-string secret1`	`HP Switch(config)# key-chain example key 1 encrypted-key U2FsdGVkX18XWadTeFN +bxHxKA/ q+s5cV1NiYvx+TuA=`
`HP Switch(config)# aaa port-access supplicant 24 secret secret1`	`HP Switch(config)# aaa port-access supplicant 24 identity id1 encrypted-secret secret1 U2FsdGVkX18XWadTeFN +bxHxKA/q+s5cV1NiYvx+TuA=`
`HP Switch(config)# sntp authentication key-id 33 authentication-mode md5 key-value secret1`	`HP Switch(config)# sntp authentication key-id 33 authentication-mode md5 encrypted-key U2FsdGVkX18XWadTeFN +bxHxKA/q+s5cV1NiYvx+TuA=`
`HP Switch(config)# password manager plaintext secret1`	`HP Switch(config)# encrypted-password manager U2FsdGVkX18XWadTeFN+bxHxKA/q +s5cV1NiYvx+TuA=`

Important operating notes

Hewlett Packard Enterprise recommends that when executing an `encrypted-<option>` command, you copy and paste the encrypted parameter from a known encrypted password that has been generated on the same switch or another switch with the same pre-shared key (whether user-specified or a default key). If an incorrectly encrypted parameter is used, it is highly likely that the decrypted version will contain incorrect characters, and neither key will function correctly or be displayed in any `show` command.

Interaction with include-credentials settings

The following table shows the interaction between `include-credentials` settings and `encrypt-credentials` settings when displaying or transferring the configuration.

Table 8: *Interactions between credential settings*

Include-credentials active	Include-credentials enabled	Encrypt-credentials enabled	Resulting behavior for sensitive data
			Hidden (default)
		Yes	Shown, encrypted
	Yes		n/a
	Yes	Yes	n/a
Yes			Hidden
Yes		Yes	Shown, encrypted
Yes	Yes		Shown, plaintext
Yes	Yes	Yes	Shown, encrypted

Front panel security

The front panel security features provide the ability to independently enable or disable some of the functions of the two buttons located on the front of the switch for clearing the password (**Clear** button) or restoring the switch to its factory default configuration (**Reset+Clear** buttons together). The ability to disable password recovery is also provided for situations which require a higher level of switch security.

The front-panel security features are designed to prevent malicious users from:

- Resetting the password(s) by pressing the **Clear** button
- Restoring the factory default configuration by using the **Reset+Clear** button combination.
- Gaining management access to the switch by having physical access to the switch itself

When security is important

Some customers require a high level of security for information. For example, the Health Insurance Portability and Accountability Act (HIPAA) of 1996 requires that systems handling and transmitting confidential medical records must be secure.

It used to be assumed that only system and network administrators would be able to get access to a network switch because switches were typically placed in secure locations under lock and key. For some customers this is no longer true. Others simply want the added assurance that even if someone did manage to get to the switch that data would still remain secure.

If you do not invoke front panel security on the switch, user defined passwords can be deleted by pushing the **Clear** button on the front panel. This function exists so that if customers forget the defined passwords they can still get back into the switch and reset the passwords. This does, however, leave the switch vulnerable when it is located in an area where non-authorized people have access to it. passwords could easily be cleared by pressing the **Clear** button. Someone who has physical access to the switch can be able to erase the passwords (and possibly configure new passwords) and take control of the switch.

As a result of increased security concerns, customers now have the ability to stop someone from removing passwords by disabling the **Clear** and/or **Reset** buttons on the front of the switch.

Front-panel button functions

Clear button

Pressing the **Clear** button alone for one second resets the password(s) configured on the switch.

Figure 3: *Clear button*
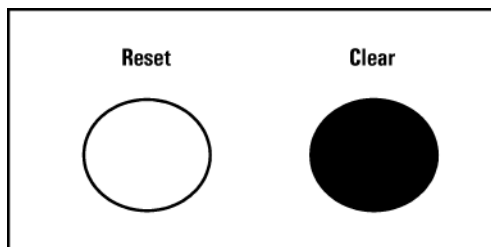
Reset button

Pressing the **Reset** button alone for one second causes the switch to reboot.

Figure 4: *Reset button*
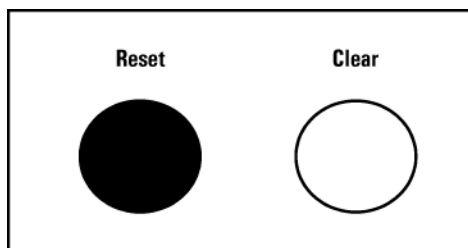
Restoring the factory default configuration

You can also use the **Reset** button **together** with the **Clear** button (**Reset+Clear**) to restore the factory default configuration for the switch.

Procedure

1. Press and hold the **Reset** button.

---
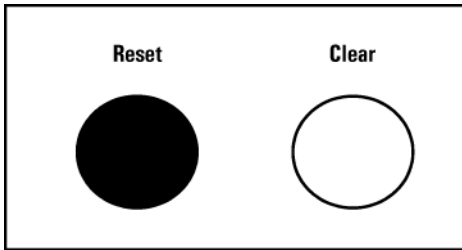
2. While holding the **Reset** button, press and hold the **Clear** button.
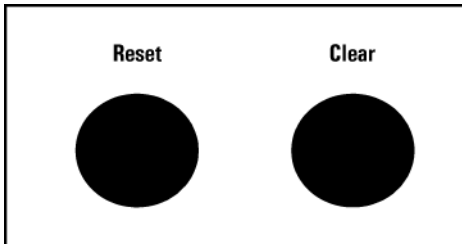
3. Release the **Reset** button.
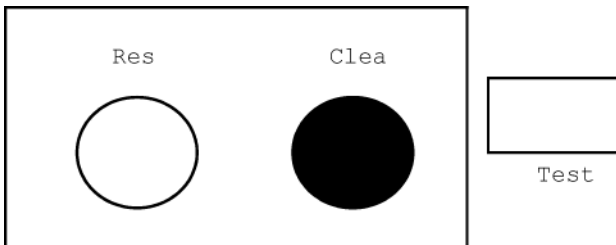
4. When the Test LED to the right of the **Clear** button begins flashing, release the **Clear** button.
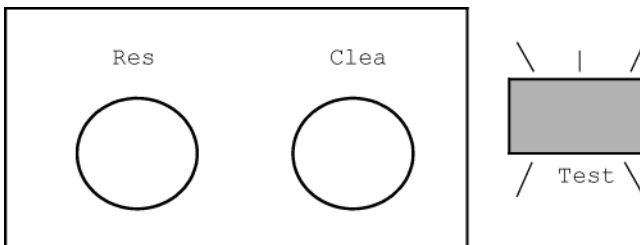
It takes approximately 20-25 seconds for the switch to reboot. This process restores the switch configuration to the factory default settings.

Configuring front panel security

Using the `front-panel-security` command from the global configuration context in the CLI you can:

- Disable or re-enable the password-clearing function of the **Clear** button. Disabling the **Clear** button means that pressing it does not remove local password protection from the switch. (This action affects the **Clear** button when used alone, but does not affect the operation of the **Reset+Clear** combination described under **Restoring the factory default configuration** on page 70.)

- Configure the **Clear** button to reboot the switch after clearing any local usernames and passwords. This provides an immediate, visual means (plus an Event Log message) for verifying that any usernames and passwords in the switch have been cleared.

- Modify the operation of the **Reset+Clear** combination so that the switch still reboots, but does **not** restore the switch's factory default configuration settings. (Use of the **Reset** button alone, to simply reboot the switch, is not affected.)

- Disable or re-enable Password Recovery.

Syntax:

`show front-panel-security`

Displays the current front panel security settings:

`clear password`:

Shows the status of the **Clear** button on the front panel of the switch. `Enabled` means that pressing the **Clear** button erases the local usernames and passwords configured on the switch (and thus removes local password protection from the switch.) `Disabled` means that pressing the **Clear** button does not remove the local usernames and passwords configured on the switch.

Default: Enabled

`reset-on-clear`:

Shows the status of the option `enabled` or `disabled`. When `reset-on-clear` is disabled and the command `clear password` is enabled, then pressing the **Clear** button erases the local usernames and passwords from the switch. When `reset-on-clear` command is enabled, pressing the **Clear** button erases the local usernames and passwords from the switch and reboots the switch. Enabling `reset-on-clear` automatically enables the `clear-password` command.

Default: Disabled.

> **NOTE**
> If you have stored security credentials (including the local manager and operator usernames and passwords) to the running config file by entering the `include-credentials` command, the `reset-on-clear` option is ignored. If you press the **Clear** button on the front panel, the manager and operator usernames and passwords are deleted from the startup configuration file, but the switch does not reboot.

`factory reset`:

Shows the status of the system **Reset** button on the front panel of the switch. Enabled means that pressing the system **Reset** button reboots the switch and also enables the system **Reset** button to be used with the **Clear** button. See **Restoring the factory default configuration** on page 70 to reset the switch to its factory-default configuration.

Default: Enabled.

`password recovery`:

Shows whether the switch is configured with the ability to recover a lost password. See **Password recovery** on page 75. Default: Enabled.

> **CAUTION**
> Disabling this option removes the ability to recover a password on the switch. Disabling this option is an extreme measure and is not recommended unless you have the most urgent need for high security. If you disable `password-recovery` and then lose the password, you will have to use the **Reset** and **Clear** buttons, see **Restoring the factory default configuration** on page 70 to reset the switch to factory default configuration and create a new password.

Executing the `show front-panel-security` command produces the following output when the switch is configured with the default front panel security settings:

The default front-panel security settings

```
HP Switch(config)# show front-panel-security
Clear Password       - Enabled
 Reset-on-clear      - Disabled
Factory Reset        - Enabled
```

Disabling the clear password function of the Clear button

Syntax:

[no] front-panel-security password-clear

In the factory default configuration, pressing the **Clear** button on the switch front panel erases any local usernames and passwords configured on the switch. This command disables the password clear function of the **Clear** button, so that pressing it has no effect on any local usernames and passwords.

For redundant management systems, this command only affects the active management module.

Default: `enabled`.

NOTE   Although the **Clear** button does not erase passwords when disabled, you can still use it with the **Reset** button, **Reset+Clear**, to restore the switch to its factory default configuration, as described under **Restoring the factory default configuration** on page 70.

This command displays a Caution message in the CLI. If you want to proceed with disabling the **Clear** button, type **[Y]**; otherwise type **[N]**. For example:

Figure 5: *Disabling the Clear button and displaying the new configuration*

```
HP Switch(config)# no front-panel-security password-clear
                        **** CAUTION ****
Disabling the clear button prevents switch passwords from being easily reset or
recovered. Ensure that you are familiar with the front panel security options
before proceeding.

Continue with disabling the clear button [y/n]? y

HP Switch(config)# show front-panel-security
Clear Password      - Disabled  ◄─────
Factory Reset       - Enabled
Password Recovery   - Enabled
```

Indicates the command has disabled the Clear button on the switch's front panel. In this case the Show command does not include the **reset-on-clear** status because it is inoperable while the Clear Password functionality is disabled, and must be reconfigured whenever Clear Password is re-enabled .

Re-enabling the Clear button and setting or changing the 'reset-on-clear' operation

Syntax:

[no] front-panel-security password-clear reset-on-clear

This command does both of the following:

• Re-enables the password clearing function of the **Clear** button on the switch front panel.

• Specifies whether the switch reboots if the **Clear** button is pressed.

To re-enable the password clear function, you must also specify whether to enable or disable the `reset-on-clear` option.

Defaults:

- `password-clear`: **Enabled**.

- `reset-on-clear`: **Disabled**.

-

To enable password-clear with reset-on-clear disabled:

Syntax:

`[no] front-panel-security password-clear reset-on-clear`

To enable password-clear with reset-on-clear also enabled:

Syntax:

`front-panel-security password-clear reset-on-clear`

Either form of the command **enables** password-clear.

For redundant management systems, this command only affects the active management module.

> **NOTE**
> If you disable `password-clear` and also disable the `password-recovery` option, you can still recover from a lost password by using the **Reset+Clear** button combination at reboot. Although the Clear button does not erase passwords when disabled, you can still use it with the **Reset** button (**Reset+Clear**) to restore the switch to its factory default configuration. You can then get access to the switch to set a new password.

For example, suppose `password-clear` is disabled and you want to restore it to its default configuration (enabled, with `reset-on-clear` disabled).

Figure 6: *Re-enabling the Clear button' default operation*
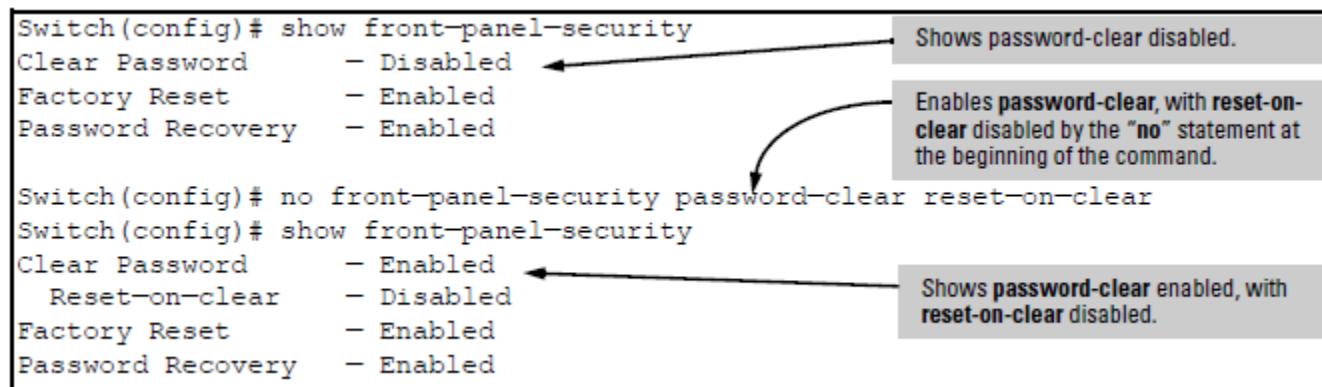
Changing the operation Reset+Clear combination

In their default configuration, using the **Reset+Clear** buttons in the combination described under **Restoring the factory default configuration** on page 70 replaces the switch current `startup-config` file with the factory default `startup-config` file, then reboots the switch and removes local password protection.

> **CAUTION**
> Anyone with physical switch access can use this button combination to replace the switch current configuration with the factory default configuration, and render the switch accessible without the need to input a username or password.

You can use the `factory-reset` command to prevent the **Reset+Clear** combination from being used for this purpose.

Syntax:

```
[no] front-panel-security factory-reset
```

Disables or re-enables the following functions associated with using the **Reset+Clear** buttons in the combination described under **Restoring the factory default configuration** on page 70:

- Replacing the current `startup-config` file with the factory default `startup-config` file

- Clearing any local usernames and passwords configured on the switch

Default: Both functions enabled.

For redundant management systems, this command only affects the active management module.

> **NOTE**
> The **Reset+Clear** button combination always reboots the switch, regardless of whether the `[no]` form of the command has been used to disable the above two functions. Also, if you disable `factory-reset`, you cannot disable the `password-recovery` option, and the reverse.

Figure 7: *Example of disabling the factory reset option*

Password recovery

The password recovery feature is enabled by default and provides a method for regaining management access to the switch (without resetting the switch to its factory default configuration) in the event that the system administrator loses the local manager username (if configured) or password. Using password recovery requires:

- `password-recovery`

  enabled (the default) on the switch prior to an attempt to recover from a lost username/password situation

- Contacting your Hewlett Packard Enterprise Customer Care Center to acquire a one-time-use password

Disabling or re-enabling the password recovery process

Disabling the password recovery process means that the only method for recovering from a lost manager username (if configured) and password is to reset the switch to its factory-default configuration, removing any non-default configuration settings.

> **CAUTION**
> Disabling password-recovery requires that factory-reset be enabled, and locks out the ability to recover a lost manager username and password on the switch. In this event, there is no way to recover from a lost manager username/password situation without resetting the switch to its factory default configuration. This can disrupt network operation and make it necessary to temporarily disconnect the switch from the network to prevent unauthorized access and other problems while it is being reconfigured. Also, with factory-reset enabled, unauthorized users can use the **Reset +Clear** button combination to reset the switch to factory default configuration and gain management access to the switch.

Syntax:

```
[no] front-panel-security password-recovery
```

Enables or disables the ability to recover a lost password.

When enabled, the switch allows management access through the password recovery process described below. This provides a method for recovering from lost manager usernames and passwords.

When disabled, the password recovery process is disabled and the only way to regain management access to the switch is to use the **Reset+Clear** button combination. See **Restoring the factory default configuration** on page 70 to restore the switch to its factory default configuration.

Default: Enabled.

> **NOTE**
> To disable password-recovery:
>
> • You must have physical access to the front panel of the switch.
>
> • Factory-reset must be enabled (the default).

Steps for disabling password-recovery:

Procedure

1. Use `show front-panel-security` to determine whether the factory-reset parameter is enabled. If it is disabled, use the `front-panel-security factory-reset` command to enable it.

2. Press and release the **Clear** button on the front panel of the switch.

3. Within 60 seconds of pressing the **Clear** button, enter the following command: `front-panel-security password-recovery`

4. Do one of the following after the `CAUTION` message appears:

   a. If you want to complete the command, press **[Y]**.

   b. If you want to abort the command, press **[N]**.

Example of disabling password recovery:

```
Switch(config)# no front-panel-security password-recovery
                      **** CAUTION ****
```

```
Disabling the clear button without password recovery prevents switch passwords
from being reset. If the switch password is lost, restoring the default factory
configuration will be required to regain access!

Continue with disabling password recovery [y/n]?
```

Password recovery process

If you lose the manager username/password, but password-recovery is enabled, then you can use the password recovery process to gain management access to the switch with an alternate password supplied by Networking Support.

<table>
<tr>
<td>NOTE</td>
<td>Disabled password-recovery locks out the ability to recover a manager username/password pair on the switch. The only way to recover from this is to use the **Reset +Clear** button combination described under **Restoring the factory default configuration** on page 70. This disrupts network operation and necessitates temporarily disconnecting the switch from the network to prevent unauthorized access and other problems while it is being reconfigured.</td>
</tr>
</table>

To recover a lost password:

Procedure

1. Note the switch base MAC address. It is shown on the label located on the upper right front corner of the switch.

2. Contact Networking Support for further assistance. Using the switch MAC address, Networking Support will generate and provide a "one-time use" alternate password to gain management access to the switch. Once you gain access, configure a new password.

<table>
<tr>
<td>NOTE</td>
<td>The alternate password provided by Networking Support is valid only for a single login attempt. You cannot use the same "one-time-use" password if you lose the password a second time. Because the password algorithm is randomized based upon your switch MAC address, the password will change as soon as you use the "one-time-use" password provided by Networking Support.</td>
</tr>
</table>

Configuring connection-rate filtering

Viewing the connection-rate configuration

Use the following command to view the basic connection-rate configuration. To view connection-rate ACLs and/or any other switch configuration details, use `show config` or `show running`.

Syntax

`show connection-rate-filter`

Displays the current global connection-rate status (enabled/disabled) and sensitivity setting, and the current per-port configuration. This command does not display the current (optional) connection-rate ACL configuration.
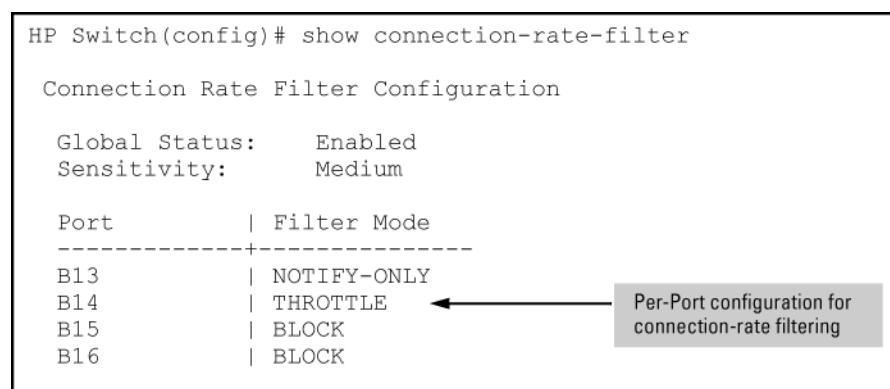
Figure 8: *Displaying the connection-rate status, sensitivity, and per-port configuration*

```
HP Switch(config)# show connection-rate-filter

 Connection Rate Filter Configuration

  Global Status:    Enabled
  Sensitivity:      Medium

  Port          | Filter Mode
  -------------+---------------
  B13           | NOTIFY-ONLY
  B14           | THROTTLE          <----  Per-Port configuration for
  B15           | BLOCK                    connection-rate filtering
  B16           | BLOCK
```

To view the complete connection-rate configuration, including any ACLs, use `show config` (for the `startup-config` file) or `show running` (for the `running-config` file). See **Applying connection-rate ACLs** on page 87.

Figure 9: *Connection-rate filtering configuration in the startup-config file*

```
HP Switch(config)# show config
Startup configuration
; J8697A Configuration Editor; Created on
hostname "HP Switch"
connection-rate-filter sensitivity medium
ip access-list connection-rate-filter "Sample"
    filter ip 13.28.234.180 0.0.15.255
    ignore ip 0.0.0.0 255.255.255.255
    exit
module 2 type J8161A
module 4 type J8161A
ip routing
logging 13.28.234.180
snmp-server community "public" Unrestricted
vlan 1
    name "DEFAULT_VLAN"
    untagged B1-B12, B19-B24, D1-D24
    no ip address
    no untagged B13-B18
    ip proxy-arp
    exit
vlan 15
    name "VLAN_15"
    untagged B13-B18
    ip address 13.28.234.181 255.255.240.0
    ip proxy-arp
    ip connection-rate-filter-access-group "Sample"
    exit
filter connection-rate B13 notify-only
filter connection-rate B14 throttle
filter connection-rate B16-B16 block
```

Entry showing that connection-rate-filtering is enabled and set to "medium" sensitivity.

Example of a connection-rate filtering ACL appearing in the configuration.

Example of a connection-rate filtering ACL appearing in a VLAN configuration.

Example of per-port connection-rate filtering policies appearing in the configuration.

Enabling global connection-rate filtering and sensitivity

Use the commands in this section to enable connection-rate filtering on the switch and to apply the filtering on a per-port basis.

Syntax

`connection-rate-filter sensitivity < low | medium | high | aggressive >`

`[no] connection-rate-filter`

This command:

- Enables connection-rate filtering.

- Sets the global sensitivity level at which the switch interprets a given host attempt to connect to a series of different devices as a possible attack by a malicious agent residing in the host.

Options for configuring sensitivity include:

low

>Sets the connection-rate sensitivity to the lowest possible sensitivity, which allows a mean of 54 destinations in less than 0.1 seconds, and a corresponding penalty time for Throttle mode (if configured) of less than 30 seconds.

medium

>Sets the connection-rate sensitivity to allow a mean of 37 destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 30 and 60 seconds.

high

>Sets the connection-rate sensitivity to allow a mean of 22 destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 60 and 90 seconds.

aggressive

>Sets the connection-rate sensitivity to the highest possible level, which allows a mean of 15 destinations in less than 1 second, and a corresponding penalty time for Throttle mode (if configured) between 90 and 120 seconds.

[no] connection-rate-filter

>This command disables connection-rate filtering on the switch.

> **NOTE** The sensitivity settings configured on the switch determine the Throttle mode penalty periods.

Configuring per-port filtering

Syntax

```
filter connection-rate <port-list> {notify-only | throttle | block}
```

```
no filter connection-rate <port-list>
```

Configures the per-port policy for responding to detection of a relatively high number of inbound IP connection attempts from a given source. The level at which the switch detects such traffic depends on the sensitivity setting configured by the `connection-rate-filter sensitivity` command.

> **NOTE** You can use connection-rate ACLs to create exceptions to the configured filtering policy.

The `no` form of the command disables connection-rate filtering on the ports in # `<port-list>`.

The `notify-only` option can be used if the switch detects a relatively high number of IP connection attempts from a specific host, `notify-only` generates an Event Log message and sends a similar message to any SNMP trap receivers configured on the switch.

The `throttle` command can be used if the switch detects a relatively high number of IP connection attempts from a specific host, this option generates the `notify-only` messaging and blocks all inbound traffic from the offending host for a penalty period. After the penalty period, the switch allows traffic from the offending host to resume, and re-examines the traffic. If the suspect behavior continues, the switch again blocks the traffic from the offending host and repeats the cycle.

The `block` command can be used if the switch detects a relatively high number of IP connection attempts from a specific host, this option generates the `notify-only` messaging and also blocks all inbound traffic from the offending host.

Table 9: *Throttle mode penalty periods*

Throttle mode (sensitivity)	Frequency of IP connection requests from the same source	Mean number of new destination hosts in the frequency period	Penalty period
Low	<0.1 second	54	<30 seconds
Medium	<1.0 second	37	30 - 60 seconds
High	<1.0 second	22	60 - 90 seconds
Aggressive	<1.0 second	15	90 - 120 seconds

Example of a Basic Connection-Rate Filtering Configuration

Figure 10: *Sample network*
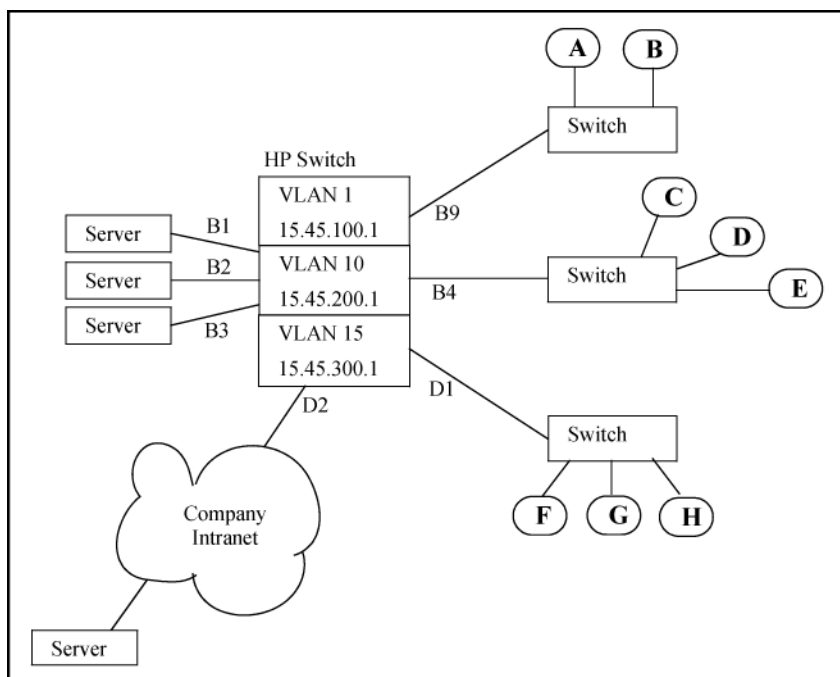
Basic configuration

Suppose that in the sample network, the administrator wanted to enable connection-rate filtering and configure the following response to high connection-rate traffic on the switch:

• Ports B1 — B3: Throttle traffic from the transmitting hosts.

• Port B4: Respond with notify-only to identify the transmitting hosts.

• Ports B9, D1, and D2: Block traffic from the transmitting hosts.

This example illustrates the configuration steps and resulting `startup-config` file:

```
HP Switch(config)# connection-rate-filter sensitivity low          ◄——  Enables connection-
                                                                         rate filtering and sets
HP Switch(config)# filter connection-rate b1-b3 throttle                 the sensitivity to "low".
HP Switch(config)# filter connection-rate b4 notify-only
HP Switch(config)# filter connection-rate b9, d1-d2 block
HP Switch(config)# write mem

HP Switch(config)# show config
Startup configuration

: J8697A Configuration Editor; Created on release #K.15.XX

hostname "HP Switch"                                               Configures the desired
connection-rate-filter sensitivity low          ◄————             responses to inbound, high
module 2 type 8702A                                               connectivity-rate traffic on the
module 4 type 8702A                                               various ports.
ip routing
snmp-server communitye "public" Unrestricted
snmp-server host 12.45.200.75 "public"
vlan 1
   name "DEFAULT_VLAN"
   untagged B5-B24
   ip address dhcp-bootp
   no untagged B1-B4, D1-D24
   ip proxy-arp
   exit
filter connection-rate B4 notify-only                             Shows the per-port configuration
filter connection-rate B1-B3 throttle           ◄————             for the currently enabled
filter connection-rate B9, D1-D2 block                            connectivity-rate filtering.
```

Blocked hosts

Listing currently-blocked hosts

Syntax

```
show connection-rate-filter < all-hosts | blocked-hosts | throttled-hosts >
```

```
all-hosts
```

Lists, by VLAN membership, all hosts currently detected in a throttling or blocking state, along with a state indicator.

```
blocked-hosts
```

Lists, by VLAN membership, the hosts currently blocked by connection-rate action.

```
throttled-hosts
```

Lists, by VLAN membership, the hosts currently in a throttling state due to connection-rate action.

Figure 11: *Example of listing hosts in any connection-rate state*

```
HP Switch(config)# show connection-rate-filter all-hosts

  VLAN ID         | Source IP Address | Filter Mode
  ------------+-------------------+-----------
  10              | 13.28.234.175     | THROTTLE
  10              | 13.28.234.179     | THROTTLE
  15              | 13.28.234.180     | BLOCK
```

Figure 12: *Example of listing hosts blocked by connection-rate filtering*

```
HP Switch(config)#show connection-rate-filter blocked-
hosts

  VLAN ID        | Source IP Address
  ------------+-------------------
```

Unblocking currently-blocked hosts

If a host becomes blocked by triggering connection-rate filtering on a port configured to block high connection rates, the host remains blocked on all ports on the switch even if you change the per-port filtering configuration. To help prevent a malicious host from automatically regaining access to the network, the source IP address block imposed by connection-rate filtering does not age-out.

When a host becomes blocked the switch generates a event log message and sends the message to any configured SNMP trap receivers. An example of an event log message is:

```
Src IP xxx.xxx.xxx.xxx blocked
```

NOTE

Before unblocking a host that was blocked by connection-rate filtering, Hewlett Packard Enterprise recommends inspecting the host with current antivirus tools and removing all potentially malicious agents.

If a trusted host frequently triggers connection-rate blocking with legitimate, high connection-rate traffic, consider either changing the sensitivity level on the associated port or configuring a connection-rate ACL to create a filtering exception for the host.

Syntax

```
connection-rate-filter unblock < all | host | ip-addr >
```

```
all
```

Unblocks all hosts currently blocked due to action by connection-rate filtering on ports where block mode has been configured.

```
host < ip-addr >
```

Unblocks the single host currently blocked due to action by connection-rate filtering on ports where block mode has been configured.

```
ip-addr < mask >
```

Unblocks traffic from any host in the specified subnet currently blocked due to action by connection-rate filtering on ports where block mode has been configured.

> **NOTE**
> There is also an option to unblock any host belonging to a specific VLAN using the `vlan <vid> connection-rate-filter unblock` command.

> **NOTE**
> For a complete list of options for unblocking hosts, see **Unblocking a currently blocked host** on page 93.

Configuring and applying connection-rate ACLs

Configuring a connection-rate ACL using source IP address criteria

Syntax

```
ip access-list connection-rate-filter < crf-list-name >
```

Creates a connection-rate-filter ACL and puts the CLI into the ACE context:

```
switch(config-crf-nacl)#
```

If the ACL already exists, this command simply puts the CLI into the ACE context.

Syntax

```
< filter | ignore > ip < any | host ip-addr | ip-addr mask-length >
```

Used in the ACE context to specify the action of the connection-rate ACE and the source IP address of the traffic that the ACE affects.

```
< filter | ignore >
```

The `filter` option assigns policy filtering to traffic with source IP address (SA) matching the source address in the ACE. The `ignore` option specifies bypassing policy filtering for traffic with an SA that matches the source address in the ACE.

```
ip < any | host ip-addr | ip-addr mask-length >
```

Specifies the SA criteria for traffic addressed by the ACE.

```
any
```

Applies the ACEs action (`filter` or `ignore`) to traffic having any SA.

```
host ip-addr
```

Applies the ACEs action (`filter` or `ignore`) to traffic having the specified host SA.

```
ip-addr mask-length
```

Applies the ACEs action (`filter` or `ignore`) to traffic having an SA within the range defined by either:

```
<src-ip-addr/cidr-mask-bits>
```

or

```
<src-ip-addr <mask>>
```

Use this criterion for traffic received from either a subnet or a group of IP addresses. The mask can be in either dotted-decimal format or CIDR format with the number of significant bits.

Configuring a connection-rate ACL using UDP/TCP criteria

Syntax

```
ip access-list connection-rate-filter crf-list-name
```

Creates a connection-rate-filter ACL and puts the CLI into the ACE context:

```
switch(config-crf-nacl)#
```

If the ACL already exists, this command simply puts the CLI into the ACE context.

Syntax

```
< filter | ignore | < udp | tcp > < any >
< filter | ignore | < udp | tcp > < host ip-addr > [ udp/tcp-options ]

< filter | ignore | < udp | tcp > ip-addr mask-length [ udp/tcp-options ]
```

Used in the ACE context (above) to specify the action of the connection-rate ACE (filter or ignore), and the UDP/TCP criteria and SA of the IP traffic that the ACE affects.

```
< filter | ignore >
```

```
filter
```

This option assigns a policy of filtering (dropping) IP traffic having an SA that matches the source address criteria in the ACE.

```
ignore
```

This option specifies a policy of allowing IP traffic having an SA that matches the source address criteria in the ACE.

```
< udp | tcp > < any | host > ip-addr | ip-addr mask-length
```

Applies the filter or ignore action to either TCP packets or UDP packets having the specified SA.

```
any
```

Applies the ACEs action (`filter` or `ignore`) to IP traffic having any SA.

```
host <ip-addr>
```

Applies the ACEs action (`filter` or `ignore`) to IP traffic having the specified host SA.

```
ip-addr <mask-length>
```

Applies the ACEs action (`filter` or `ignore`) to IP traffic having an SA within the range defined by either:

<src-ip-addr/cidr-mask-bits>

or

<src-ip-addr <mask>>

Use this criterion for traffic received from either a subnet or a group of IP addresses. The mask can be in either dotted-decimal format or CIDR format with the number of significant bits.

[udp/tcp-options]

```
destination-port <tcp-data> [ source-port <tcp-data> ]
```

```
source-port <tcp-data> [ destination-port <tcp-data> ]
```

```
destination-port <udp-data> [ source-port <udp-data> ]
```

```
source-port <udp-data> [ destination-port <udp-data> ]
```

tcp-data: *operator tcp-port-#*

udp-data: *operator udp-port-#*

```
operator < eq | gt | lt | neq | range >
```

```
eq <port-nbr-or-name>
```

"Equal To": To have a match with the ACE entry, the TCP or UDP source-port number in a packet must be equal to the specified port number.

```
gt <port-nbr-or-name>
```

"Greater Than": To have a match with the ACE entry, the TCP or UDP source-port number in a packet must be greater than the specified port number.

```
lt <port-nbr-or-name>
```

"Less Than": To have a match with the ACE entry, the TCP or UDP source-port number in a packet must be less than the specified port number.

```
neq <port-nbr-or-name>
```

"Not Equal": To have a match with the ACE entry, the TCP or UDP source-port number in a packet must not be equal to the specified port number.

```
range <start-port-nbr/name> <end-port-nbr/name>
```

To have a match with the ACE entry, the TCP or UDP source-port number in a packet must be in the range `<start-port-nbr/name> <end-port-nbr/name>`.

tcp-data or *udp-data*

TCP or UDP Port Number or (Well-Known) Port Name:

Use the TCP or UDP port number required for the desired match. The switch also accepts certain well-known TCP or UDP port names as alternates to their corresponding port numbers:

```
TCP/UDP-PORT:
```

Specify port by number.
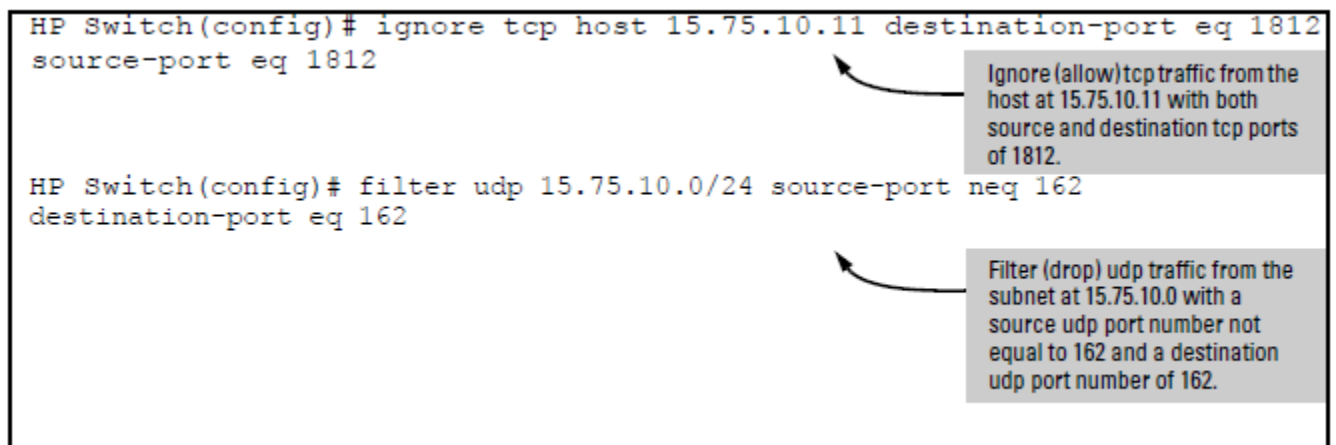
```
bootpc
```

Bootstrap Protocol, client (68)

```
bootps
```

Bootstrap Protocol, server (67)

`dns`

Domain Name Service (53)

`ntp`

Network Time Protocol (123)

`radius`

Remote Authentication Dial-In User Service (1812)

`radius-old`

Remote Authentication Dial-In User Service (1645)

`rip`

Routing Information Protocol (520)

`snmp`

Simple Network Management Protocol (161)

`snmp-trap`

Simple Network Management Protocol (162)

`tftp`

Trivial File Transfer Protocol (69)

Figure 13: *Examples of connection-rate ACEs using UDP/TCP criteria*

```
HP Switch(config)# ignore tcp host 15.75.10.11 destination-port eq 1812
source-port eq 1812
```

> Ignore (allow) tcp traffic from the host at 15.75.10.11 with both source and destination tcp ports of 1812.

```
HP Switch(config)# filter udp 15.75.10.0/24 source-port neq 162
destination-port eq 162
```

> Filter (drop) udp traffic from the subnet at 15.75.10.0 with a source udp port number not equal to 162 and a destination udp port number of 162.

Applying connection-rate ACLs

To apply a connection-rate ACL, use the access group command described below.

NOTE

This command differs from the access group command for non-connection-rate ACLs.

Syntax

```
[no] vlan <vid> ip access-group <crf-list-name> connection-rate-filter
```

This command applies a connection-rate access control list (ACL) to inbound traffic on ports in the specified VLAN that are configured for connection-rate filtering. A connection-rate ACL does not apply to ports in the VLAN that are not configured for connection-rate filtering.

The `[no]` form of the command removes the connection-rate ACL assignment from the VLAN.

> **NOTE**
> - The switch allows only one connection-rate ACL assignment per VLAN. If a connection-rate ACL is already assigned to a VLAN, assigning another to the same VLAN overwrites the first ACL with the second.
> - A connection-rate ACL can be in addition to any standard or extended ACLs already assigned to the VLAN.

Using an ACL in a connection-rate configuration example

This example adds connection-rate ACLs to the example in **Viewing the connection-rate configuration** .

Figure 14: *Sample network*
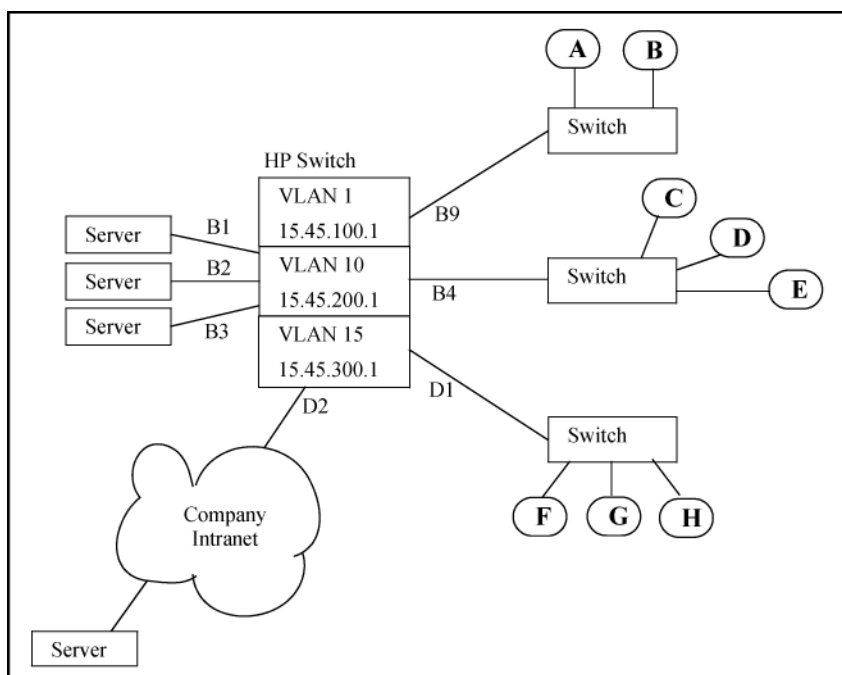
In the basic example, the administrator configured connection-rate blocking on port D2. However:

- The administrator has elevated the connection-rate sensitivity to `high`.

- The server at IP address 15.45.50.17 frequently transmits a relatively high rate of legitimate connection requests, which now triggers connection-rate blocking of the server's IP address on port D2. This causes periodic, unnecessary blocking of access to the server.

The administrator needs to maintain blocking protection from the "Company Intranet" while allowing access to the server at 15.45.50.17. Because the server is carefully maintained as a trusted device, the administrator's solution is to configure a connection-rate ACL that causes the switch to ignore (circumvent) connection-rate filtering for inbound traffic from the server, while maintaining the filtering for all other inbound traffic on port D2.

The configuration steps include:

Procedure

1. Create the connection-rate ACL with a single entry:

   a. Use the IP address of the desired server.

   b. Include a CIDR notation of "32" for the ACL mask. (Which means the mask allows only traffic whose source IP address (SA) exactly matches the specified IP address.)

   c. The ACL automatically includes the implicit
      `filter`

ACE as the last entry, which means that any traffic that is not from the desired server is subject to filtering by the connection-rate policy configured on port D2.

2. Assigning the ACL to the VLAN through which traffic from the server enters the switch.

Figure 15: *Creating and assigning a connection rate ACL*

```
                                              Configures the action to allow traffic from the server at
   Enters the connection-                     15.45.50.17 without filtering through the per-port connection-
   rate ACL context and                       rate policy configured on port D2.
   names the ACL.


HP Switch(config)# ip access-list connection-rate-filter 17-server
HP Switch(config-crf-nacl)# ignore ip host 15.45.50.17
HP Switch(config-crf-nacl)# exit

HP Switch(config)# vlan 15
HP Switch(vlan-15)# ip access-group 17-server connection-rate-filter
HP Switch(vlan-15)# exit                                      Assigns the new
HP Switch(config)# write mem                                  ACL to VLAN 15,
                                                              where traffic
                                                              from the desired
                                                              server enters the
                                                              switch.
```

Figure 16: *Example of switch configuration display with a connection-rate ACL*

```
HP Switch(config)# show config
Startup configuration

; J8697A Configuration Editor; Created on release #K.15.XX
hostname "HP Switch"
connection-rate-filter sensitivity high
ip access-list connection-rate-filter "17-server"
    ignore ip 15.45.50.17  0.0.0.0
    exit                               The new switch configuration
module 2 type J8702A                   includes the ACL configured in
module 4 type J8702A                   figure 3-10.
ip routing
logging 13.28.234.180
snmp-server community "public" Unrestricted
snmp-server host 15.45.200.75 "public"
vlan 1
    name "DEFAULT_VLAN"
    untagged B5-B24
    no ip address
    no untagged B1-B4, D1-D24
    ip proxy-arp
    exit
vlan 10
    name "VLAN10"
    untagged B1-B4
    no ip address
    ip proxy-arp
    exit                               Shows the assignment of the
vlan 15                                above connection-rate ACL to
    name "VLAN_15"                      VLAN 15.
    untagged D1-D24
    no ip address
    ip proxy-arp
    ip access-group "17-server" connection-rate-filter
    exit
filter connection-rate B4 notify-only
filter connection-rate B1-B3 throttle
filter connection-rate B9, D1-D2 block
```

Connection-rate filtering

Features and benefits

Connection-rate filtering is a countermeasure tool you can use in your incident-management program to help detect and manage worm-type IT security threats received in inbound IP traffic. Major benefits of this tool include:

- Behavior-based operation that does not require identifying details unique to the code exhibiting the worm-like operation.

- Handles unknown worms.

- Needs no signature updates.

- Protects network infrastructure by slowing or stopping IP traffic from hosts exhibiting high connection-rate behavior.

- Allows network and individual switches to continue to operate, even when under attack.

- Provides Event Log and SNMP trap warnings when worm-like behavior is detected.

- Gives IT staff more time to react before the threat escalates to a crisis.

> **NOTE**
> When configured on a port, connection-rate filtering is triggered by IPv4 traffic received inbound with a relatively high rate of IP connection attempts.

> **NOTE**
> As stated previously, connection-rate filtering is triggered by inbound IP traffic exhibiting a relatively high-incidence of IP connection attempts from a single source.

Figure 17: *Example of protecting a network from agents using a high IP connection rate to propagate*
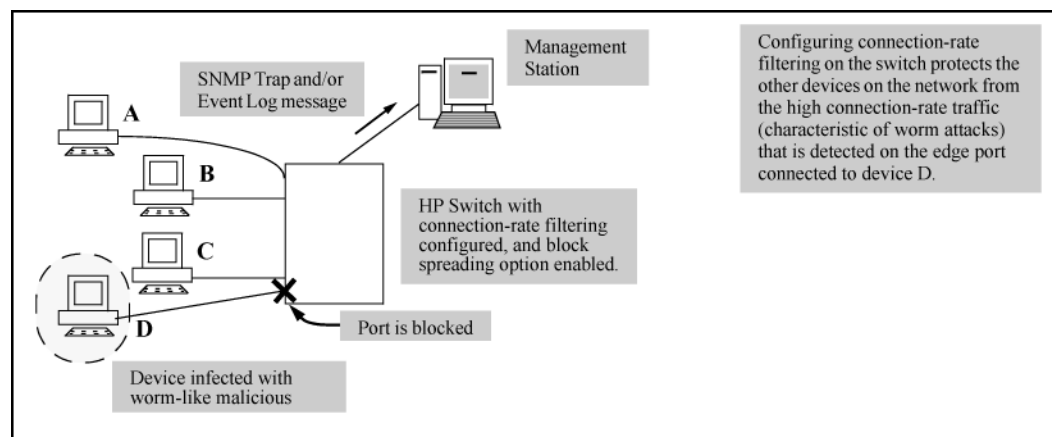
General operation

Connection-rate filtering enables notification of worm-like behavior detected in inbound IP traffic and, depending on how you configure the feature, also throttles or blocks such traffic. This feature also provides a method for allowing legitimate, high connection-rate traffic from a given host while still protecting your network from possibly malicious traffic from other hosts.

Filtering options

In the default configuration, connection-rate filtering is disabled. When enabled on a port, connection-rate filtering monitors inbound IP traffic for a high rate of connection requests from any given host on the port. If a host appears to exhibit the worm-like behavior of attempting to establish a large number of outbound IP connections in a short period of time, the switch responds in one of the following ways, depending on how connection-rate filtering is configured:

- **Notify only** (of potential attack): While the apparent attack continues, the switch generates an Event Log notice identifying the offending host's source IP address and (if a trap receiver is configured on the switch) a similar SNMP trap notice).

- **Throttle**: In this case, the switch temporarily blocks inbound IP traffic from the offending host source IP address for a "penalty" period and generates an Event Log notice of this action and (if a trap receiver is configured on the switch) a similar SNMP trap notice. When the "penalty" period expires the switch re-evaluates the traffic from the host and continues to block this traffic if the apparent attack continues. (During the re-evaluation period, IP traffic from the host is allowed.)

- **Block**: This option blocks all IP traffic from the host. When a block occurs, the switch generates an Event Log notice and (if a trap receiver is configured on the switch) a similar SNMP trap notice. Note that a network administrator must explicitly re-enable a host that has been previously blocked.

Sensitivity to connection rate detection

The switch includes a global sensitivity setting that enables adjusting the ability of connection-rate filtering to detect relatively high instances of connection-rate attempts from a given source.

Application options

For the most part, normal network traffic is distinct from the traffic exhibited by malicious agents. However, when a legitimate network host generates multiple connections in a short period of time, connection-rate filtering can generate a "false positive" and treat the host as an infected client. Lowering the sensitivity or changing the filter mode can reduce the number of false positives. Conversely, relaxing filtering and sensitivity provisions lowers the switch ability to detect worm-generated traffic in the early stages of an attack, and should be carefully investigated and planned to ensure that a risky vulnerability is not created. As an alternative, you can use connection-rate ACLs (access control lists) or selective enabling to allow legitimate traffic.

Selective enable

This option involves applying connection-rate filtering only to ports posing a significant risk of attack. For ports that are reasonably secure from attack, then there can be little benefit in configuring them with connection-rate filtering.

Connection-rate Access Control Lists (ACLs)

The basic connection-rate filtering policy is configured per-port as `notify-only`, `throttle`, and `block`. A connection-rate ACL creates exceptions to these per-port policies by creating special rules for individual hosts, groups of hosts, or entire subnets. Thus, you can adjust a connection-rate filtering policy to create and apply an exception to configured filters on the ports in a VLAN. Note that connection-rate ACLs are useful only if you need to exclude inbound traffic from your connection-rate filtering policy. For example, a server responding to network demand can send a relatively high number of legitimate connection requests. This can generate a false positive by exhibiting the same elevated connection-rate behavior as a worm. Using a connection-rate ACL to apply an exception for this server allows you to exclude the trusted server from connection-rate filtering and thereby keep the server running without interruption.

> **NOTE**
> Use connection-rate ACLs only when you need to exclude an IP traffic source (including traffic with specific UDP or TCP criteria) from a connection-rate filtering policy. Otherwise, the ACL is not necessary.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Operating rules

- Connection-rate filtering does not operate on IPv6 traffic.

- Connection-rate filtering is triggered by inbound IP traffic exhibiting high rates of IP connections to new hosts. After connection-rate filtering has been triggered on a port, all traffic from the suspect host is subject to the configured connection-rate policy (`notify-only`, `throttle`, or `block`).

- When connection-rate filtering is configured on a port, the port cannot be added to, or removed from, a port trunk group. Before this can be done, connection-rate filtering must be disabled on the port.

- Where the switch is throttling or blocking inbound IP traffic from a host, any outbound traffic destined for that host is still permitted.

- Once a throttle has been triggered on a port—temporarily blocking inbound IP traffic—it cannot be undone during operation: the penalty period must expire before traffic is allowed from the host.

Unblocking a currently blocked host

A hostblocked by connection-rate filtering remains blocked until explicitly unblocked by one of the following methods:

- Using the `connection-rate-filter unblock` command, see **Listing currently-blocked hosts** on page 82.

- Rebooting the switch.

- Disabling connection-rate filtering using the `no connection-rate-filter` command.

- Deleting a VLAN removes blocks on any hosts on that VLAN.

> **NOTE:** Changing a port setting from `block` to `throttle`, `notify-only`, or to `no filter connection-rate`, does not unblock a currently blocked host. Similarly, applying a connection-rate ACL does not unblock a currently blocked host. See the above list for the correct methods to use to unblock a host.

Applying connection-rate ACLs

A host sending legitimate traffic can trigger connection-rate filtering in some circumstances. If you can verify that such a host is indeed sending valid traffic and is not a threat to your network, you can want to configure a connection-rate ACL (access control list) that allows this traffic to bypass the configured connection-rate filtering.

A connection-rate ACL is an optional tool that consists of one or more explicitly configured Access Control Entries (ACEs) used to specify whether to enforce the configured connection-rate policy on traffic from a particular source.

Use of connection-rate ACLs provides the option to apply exceptions to the configured connection-rate filtering policy. This enables you to allow legitimate traffic from a trusted source, and apply connection-rate filtering only to inbound traffic from untrusted sources. For example, where a connection-rate policy has been configured, you can apply a connection-rate ACL that causes the switch bypass connection-rate policy filtering on traffic from:

- A trusted server exhibiting a relatively high IP connection rate due to heavy demand

- A trusted traffic source on the same port as other, untrusted traffic sources.

The criteria for an exception can include the source IP address of traffic from a specific host, group of hosts, or a subnet, and can also include source and destination TCP/UDP criteria. This allows you to apply a notify-only, throttling, or blocking policy while allowing exceptions for legitimate traffic from specific sources. You can also allow exceptions for traffic with specific TCP or UDP criteria.

For more information on when to apply connection-rate ACLs, see **Application options** on page 92.

> **NOTE**
>
> Connection-rate ACLs are a special case of the switch ACL feature. If you need information on other applications of ACLs or more detailed information on how ACLs operate, see **IPv4 Access Control Lists (ACLs)** on page 398.

Connection-rate ACL operation

A connection-rate ACL applies to inbound traffic on all ports configured for connection-rate filtering in the assigned VLAN, and creates an exception to the connection-rate filter policy configured on each port. A connection-rate ACL has no effect on ports in the VLAN that are not configured for connection-rate filtering.

A connection-rate ACL accepts inbound, legitimate traffic from trusted sources without filtering the traffic for the configured connection-rate policy. You can configure an ACL to assign policy filtering (`filter`) for traffic from some sources and no policy filtering (`ignore`) for traffic from other sources. However, the implicit `filter` invoked as the last entry in any connection-rate ACL ensures that any traffic not specifically excluded from policy filtering (by the `ignore` command) is filtered by the configured policy for the port on which that traffic entered the switch.

Figure 18: *Connection-rate ACL applied to traffic received through a given port*
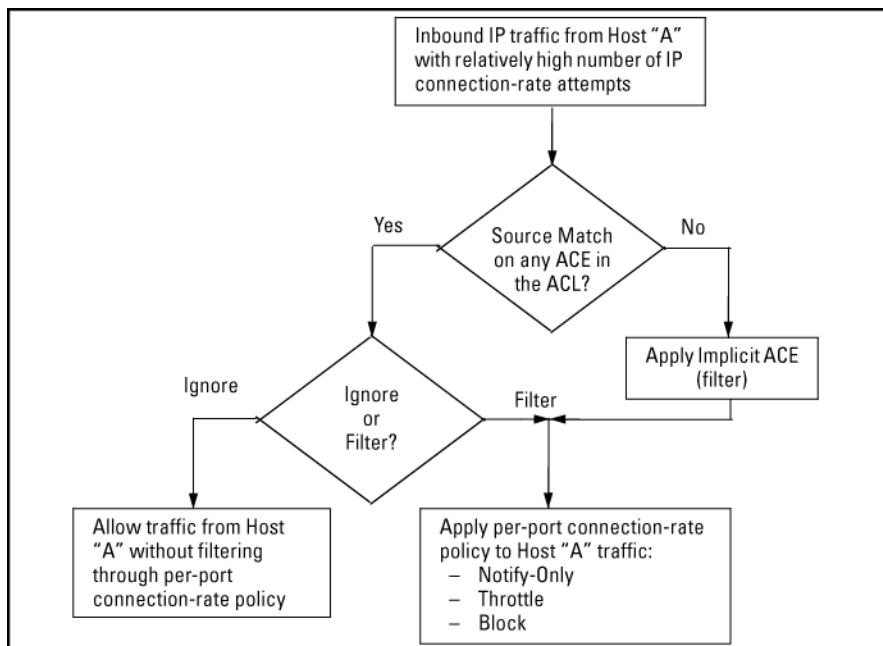
Connection-Rate ACL operating notes

- ACE Types:

   A connection-rate ACL allows you to configure two types of ACEs (Access Control Entries):

   ◦ `ignore <source-criteria>`

      This ACE type directs the switch to permit all inbound traffic meeting the configured `<source-criteria>` without filtering the traffic through the connection-rate policy configured on the port through which the traffic entered the switch. For example, `ignore host 15.45.120.70` tells the switch to permit traffic from the host at 15.45.120.70 without filtering this host's traffic through the connection-rate policy configured for the port on which the traffic entered the switch.

   ◦ `filter <source-criteria>`

This ACE type does the opposite of an `ignore` entry. That is, all inbound traffic meeting the configured *source-criteria* must be filtered through the connection-rate policy configured for the port on which the traffic entered the switch. This option is most useful in applications where it is easier to use `filter` to specify suspicious traffic sources for screening than to use `ignore` to specify exceptions for trusted traffic sources that don't need screening. For example, if the host at 15.45.127.43 requires connection-rate screening, but all other hosts in the VLAN do not, you would configure and apply a connection-rate ACL with `filter ip host 15.45.127.43` as the first ACE and `ignore ip any` as the second ACE. In this case, the traffic from host 15.45.127.43 would be screened, but traffic from all other hosts on the VLAN would be permitted without connection-rate screening.

- Implicit ACE

  A connection-rate ACL includes a third, implicit `filter ip any` ACE which is automatically the last ACE in the ACL. This implicit ACE does not appear in displays of the ACL configuration, but is always present in any connection-rate ACL you configure. For example, assume that a port is configured with a connection-rate policy and is in a VLAN configured with a connection-rate ACL. If there is no match between an incoming packet and the ACE criteria in the ACL, then the implicit `filter ip any` sends the packet for screening by the connection-rate policy configured on that port. To preempt the implicit `filter ip any` in a given connection-rate ACL, you can configure `ignore IP any` as the last explicit ACE in the connection-rate ACL. The switch then ignores (permit) traffic that is not explicitly addressed by other ACEs configured sequentially earlier in the ACL without filtering the traffic through the existing connection-rate policy.

- Monitoring Shared Resources

  Active instances of throttling or blocking a client that is generating a high rate of connection requests uses internal routing switch resources that are shared with several other features. The routing switch provides ample resources for all features. However, if the internal resources become fully subscribed, new instances of throttling or blocking cannot be initiated until the necessary resources are released from other uses. (Event Log messages and SNMP traps are not affected.) For information on determining current resource availability and usage, see the appendix titled "Monitoring Resources" in the management and configuration guide for your switch.

Using CIDR notation to enter the ACE mask

You can use Classless Inter-Domain Routing (CIDR) notation to enter ACE masks. The switch interprets the bits specified with CIDR notation as the IP address bits in an ACE and the corresponding IP address bits in a packet. The switch then converts the mask to inverse notation for ACE use.

Table 10: *CIDR notation for masks*

IP address used in an ACL with CIDR notation	Resulting ACL mask	Meaning
10.38.240.125/15	0.1.255.255	The leftmost 15 bits must match; the remaining bits are wildcards.
10.38.240.125/20	0.0.15.255	The leftmost 20 bits must match; the remaining bits are wildcards.
10.38.240.125/21	0.0.7.255	The leftmost 21 bits must match; the remaining bits are wildcards.

Table Continued

IP address used in an ACL with CIDR notation	Resulting ACL mask	Meaning
10.38.240.125/24	0.0.0.255	The leftmost 24 bits must match; the remaining bits are wildcards.
10.38.240.125/32	0.0.0.0	All bits must match.

Connection-rate log and trap messages

See the Event Log message reference guide for information about Event Log messages.

Overview

The spread of malicious agents in the form of worms has severe implications for network performance. Damage can be as minimal as slowing down a network with excessive, unwanted traffic, or as serious as putting attacker-defined code on a system to cause any type of malicious damage.

Current methods to stop the propagation of malicious agents rely on signature recognition to prevent hosts from being infected. However, the latency between the introduction of a new virus or worm into a network, and the implementation and distribution of a signature-based patch can be significant. Within this period, a network can be crippled by the abnormally high rate of traffic generated by infected hosts.

Connection-rate filtering based on virus throttling technology is recommended for use on the edge of a network. It is primarily concerned with the class of worm-like malicious code that tries to replicate itself by using vulnerabilities on other hosts (weaknesses in network applications behind unsecured ports). Agents of this variety operate by choosing a set of hosts to attack based on an address range (sequential or random) that is exhaustively searched, either by blindly attempting to make connections by rapidly sending datagrams to the address range, or by sending individual ICMP ping messages to the address range and listening for replies.

Connection-rate filtering detects the network behavior of malicious code that tries to create a large number of outbound IP connections on an interface in a short time. When a host exhibits this behavior, warnings can be sent, and connection requests can be either throttled or dropped to minimize the barrage of subsequent traffic from the host. When enabled on the switch, connection-rate filtering can help reduce the impact of worm-like malicious code and give system administrators more time to isolate and eradicate the threat. Thus, while traditional worm and virus-signature updates still need to be deployed to hosts, the network remains functional and the overall distribution of the malicious code is limited.

Configuring connection-rate filtering for low risk networks

As stated earlier, connection-rate filtering is triggered only by inbound IP traffic generating a relatively high number of new IP connection requests from the same host.

Procedure

1. Enable `notify-only` mode on the ports you want to monitor.

2. Set global sensitivity to `low`.

3. If SNMP trap receivers are available in your network, use the `snmp-server` command to configure the switch to send SNMP traps.

4. Monitor the Event Log or (if configured) the available SNMP trap receivers to identify hosts exhibiting high connection rates.

5. Check any hosts that exhibit relatively high connection rate behavior to determine whether malicious code or legitimate use is the cause of the behavior.

6. Hosts demonstrating high, but legitimate connection rates, such as heavily used servers, can trigger a connection-rate filter. Configure connection rate ACLs to create policy exceptions for trusted hosts. (Exceptions can be configured for these criteria:

   a. A single source host or group of source hosts

   b. A source subnet

   c. Either of the above with TCP or UDP criteria

      For more on connection rate ACLs, see **Application options** on page 92.

7. Increase the sensitivity to `Medium` and repeat steps 5 and 6.

   > **NOTE** On networks that are relatively infection-free, sensitivity levels above
   >
   > `Medium`
   >
   > are not recommended.

8. (Optional.) Enable `throttle` or `block` mode on the monitored ports.

   > **NOTE** On a given VLAN, to unblock the hosts that have been blocked by the connection-rate feature, use the
   >
   > `vlan <vid> connection-rate filter unblock`
   >
   > command.

9. Maintain a practice of carefully monitoring the Event Log or configured trap receivers for any sign of high connectivity-rate activity that could indicate an attack by malicious code.

Configuring connection-rate filtering for high risk networks

This procedure is similar to the general steps required for a relatively attack free network, except for policies suggested for managing hosts exhibiting high connection rates. This allows better network performance for unaffected hosts and helps to identify hosts that can require updates or patches to eliminate malicious code.

Procedure

1. Configure connection-rate filtering to `throttle` on all ports.

2. Set global sensitivity to `medium`.

3. If SNMP trap receivers are available in your network, use the `snmp-server` command to configure the switch to send SNMP traps.

4. Monitor the Event Log or the available SNMP trap receivers (if configured on the switch) to identify hosts exhibiting high connection rates.

5. Check any hosts that exhibit relatively high connection rate behavior to determine whether malicious code or legitimate use is the cause of the behavior.

6. On hosts you identify as needing attention to remove malicious behavior:

a. To immediately halt an attack from a specific host, group of hosts, or a subnet, use the per-port block mode on the appropriate ports.

b. After gaining control of the situation, you can use connection-rate ACLs to more selectively manage traffic to allow receipt of normal traffic from reliable hosts.

Overview

Web and MAC authentication are designed for employment on the "edge" of a network to provide port-based security measures for protecting private networks and a switch from unauthorized access. Because neither method requires clients to run special supplicant software (unlike 802.1X authentication), both web and MAC authentication are suitable for legacy systems and temporary access situations where introducing supplicant software is not an attractive option. Only a web browser (for web-based authentication) or a MAC address (for MAC authentication) is required.

Both web and MAC authentication methods rely on a RADIUS server to authenticate network access. This simplifies access security management by allowing the control of access from a master database in a single server. Up to three RADIUS servers can be used for backup in case access to the primary server fails. It also means the same credentials can be used for authentication, regardless of which switch or switch port is the current access point into the LAN.

On a port configured for web-based or MAC authentication, the switch operates as a port-access authenticator using a RADIUS server and the CHAP protocol. Inbound traffic is processed by the switch alone, until authentication occurs. Some traffic from the switch to an unauthorized client is supported (for example, broadcast or unknown destination packets) before authentication occurs.

Web-based authentication

The web-based authentication method uses a web page login to authenticate users for access to the network. When a client connects to the switch and opens a web browser, the switch automatically presents a login page.

> **NOTE:** A proxy server is not supported for use by a browser on a client device that accesses the network through a port configured for web-based authentication.

In the login page, a client enters a username and password, which the switch forwards to a RADIUS server for authentication. After authenticating a client, the switch grants access to the secured network. Besides a web browser, the client needs no special supplicant software.

MAC authentication

The MAC authentication method grants access to a secure network by authenticating devices for access to the network. When a device connects to the switch, either by direct link or through the network, the switch forwards the device's MAC address to the RADIUS server for authentication. The RADIUS server uses the device MAC address as the username and password, and grants or denies network access in the same way that it does for clients capable of interactive logons. The process does not use either a client device configuration or a logon session. MAC authentication is well-suited for clients not capable of providing interactive logons, such as telephones, printers, and wireless access points. Also, because most RADIUS servers allow for authentication to depend on the source switch and port through which the client connects to the network, you can use MAC authentication to "lock" a particular device to a specific switch and port.

NOTE

802.1X port-access, web-based authentication, and MAC authentication can be configured at the same time on the same port. The client limit is 256 clients per port for MAC-auth and Web-auth; the client limit for 802.1X is 32 clients per port. The MAC-auth and Web-auth limit of 256 clients only applies when there are fewer than 16,384 authentication clients on the entire switch. After the limit of 16,384 clients is reached, no additional authentication clients are allowed on any port for any method. The default is one client.

Web-based and/or MAC authentication and MAC lockdown, MAC lockout, and port-security are mutually exclusive on a given port. If you configure any of these authentication methods on a port, you must disable LACP on the port.

Concurrent web-based and MAC authentication

Web-based authentication and MAC authentication can be configured at the same time on a port. It is assumed that MAC authentication will use an existing MAC address. The following conditions apply for concurrent authentication:

- A specific MAC address cannot be authenticated by both web and MAC authentication at the same time.

- Each new web-based/MAC authentication client always initiates a MAC authentication attempt. This same client can also initiate web-based authentication at any time before the MAC authentication succeeds. If either authentication succeeds then the other authentication (if in progress) is ended. No further web-based/MAC authentication attempts are allowed until the client is de-authenticated.

- Web-based and MAC authentications are not allowed on the same port if an unauthenticated (guest) VLAN is enabled for MAC authentication. An unauthenticated VLAN cannot be enabled for MAC authentication if web-based and MAC authentication are both enabled on the port.

- Hitless reauthentication must be of the same type (MAC) that was used for the initial authentication. Non-hitless reauthentication can be of any type.

The remaining web-based/MAC functionality, including interactions with 802.1X, remains the same. web and MAC authentication can be used for different clients on the same port.

Normally, MAC authentication finishes much sooner than web authentication. However, if web authentication completes first, MAC authentication ceases, even though MAC authentication could succeed. There is no guarantee that MAC authentication ends before web-based authentication begins for the client.

Concurrent web-based and MAC authentication is backward compatible with all existing user configurations.

Authorized and unauthorized client VLANs

Web-based and MAC Authentication provides a port-based solution in which a port belongs to one untagged VLAN at a time. The switch supports up to 32 simultaneous client sessions per port.

All authenticated client sessions operate in the same untagged VLAN. To simultaneously support multiple client sessions in different VLANs for a network application, design the system so clients request network access on different switch ports.

In the default configuration, the switch blocks access to all clients that the RADIUS server does not authenticate. However, you can configure an individual port to provide limited network services and access to unauthorized clients by using an "unauthorized" VLAN for each session. The unauthorized VLAN ID assignment can be the same for all ports, or different, depending on the services and access you plan to allow for unauthenticated clients.

You configure access to an optional, unauthorized VLAN when you configure web-based and MAC authentication on a port.

RADIUS-based authentication

In web-based and MAC authentication, you use a RADIUS server to temporarily assign a port to a static VLAN to support an authenticated client. When a RADIUS server authenticates a client, the switch-port membership during the client's connection is determined according to the following hierarchy:

Procedure

1. A RADIUS-assigned VLAN.

2. An authorized VLAN specified in the web-based or MAC authentication configuration for the subject port.

3. A static, port-based, untagged VLAN to which the port is configured. A RADIUS-assigned VLAN has priority over switch-port membership in any VLAN.

Wireless clients

You can allow wireless clients to move between switch ports under web-based/MAC authentication control. Clients can move from one web-authorized port to another or from one MAC-authorized port to another. This capability allows wireless clients to move from one access point to another without having to reauthenticate.

How web-based and MAC authentication operate
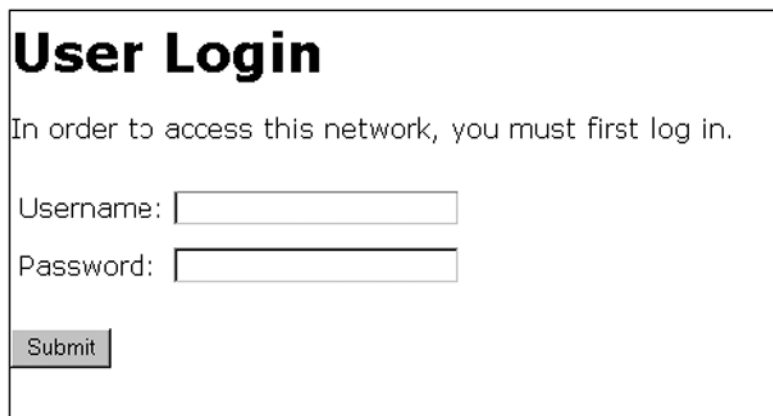
Before gaining access to the network, a client first presents authentication credentials to the switch. The switch then verifies the credentials with a RADIUS authentication server. Successfully authenticated clients receive access to the network, as defined by the System Administrator. Clients who fail to authenticate successfully receive no network access or limited network access as defined by the System Administrator.

Web-based authentication

When a client connects to a web-based authentication enabled port, communication is redirected to the switch. The switch assigns a temporary IP address and a login screen is presented for the client to enter their username and password.

The default User Login screen is shown in **Figure 19: Default user login screen** on page 101.

Figure 19: *Default user login screen*

When a client connects to the switch, it sends a DHCP request to receive an IP address to connect to the network. To avoid address conflicts in a secure network, you can specify a temporary IP address pool to be used by DHCP by configuring the `dhcp-addr` and `dhcp-lease` options when you enable web-based authentication with the `aaa port-access web-based` command.

The Secure Sockets Layer (SSLv3/TLSv1) feature provides remote web-based access to the network through authenticated transactions and encrypted paths between the switch and management station clients capable of SSL/TLS. If you have enabled SSL on the switch, you can specify the `ssl-login` option when you configure web-based authentication so that clients who log in to specified ports are redirected to a secure login page (https://...) to enter their credentials.

The switch passes the supplied username and password to the RADIUS server for authentication and displays the following progress message:

Figure 20: *Progress message during authentication*
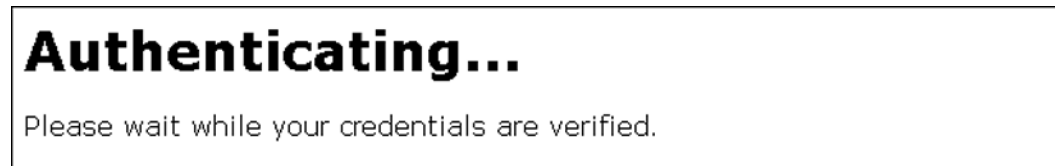
If the client is authenticated and the maximum number of clients allowed on the port (`client-limit`) has not been reached, the port is assigned to a static, untagged VLAN for network access. After a successful login, a client can be redirected to a URL if you specify a URL value (`redirect-url`) when you configure web-based authentication.

Figure 21: *Authentication completed*
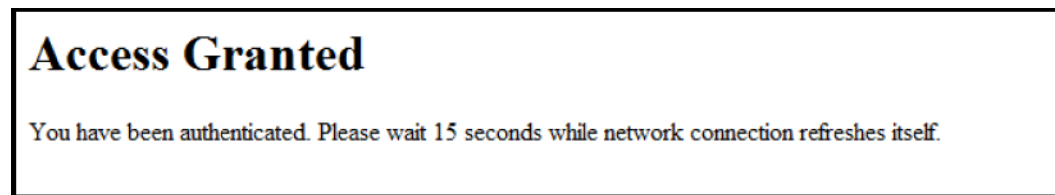
Order of priority for assigning VLANs

The assigned VLAN is determined, in order of priority, as follows:

Procedure

1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.

2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the authorized VLAN (`auth-vid` if configured) and temporarily drops all other VLAN memberships.

3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.

4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.

The assigned port VLAN remains in place until the session ends. Clients can be forced to reauthenticate after a fixed period of time (`reauth-period`) or at any time during a session (`reauthenticate`). An implicit logoff period can be set if there is no activity from the client after a given amount of time (`logoff-period`). In addition, a session ends if the link on the port is lost, requiring reauthentication of all clients. Also, if a client moves from one port to another and client moves have not been enabled (`client-moves`) on the ports, the session ends and the client must reauthenticate for network access. At the end of the session the port returns to its pre-authentication state. Any changes to the port's VLAN memberships made while it is an authorized port take affect at the end of the session.

A client can not be authenticated due to invalid credentials or a RADIUS server timeout. The `max-retries` parameter specifies how many times a client can enter their credentials before authentication fails. The `server-timeout` parameter sets how long the switch waits to receive a response from the RADIUS server before timing out. The `max-requests` parameter specifies how many authentication attempts can result in a RADIUS server timeout before authentication fails. The switch waits a specified amount of time (`quiet-period`) before processing any new authentication requests from the client.

Network administrators can assign unauthenticated clients to a specific static, untagged VLAN (`unauth-vid`), to provide access to specific (guest) network resources. If no VLAN is assigned to unauthenticated clients the port is blocked and no network access is available. Should another client successfully authenticate through that port any unauthenticated clients on the `unauth-vid` are dropped from the port.

MAC-based authentication

When a client connects to a MAC authentication enabled port traffic is blocked. The switch immediately submits the client's MAC address (in the format specified by the `addr-format`) as its certification credentials to the RADIUS server for authentication.

If the client is authenticated and the maximum number of MAC addresses allowed on the port (`addr-limit`) has not been reached, the port is assigned to a static, untagged VLAN for network access.

The assigned VLAN is determined, in order of priority, as follows:

Procedure

1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.

2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the authorized VLAN (if configured) and temporarily drops all other VLAN memberships.

3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.

4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.

The assigned port VLAN remains in place until the session ends. Clients may be forced to reauthenticate after a fixed period of time (`reauth-period`) or at any time during a session (`reauthenticate`). An implicit logoff period can be set if there is no activity from the client after a given amount of time (`logoff-period`). In addition, a session ends if the link on the port is lost, requiring reauthentication of all clients. Also, if a client moves from one port to another and client moves have not been enabled (`addr-moves`) on the ports, the session ends and the client must reauthenticate for network access. At the end of the session the port returns to its pre-authentication state. Any changes to the port's VLAN memberships made while it is an authenticated port take affect at the end of the session.

A client may not be authenticated due to invalid credentials or a RADIUS server timeout. The `server-timeout` parameter sets how long the switch waits to receive a response from the RADIUS server before timing out. The `maxrequests` parameter specifies how many authentication attempts may result in a RADIUS server timeout before authentication fails. The switch waits a specified amount of time (`quiet-period`) before processing any new authentication requests from the client.

Network administrators may assign unauthenticated clients to a specific static, untagged VLAN (`unauth-vid`), to provide access to specific (guest) network resources. If no VLAN is assigned to unauthenticated clients the port remains in its original VLAN configuration. Should another client successfully authenticate through that port any unauthenticated clients are dropped from the port.

Operating rules and notes

- The switch supports concurrent 802.1X, web andMAC authentication operation on a port (with up to 32 clients allowed). However, concurrent operation of web and MAC authentication with other types of authentication on the same port is not supported. That is, the following authentication types are mutually exclusive on a given port:

  ◦ Web-based and/or MAC authentication (with or without 802.1X)

  ◦ MAC lockdown

  ◦ MAC lockout

  ◦ Port-Security

- Order of Precedence for Port Access Management (highest to lowest):

  1. MAC lockout

  2. MAC lockdown or Port Security

  3. Port-based Access Control (802.1X) or web-based authentication or MAC authentication

---

NOTE

When configuring a port for web-based or MAC authentication, be sure that a higher precedent port access management feature is not enabled on the port. For example, be sure that Port Security is disabled on a port before configuring the port for web-based or MAC authentication. If Port Security is enabled on the port this misconfiguration does not allow web-based or MAC authentication to occur.

- VLANs: If your LAN does not use multiple VLANs, then you do not need to configure VLAN assignments in your RADIUS server or consider using either authorized or unauthorized VLANs. If your LAN does use multiple VLANs, then some of the following factors can apply to your use of web-based authentication and MAC-authentication:

  ◦ Web-based authentication and MAC-authentication operate only with port-based VLANs. Operation with protocol VLANs is not supported, and clients do not have access to protocol VLANs during web-based authentication and MAC authentication sessions.

  ◦ A port can belong to one, untagged VLAN during any client session. Where multiple authenticated clients can simultaneously use the same port, they must all be capable of operating on the same VLAN.

  ◦ During an authenticated client session, the following hierarchy determines a port's VLAN membership:

    1. If there is a RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to this VLAN and temporarily drops all other VLAN memberships.

    2. If there is no RADIUS-assigned VLAN, then, for the duration of the client session, the port belongs to the authorized VLAN (if configured) and temporarily drops all other VLAN memberships.

    3. If neither 1 or 2, above, apply, but the port is an untagged member of a statically configured, port-based VLAN, then the port remains in this VLAN.

    4. If neither 1, 2, or 3, above, apply, then the client session does not have access to any statically configured, untagged VLANs and client access is blocked.

◦ After an authorized client session begins on a given port, the port's VLAN membership does not change. If other clients on the same port become authenticated with a different VLAN assignment than the first client, the port blocks access to these other clients until the first client session ends.

◦ The optional "authorized" VLAN (`auth-vid`) and "unauthorized" VLAN (`unauth-vid`) you can configure for web-based or MAC authentication must be statically configured VLANs on the switch. Also, if you configure one or both of these options, any services you want clients in either category to access must be available on those VLANs.

• Where a given port's configuration includes an unauthorized client VLAN assignment, the port will allow an unauthenticated client session only while there are no requests for an authenticated client session on that port. In this case, if there is a successful request for authentication from an authorized client, the switch terminates the unauthorized-client session and begins the authorized-client session.

• When a port on the switch is configured for web-based or MAC authentication and is supporting a current session with another device, rebooting the switch invokes a re-authentication of the connection.

• When a port on the switch is configured as a web-based or MAC authenticator, it blocks access to a client that does not provide the proper authentication credentials. If the port configuration includes an optional, unauthorized VLAN (`unauth-vid`), the port is temporarily placed in the unauthorized VLAN if there are no other authorized clients currently using the port with a different VLAN assignment. If an authorized client is using the port with a different VLAN or if there is no unauthorized VLAN configured, the unauthorized client does not receive access to the network.

• Web-based or MAC authentication and LACP cannot both be enabled on the same port.

• Web-based/MAC authentication and LACP are not supported at the same time on a port. The switch automatically disables LACP on ports configured for web or MAC authentication.

• Use the `show port-access web-based` command to display session status, port-access configuration settings, and statistics for web-based authentication sessions.

• When spanning tree is enabled on a switch that uses 802.1X, web-based authentication, or MAC authentication, loops can go undetected. For example, spanning tree packets that are looped back to an edge port will not be processed because they have a different broadcast/multicast MAC address from the client-authenticated MAC address. To ensure that client-authenticated edge ports get blocked when loops occur, you should enable loop protection on those ports. See "Multiple instance spanning-tree operation" in the advanced traffic management guide for your switch.

Setup procedure for web-based/MAC authentication

Before you configure web-based/MAC authentication, follow these guidelines:

Procedure

1. Configure a local username and password on the switch for both the operator (login) and manager (enable) access levels. (While this is not required for a Web- or MAC-based configuration, Hewlett Packard Enterprise recommends that you use a local user name and password pair, at least until your other security measures are in place, to protect the switch configuration from unauthorized access.)

2. Determine the switch ports that you want to configure as authenticators. Note that before you configure web-based or MAC authentication on a port operating in an LACP trunk, you must remove the port from the trunk.

3. To display the current configuration of 802.1X, web-based, and MAC authentication on all switch ports, enter the `show port-access config` command.
   Output for the show port-access config command

```
HP Switch (config)# show port-access config

Port-access authenticator activated [No] : No
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

     Supplicant Authenticator Web-Auth Mac-Auth LMA-Auth Ctrl  Mixed Speed
Port Enabled    Enabled      Enabled  Enabled  Enabled  Dir   Mode  VSA  MBV
---- --------- ------------ -------- -------- -------- ----- ---- ---- ---
C1   No        Yes          No       No       No       In    No   Yes  Yes
C2   No        Yes          No       No       No       Both  Yes  Yes  Yes
C3   No        Yes          No       No       No       Both  No   No   Yes
C4   No        Yes          No       No       Yes      Both  No   Yes  Yes
...
```

4. Determine whether any VLAN assignments are needed for authenticated clients.

   a. If you configure the RADIUS server to assign a VLAN for an authenticated client, this assignment overrides any VLAN assignments configured on the switch while the authenticated client session remains active. The VLAN must be statically configured on the switch.

   b. If there is no RADIUS-assigned VLAN, the port can join an "Authorized VLAN" for the duration of the client session. This must be a port-based, statically configured VLAN on the switch.

   c. If there is neither a RADIUS-assigned VLAN or an "authorized VLAN" for an authenticated client session on a port, the port's VLAN membership remains unchanged during authenticated client sessions. Configure the port for the VLAN in which you want it to operate during client sessions.

   > **NOTE**
   >
   > When configuring a RADIUS server to assign a VLAN, you can use either the VLAN's name or VID. For example, if a VLAN configured in the switch has a VID of 100 and is named **vlan100**, you could configure the RADIUS server to use either "100" or "vlan100" to specify the VLAN.

5. For clients that the RADIUS server does not authenticate, determine whether to use the optional "unauthorized VLAN" mode. This VLAN must be statically configured on the switch. If you do not configure an "unauthorized VLAN", the switch simply blocks access to unauthenticated clients trying to use the port.

6. Determine the authentication policy you want on the RADIUS server and configure the server. Based on your switches RADIUS application information, include the following in the policy for each client or client device:

   • The CHAP-RADIUS authentication method

   • An encryption key

   • One of the following:

     ◦ Include the user name and password for each authorized client if you are configuring web-based authentication.

     ◦ Enter the device MAC address in both the username and password fields of the RADIUS policy configuration for that device if you are configuring MAC authentication. To allow a particular device to receive authentication only through a designated port and switch, include this in your policy.

7. Determine the IP address of the RADIUS server(s) you choose to support web-based or MAC authentication.

Configuring the RADIUS server to support MAC authentication

On the RADIUS server, configure the client device authentication in the same way that you would any other client, except:

- Configure the client device's (hexadecimal) MAC address as both username and password. Be careful to configure the switch to use the same format that the RADIUS server uses. Otherwise, the server will deny access. The switch provides four format options:

  ◦ aabbccddeeff (the default format)

  ◦ aabbcc-ddeeff

  ◦ aa-bb-cc-dd-ee-ff

  ◦ aa:bb:cc:dd:ee:ff

  ◦ AABBCCDDEEFF

  ◦ AABBCC-DDEEFF

  ◦ AA-BB-CC-DD-EE-FF

  ◦ AA:BB:CC:DD:EE:FF

- If the device is a switch or other VLAN capable device, use the base MAC address assigned to the device, and not the MAC address assigned to the VLAN through which the device communicates with the authenticator switch. The switch applies a single MAC address to all VLANs configured in the switch. Thus, for a given switch, the MAC address is the same for all VLANs configured on the switch. (See "Static Virtual LANs (VLANs)" in the advanced traffic management guide for your switch.)

Configuring the switch to access a RADIUS server

Configuring a RADIUS server to support web-based authentication and MAC Authentication require the following minimal commands:

(See **RADIUS Authentication, Authorization, and Accounting** on page 194 for information on other RADIUS command options.)

Syntax:

```
[no] radius-server
```

```
[host <ip-addresss>]
```

Adds a server to the RADIUS configuration or, when [no] is used, deletes a server from the configuration. You can configure up to three RADIUS server addresses. The switch uses the first server it successfully accesses. (See **RADIUS Authentication, Authorization, and Accounting** on page 194).

```
[key <global-key-string>]
```

Specifies the global encryption key the switch uses with servers for which the switch does not have a server specific key assignment (below). This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key. The tilde (~) character is allowed in the string, for example, radius server key hp~switch. It is not backward compatible; the "~" character is lost if you use a software version that does not support the "~" character.

(Default: Null.)

Syntax:

```
radius-server host <ip-address> key <server-specific key-string>
```

```
[no] radius-server host <ip-address> key
```

Optional.

Specifies an encryption key for use during authentication (or accounting) sessions with the specified server. This key must match the encryption key used on the RADIUS server. Use this command only if the specified server requires a different encryption key than configured for the global encryption key, above. The tilde (~) character is allowed in the string. It is not backward compatible; the "~" character is lost if you use a software version that does not support the "~" character.

The `[no]` form of the command removes the key configured for a specific server.

For example, to configure the switch to access a RADIUS server at IP address 192.168.32.11 using a server specific shared secret key of '1A7rd':

Configure the switch to access a RADIUS server

```
HP Switch(config)# radius-server host 192.168.32.11
HP Switch(config)# radius-server host 192.168.32.11 key 1A7rd

HP Switch(config)# show radius

 Status and Counters - General RADIUS Information

  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key :
  Dynamic Authorization UDP Port :


                Auth Acct DM/ Time
  Server IP Addr  Port Port CoA Window Encryption Key
  --------------- ---- ---- --- ------ ---------------------
  192.168.32.11   1812 1813          1A7rd
```

RADIUS service tracking

RADIUS service tracking determines the availability of RADIUS servers configured on the switch. RADIUS service tracking minimizes the wait time for clients when authentication has failed and they are placed in the unauth-vid (Guest VLAN). The wait time for both unauthenticated clients and previously authenticated clients is reduced when RADIUS service tracking is enabled.

This feature is disabled by default. For more information, see the *HPE ArubaOS-Switch Access Security Guide* for your switch.

radius-server tracking

Syntax

radius-server tracking [enable | disable] [interval *<SECONDS>*]

Description

Enables or disables the RADIUS server tracking timer in minutes.

Command context

config

Parameters

interval *<SECONDS>*

Specifies the number of seconds assigned as a RADIUS tracker. Range for this interval is 60-86400 seconds.

show radius

Use the `show radius` command to display RADIUS tracking information.

```
switch# show radius
Status and Counters - General RADIUS Information
*Server once reachable will be not be considered alive
Deadtime (minutes) : Infinite
Timeout (seconds) : 5
Retransmit Attempts : 3
Global Encryption Key :
Dynamic Authorization UDP Port : 3799
Source IP Selection : Outgoing Interface
Tracking : Enabled

CPPM Identity :
Auth Acct DM/ Time | Server IP Addr Port Port CoA Window |Encryption Key OOBM
--------- --------- ------- ------- ---- ---- ----------- -------------- -------

#switch show runnig-config
radius-server dead-time infinite
radius-server tracking interval 60
```

show radius (stacked switch)

Use the `show radius` command to display RADIUS tracking information on a stacked switch.

```
Stack(switch)# show radius

 Status and Counters - General RADIUS Information

 Dead RADIUS server are preceded by *

  Deadtime (minutes)            : Infinite
  Timeout (seconds)             : 5
  Retransmit Attempts           : 2
  Global Encryption Key         :

  Dynamic Authorization UDP Port : 3799
  Source IP Selection           : Outgoing Interface
  Tracking                      : Enabled
  Tracking Period (seconds)     : 300
  CPPM Identity                 :

                Auth  Acct  DM/ Time   | Encryption Key  OOBM
  Server IP Addr  Port  Port  CoA Window |
---------------- ----- ----- ---- ------  --------------- ----- +
   20.1.1.100     1812  1813  No  300    | procurve          No

  *20.1.1.129     1812  1813  No  300    | procurve          No
```

radius-server tracking user-name

Syntax

```
[no] radius-server tracking user-name <USER-NAME>
```

Description

Configures the dummy user-name used to track the availability of the RADIUS servers. By default, the `radius-tracking-user` user name is used.

Options

`tracking`

Tracks the availability of the RADIUS servers.

`user-name`

Sets the username that will be used to track the RADIUS servers.

Specifiers

`<USER-NAME>`

Enter the username.

RADIUS server dead time

The RADIUS server dead time feature helps the client access request be serviced faster. RADIUS server dead time functions by skipping the client access request to a dead server then requesting access to the next active RADIUS server. The longevity to skip this dead server is controlled by a configuration value in minutes. Once the RADIUS server is marked as dead by the switch, it will not send any access requests to the RADIUS server, even if the server comes line within the configured time. Using the command `radius-server dead-tome infinite`, once the dead server comes alive, if tracked through radius tracking, switch will send access request to that time again.

radius-server dead-time

Syntax

`radius-server dead-time [infinite | `*`minutes`*`]`

Description

Enables the RADIUS Server dead-time to a specific time (minutes) or to infinity.

Restrictions

If radius tracking is enabled, only infinite is allowed for dead-time. For example, the command `(config)# radius-server dead-time 120` returns an error message `Dead-time value can't be configured when RADIUS service tracking feature is enabled`. RADIUS tracking must be configured to track infinitely unavailable RADIUS server.

Command context

`manager`

Parameters

INFINITE

   Selects infinite dead time.

MINUTES

   Specifies the dead time value in minutes. Length: 1 mins to 1440 mins.

Examples

Use the `show radius` command to display dead time.

```
show radius

 Status and Counters - General RADIUS Information
*
```

```
Server once reachable will be not be considered alive
  Deadtime (minutes)            : Infinite
  Timeout (seconds)             : 5
  Retransmit Attempts           : 3
  Global Encryption Key         :

  Dynamic Authorization UDP Port : 3799
  Source IP Selection           : Outgoing Interface
  Tracking                      : Enabled
  CPPM Identity                 :

                Auth  Acct  DM/ Time   |
  Server IP Addr Port  Port  CoA Window | Encryption Key |   OOBM
  --------- ---- ------ ----- --- ------ + --------------- --------
```

show radius

Syntax

show radius

Description

Shows the RADIUS-Server detail confirming dead time configuration.

Command context

manager

show RADIUS

```
switch# show radius
Status and Counters - General RADIUS Information
*Server once reachable will be not be considered alive
Deadtime (minutes) : Infinite
Timeout (seconds) : 5
Retransmit Attempts : 3
Global Encryption Key :
Dynamic Authorization UDP Port : 3799
Source IP Selection : Outgoing Interface
Tracking : Enabled
CPPM Identity :

Auth Acct DM/ Time Server IP Addr Port Port CoA Window Encryption Key OOBM
--------- -------- ------ ------- ---- -------- ------ -------------- ----

#switch show runnig-config
radius-server dead-time infinite
radius-server tracking interval 60
```

Configuring web-based authentication

Overview

1. If you have not already done so, configure a local username and password pair on the switch.

2. Identify or create a redirect URL for use by authenticated clients. Hewlett Packard Enterprise recommends that you provide a redirect URL when using web-based authentication. If a redirect URL is not specified, web browser behavior following authentication cannot be acceptable.

3. If you plan to use multiple VLANs with web-based authentication, ensure that these VLANs are configured on the switch and that the appropriate port assignments have been made. Confirm that the VLAN used by authorized clients can access the redirect URL.

4. Ping the switch console interface to ensure that the switch can communicate with the RADIUS server you have configured to support web-based authentication on the switch.

5. Configure the switch with the correct IP address and encryption key for RADIUS server access

6. (Optional) To use SSL encryption for web-based authentication login, configure and enable SSL on the switch.

7. Enable web-based authentication on the switch ports you want to use.

8. Configure the optional settings that you want to use for web-based authentication; for example:

   • To avoid address conflicts in a secure network, configure the base IP address and mask to be used by the switch for temporary DHCP addresses. You can also set the lease length for these temporary IP addresses.

   • To use SSL encryption for web-based authentication login, configure the SSL option.

   • To redirect authorized clients to a specified URL, configure the Redirect URL option.

9. Configure how web-based authenticator ports transmit traffic before they successfully authenticate a client and enter the authenticated state:

   • You can block incoming and outgoing traffic on a port before authentication occurs.

   • You can block only incoming traffic on a port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated ports configured for web-based authentication. For example, Wake-on-LAN traffic is transmitted on a web-based Authenticated egress port that has not yet transitioned to the authenticated state.

10. To ensure that web-based authentication works properly on the ports you have configured for port-access using web-based authentication, test both authorized and unauthorized access to your system.

NOTE Client web browsers cannot use a proxy server to access the network.

Configuration commands for web-based authentication

Controlled directions

Syntax

```
aaa port-access <port-list> [controlled-directions <both|in> mixed-mode|port-speed-vsa|mbv <enable|disable>]
```

After you enable web-based-based authentication on specified ports, you can use the `aaa port-access controlled-directions` command to configure how a port transmits traffic before it successfully authenticates a client and enters the authenticated state.

`both`: (Default) Incoming and outgoing traffic is blocked on a port configured for web-based authentication before authentication occurs.

`in`: Incoming traffic is blocked on a port configured for web-based authentication before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated ports configured for web-based authentication.

`mixed-mode`: Set if unauthenticated and authenticated users are allowed on the same port.

`port-speed-vsa`: Determines if the port speed HPE VSA is allowed and used on a port.

`mbv <enable|disable>`: Allows configuration of MBV (MAC-based VLANs) on a port. MBV allows multiple clients on different untagged VLANs to authenticate on the same port.

Requirements

As implemented in 802.1X authentication, the disabling of incoming traffic and transmission of outgoing traffic on a web-based Authenticated egress port in an unauthenticated state (using the `aaa port-access controlled-direction in` command) is supported only if the 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.

The port is configured as an edge port in the network using the `spanning-tree edge-port` command.

Notes

- For information on how to configure the prerequisites for using the `aaa port-access controlled-direction in` command, see "Multiple instance spanning-tree operations" in the advanced traffic management guide for your switch.

- To display the currently configured controlled direction value for web-based authenticated ports, enter the `show port-access web-based config` command.

- The `aaa port-access controlled-direction in` command allows Wake-on-LAN traffic to be transmitted on a web-based authenticated egress port that has not yet transitioned to the authenticated state; the `controlled-direction both` setting prevents Wake-on-LAN traffic to be transmitted on a web-based authenticated egress port until authentication occurs.

  The Wake-on-LAN feature is used by network administrators to remotely power on a sleeping workstation (for example, during early morning hours to perform routine maintenance operations, such as patch management and software updates).

- Using the `aaa port-access controlled-direction in` command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated egress ports that are configured for any of the following port-based security features:

  ◦ 802.1X authentication

  ◦ MAC authentication

  ◦ Web-based authentication

  Because a port can be configured for more than one type of authentication to protect the switch from unauthorized access, the last setting you configure with the `aaa port-access controlled-direction` command is applied to all authentication methods configured on the switch.

  For information about how to configure and use 802.1X authentication, see **Configuring Port and User-Based Access Control (802.1X)** on page 518.

- When a web-based authenticated port is configured with the `controlled-direction in` setting, eavesdrop prevention is not supported on the port.

Per Port Initial Role

Use the command `aaa port-access <port-list> initial-role <role-name>` to configure initial role at per port level.

aaa port-access initial-role

Syntax

```
aaa port-access <port-list> initial-role <role-name>
```

Description

Configure initial role at per port level.

Command context

```
config
```

Parameters

`port-list`

   Specify the list of ports.

`role-name`

   Provide the name for the initial role.

Usage

The priority is as follows:

• Per port initial role takes the highest priority (when configured)

• Global initial role (when configured)

• Default `denyall` user role

Examples

```
switch(config)# aaa port-access 1/1
 controlled-direction  Configure how traffic is controlled on non-authenticated
                       ports: in both directions (ingress and egress) or
                       ingress only.
 critical-auth         Configure critical-auth VLAN for authentication failed
                       ports due to non-reachable authentication server.
 initial-role          Configure an initial-role for a given port-range.
 mbv                   Allows configuration of MBV (MAC-based VLANs) on a port.
 mixed                 Set if unauthenticated and authenticated users are
                       allowed on the same port.
 open-auth             Configure open-auth VLAN for non-authenticated ports.
 port-speed-vsa        Set if port speed VSA processing is enabled on the port.
```

Disable web-based authentication

Syntax:

```
[no] aaa port-access web-based <port-list>
```

Enables web-based authentication on the specified ports. Use the `no` form of the command to disable web-based authentication on the specified ports.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Specifying the VLAN

Syntax:

```
aaa port-access web-based <port-list> [auth-vid <vid>]
```

```
[no] aaa port-access web-based <port-list> [auth-vid <vid>]
```

Specifies the VLAN to use for an authorized client. The Radius server can override the value (accept-response includes a vid). If `auth-vid` is `0`, no VLAN changes occur unless the RADIUS server supplies one.

Use the `[no]` form of the command to set the `auth-vid` to `0`. (Default: `0`)

Maximum authenticated clients

Syntax

```
aaa port-access web-based <port-list> [client-limit <1-256>]
```

Specifies the maximum number of authenticated clients to allow on the port. (Default: `1`)

> **NOTE**
> On switches where Web-based authentication and 802.1X can operate concurrently, this limit includes the total number of clients authenticated through both methods. The limit of 256 clients only applies when there are fewer than 16,384 authentication clients on the entire switch. After the limit of 16,384 clients is reached, no additional authentication clients are allowed on any port for any method.

Specifies base address

Syntax:

```
aaa port-access web-based [dhcp-addr <ip-address/mask>]
```

Specifies the base address/mask for the temporary IP pool used by DHCP. The base address can be any valid IP address (not a multicast address). Valid mask range value is `<255.255.240.0 - 255.255.255.0>`. (Default: `192.168.0.0/255.255.255.0`)

Specifies lease length

Syntax:

```
aaa port-access web-based [dhcp-lease <5-25>]
```

Specifies the lease length, in seconds, of the temporary IP address issued for Web-Auth login purposes. (Default: 10 seconds)

Specifying the period

Syntax:

```
aaa port-access web-based <port-list> [logoff-period]<60-9999999>]
```

Specifies the period, in seconds, that the switch enforces for an implicit logoff. This parameter is equivalent to the MAC age interval in a traditional switch sense. If the switch does not see activity after a logoff-period interval, the client is returned to its pre-authentication state. (Default: 300 seconds)

Specifying the number of authentication attempts

Syntax:

```
aaa port-access web-based <port-list> [max-requests] <1-10>]
```

Specifies the number of authentication attempts that must time-out before authentication fails. (Default: 2)

Specifying maximum retries

Syntax:

```
aaa port-access web-based <port-list> [max-retries] <1-10>]
```

Specifies the number of the number of times a client can enter their user name and password before authentication fails. This allows the reentry of the user name and password if necessary. (Default: 3)

Specifying the time period

Syntax:

```
aaa port-access web-based <port-list> [quiet-period] <1-65535>]
```

Specifies the time period (in seconds) the switch uses before sending an authentication request for a client that failed authentication. (Default: 60 seconds)

Specifying the re-authentication period

Syntax:

```
aaa port-access web-based <port-list> [reauth-period] <0-9999999>]
```

Specifies the time period, in seconds, the switch enforces on a client to re-authenticate. When set to 0, reauthentication is disabled. (Default: 300 seconds)

Specifying a forced reauthentication

Syntax:

```
aaa port-access web-based <port-list> [reauthenticate]
```

Forces a re-authentication of all attached clients on the port.

Specifying the URL

Syntax:

```
aaa port-access web-based <port-list> [redirect-url <url>]
```

```
[no] aaa port-access web-based <port-list> [redirect-url]
```

Specifies the URL that a user is redirected to after a successful login. Any valid, fully-formed URL can be used, for example, http://welcome-server/welcome.htm or http://192.22.17.5. Hewlett Packard Enterprise recommends that you provide a redirect URL when using web authentication.

> **NOTE**
>
> The `redirect-url` command accepts only the first 103 characters of the allowed 127 characters.

Use the `[no]` form of the command to remove a specified redirect URL.

(Default: There is no default URL. Browser behavior for authenticated clients can not be acceptable.)

Specifying the timeout

Syntax:

```
aaa port-access web-based <port-list> [server-timeout <1-300>]
```

Specifies the period, in seconds, the switch waits for a server response to an authentication request. Depending on the current `max-requests` value, the switch sends a new attempt or ends the authentication session. (Default: 30 seconds)

Show commands for web-based authentication

show port-access web-based command

Syntax:

```
show port-access web-based [port-list]
```

Displays the status of all ports or specified ports that are enabled for web-based authentication. The information displayed for each port includes:

- Number of authorized and unauthorized clients.

- VLAN ID number of the untagged VLAN used. If the switch supports MAC (untagged) VLANs, `MACbased` is displayed to show that multiple untagged VLANs are configured for authentication sessions.

- If tagged VLANs (statically configured or RADIUS-assigned) are used (`Yes` or `No`.)

- If client-specific per-port CoS (Class of Service) values are configured (`Yes` or `No`) or the numerical value of the CoS (802.1p priority) applied to all inbound traffic. For client-specific per-port CoS values, enter the `show port-access web-based clients detailed` command.

- If per-port rate-limiting for inbound traffic is applied (`Yes` or `No`) or the percentage value of the port's available bandwidth applied as a rate-limit value.

- If RADIUS-assigned ACLs are applied.

Information on ports not enabled for web authentication is not displayed.

Example:

```
Switch (config)# show port-access web-based

 Port Access Web-Based Status

      Auth      Unauth   Untagged Tagged Port     % In   RADIUS
Port  Clients   Clients  VLAN     VLANs  COS      Limit  ACL
----- --------  -------- -------- ------ -------- ------ ------
1     1         1        4006     Yes    70000000 100    Yes
2     2         0        MACbased No     Yes      Yes    Yes
3     4         0        1        Yes    No       No     No
```

show port-access web-based clients command

Syntax:

```
show port-access web-based clients [port-list]
```

Displays the session status, name, and address for each web-authenticated client on the switch. The IP address displayed is taken from the DHCP binding table (learned through the DHCP Snooping feature).

If DHCP snooping is not enabled on the switch, n/a (not available) is displayed for a client's IP address.

If a web-authenticated client uses an IPv6 address, n/a - IPv6 is displayed.

If DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table, n/a - no info is displayed.

Example:

```
HP Switch (config)# show port-access web-based clients

 Port Access Web-Based Client Status

Port  Client Name  MAC Address    IP Address      Session Status
----- ------------ -------------- --------------- -------------
1     webuser1     0010b5-891a9e  192.192.192.192 Authenticated
1     webuser2     001560-b3ea48  n/a - no info   Authenticating
1     webuser3     000000-111111  n/a - IPv6      Authenticating
3     webuser4     000000-111112  n/a             Authenticating
```

show port-access web-based clients detailed command

Syntax:

```
show port-access web-based clients <port-list> detailed
```

Displays detailed information on the status of web-based authenticated client sessions on specified switch ports.

This syntax shows session status, name, and address for each web-based authenticated client on the switch. The IP address displayed is taken from the DHCP binding table, learned through DHCP snooping. The following can appear if the client's IP address is not available:

`n/a` — DHCP snooping is not enabled on the switch; `n/a` is displayed for a client's IP address.

`n/a-IPv6` — a web-based authenticated client uses an IPv6 address.

`n/a-no info` — DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table.

Example:

```
HP Switch (config)# show port-access web-based clients 1 detailed

 Port Access Web-Based Client Status Detailed

  Client Base Details :
   Port            : 1
   Session Status : authenticated  Session Time(sec) : 6
   Username        : webuser1       MAC Address        : 0010b5-891a9e
   IP              : n/a

 Access Policy Details :
  COS Map         : 11111111        In Limit %         : 98
  Untagged VLAN   : 4006            Out Limit %        : 100
  Tagged VLANs    : 1, 3, 5, 6, 334, 2566
  RADIUS-ACL List :
```

```
        deny in udp from any to 10.2.8.233 CNT
            Hit Count: 0
        permit in udp from any to 10.2.8.233 CNT
            Hit Count: 0
        deny in tcp from any to 10.2.8.233 CNT
            Hit Count: 0
        permit in tcp from any to 10.2.8.233 CNT
            Hit Count: 0
        permit in tcp from any to 0.0.0.0/0 CNT
            Hit Count: 0
```

show port-access web-based config command

Syntax:

```
show port-access web-based config [port-list]
```

Displays the currently configured web-based authentication settings for all switch ports or specified ports, including:

- Temporary DHCP base address and mask.

- Support for RADIUS-assigned dynamic VLANs (`Yes` or `No`).

- Controlled direction setting for transmitting Wake-on-LAN traffic on egress ports.

- Authorized and unauthorized VLAN IDs.

If the authorized or unauthorized VLAN ID value is `0`, the default VLAN ID is used unless overridden by a RADIUS-assigned value.

Example:

```
HP Switch (config)# show port-access web-based config

Port Access Web-Based Configuration

 DHCP Base Address : 192.168.0.0
 DHCP Subnet Mask  : 255.255.255.0
 DHCP Lease Length : 10
 Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No
 Access Denied Message  : System Default

              Client Client Logoff  Re-Auth Unauth   Auth     Cntrl
Port  Enabled Limit  Moves  Period  Period  VLAN ID  VLAN ID  Dir
----- -------- ------ ------ ------- ------- -------- -------- -----
1     Yes      1      No     300     0       0        0        both
2     Yes      1      No     300     0       0        0        in
```

show port-access web-based config detailed command

Syntax:

```
show port-access web-based config <port-list> detailed
```

Displays more detailed information on the currently configured web-based authentication settings for specified ports.

Example:

```
Switch (config)# show port-access web-based config 1 detailed

 Port Access Web-Based Detailed Configuration
```

```
Port          : 1         Web-based enabled : Yes
Client Limit  : 1         Client Moves      : No
Logoff Period : 300       Re-Auth Period    : 0

Unauth VLAN ID : 0        Auth VLAN ID      : 0

Max Requests  : 3         Quiet Period      : 60
Server Timeout : 30

Max Retries   : 3         SSL Enabled       : No
Redirect URL :
```

show port-access web-based config auth-server command

Syntax:

```
show port-access web-based config [port-list] auth-server
```

Displays the currently configured web authentication settings for all switch ports or specified ports and includes RADIUS server-specific settings, such as:

- Timeout waiting period.

- Number of timeouts supported before authentication login fails.

- Length of time (quiet period) supported between authentication login attempts.

Example:

```
Switch (config)# show port-access web-based config auth-server

Port Access Web-Based Configuration

              Client Client Logoff  Re-Auth  Max  Quiet   Server
 Port  Enabled Limit  Moves  Period  Period   Req  Period  Timeout
 ----- -------- ------ ------ ------- -------- ---- ------- --------
 1     Yes      1      No     300     0        3    60      30
 2     No       1      No     300     0        3    60      30
 ...
```

show port-access web-based config web-server command

Syntax:

```
show port-access web-based config [port-list] web-server
```

Displays the currently configured Web Authentication settings for all ports or specified ports, including web-specific settings for password retries, SSL login status, and a redirect URL, if specified.

Configuring MAC authentication

Preparation for configuring MAC authentication

Before you configure MAC authentication:

Procedure

1. Configure a local username and password on the switch.

2. Ensure that the VLANs are configured on the switch and that the appropriate port assignments have been made if you plan to use multiple VLANs with MAC authentication.

3. Ping the switch console interface to ensure that the switch is able to communicate with the RADIUS server you are configuring to support MAC authentication.

4. Configure the switch with the correct IP address and encryption key to access the RADIUS server.

5. Configure the switch for MAC authentication with the ports you will be using.

6. Test both the authorized and unauthorized access to your system to ensure that MAC authentication works properly on the ports you have chosen to configure for port-access.

Configuration commands for MAC authentication

Configuring the global MAC authentication password

MAC authentication requires that only a single entry containing the username and password is placed in the user database with the device's MAC address. This creates an opportunity for malicious device spoofing. The global password option configures a common MAC authentication password to use for all MAC authentications sent to the RADIUS server. This makes spoofing more difficult.

It is important that when implementing the global MAC authentication password option, that the user database on the RADIUS server has this password as the password for each device performing MAC authentication.

Syntax:

```
[no] aaa port-access mac-based password <password-value>
```

Specifies the global password to be used by all MAC authenticating devices.

The [no] form of the command disables the feature.

Configuring a global MAC authentication password

```
HP Switch(config)# aaa port-access mac-based password secretMAC1

HP Switch(config)# show port-access mac-based config

Port Access MAC-Based Configuration

MAC Address Format : no-delimiter
Password           : secretMAC1

Unauth Redirect Configuration URL :

Unauth Redirect Client Timeout (sec) : 1800
Unauth Redirect Restrictive Filter : Disabled
Total Unauth Redirect Client Count : 0

            Client Client Logoff  Re-Auth Unauth  Auth    Cntrl
Port  Enabled Limit  Moves  Period  Period  VLAN ID VLAN ID Dir
----- ------- ------ ------ ------- ------- ------- ------- -----
1     No      1      No     300     0       0       0       both
2     No      1      No     300     0       0       0       both
3     No      1      No     300     0       0       0       both
4     No      1      No     300     0       0       0       both
```

```
5       No      1       No      300     0       0       0       both
6       No      1       No      300     0       0       0       both
7       No      1       No      300     0       0       0       both
8       No      1       No      300     0       0       0       both
```

NOTE: The password value will display in an exported config file when `include-credentials` is enabled.

Configuring a MAC-based address format

Syntax:

```
aaa port-access mac-based addr-format <no-delimiter|single-dash|multi-dash|multi-
colon| no-delimiter-uppercase|single-dash-uppercase| multi-dash-uppercase|multi-
colon-uppercase>
```

Specifies the MAC address format used in the RADIUS request message. This format must match the format used to store the MAC addresses in the RADIUS server.

Default: `no-delimiter`

`no-delimiter`: specifies an aabbccddeeff format.

`single-dash`: specifies an aabbcc-ddeeff format.

`multi-dash`: specifies an aa-bb-cc-dd-ee-ff format.

`multi-colon`: specifies an aa:bb:cc:dd:ee:ff format.

`no-delimiter-uppercase`: specifies an AABBCCDDEEFF format.

`single-dash-uppercase`: specifies an AABBCC-DDEEFF format

`multi-dash-uppercase`: specifies an AA-BB-CC-DD-EE-FF format

`multi-colon-uppercase`: specifies an AA:BB:CC:DD:EE:FF format.

Creating a custom delimiter for a MAC address

A custom MAC delimiter can be configured which instructs all switches to accept MAC addresses only in the specified format.

mac-delimiter

Syntax

```
mac-delimiter {default | colon | hyphen | oui-nic | none}

no mac-delimiter {colon | hyphen | oui-nic}
```

Description

Configures a custom delimiter for all MAC address.

The `no` form of the command returns the MAC delimiter to its default which is to use no delimiter, displaying MAC addresses as XXXXXXXXXXXX.

Command context

```
config
```

Parameters

default

    MAC will be in traditional `PVOS` style which is XXXXXXXXXXXX.

colon

    Specifies the MAC delimiter format as XX:XX:XX:XX:XX:XX.

hyphen

    Specifies the MAC delimiter format as XX-XX-XX-XX-XX-XX.

oui-nic

    Specifies the MAC delimiter format as XXXXXX-XXXXXX.

none

    Specifies the MAC address without a delimiter.

Usage

The MAC address is case insensitive. The format of the MAC delimiter can be accepted in lowercase, upper case, or mixed case.

Use the `show running-config` command to display the MAC address in the configured format.

Examples

By using the command `mac-delimiter hyphen`, the MAC delimiter is now configured as hyphen. To verify the change to MAC delimiter, use `show mac-address`.

```
switch(config)# mac-delimiter hyphen

switch(config)# show mac-address
Status and Counters - Port Address Table
MAC Address Port VLAN
---------------- ----------------------------- ----
68-b5-99-a2-91-80 1 1
68-b5-99-a2-91-98 1 1
```

By using the command `mac-delimiter colon`, the MAC delimiter is now configured as colon. To verify the change to MAC delimiter, use `show mac-address`.

```
switch(config)# mac-delimiter colon

switch(config)# show mac-address
Status and Counters - Port Address Table
MAC Address Port VLAN
---------------- ----------------------------- ----
68:b5:99:a2:91:80 1 1
68:b5:99:a2:91:98 1 1
```

By using the command `mac-delimiter oui-nic`, the MAC delimiter is now configured with one hyphen (XXXXXX-XXXXXX). To verify the change to MAC delimiter, use `show mac-address`.

```
switch(config)# mac-delimiter oui-nic

switch(config)# show mac-address
Status and Counters - Port Address Table
MAC Address Port VLAN
---------------- ----------------------------- ----
```

```
68b599-a29180 1 1
68b599-a29198 1 1
```

By using the command `mac-delimiter default`, the MAC delimiter configuration is set to default. To verify the change to MAC delimiter, use `show mac-address`.

```
switch(config)# mac-delimiter default

switch(config)# show mac-address
Status and Counters - Port Address Table
MAC Address Port VLAN
---------------- ------------------------------ ----
68b599a29180 1 1
68b599a29198 1 1
```

By using the command `mac-delimiter none`, the MAC delimiter configuration is set to default. To verify the change to MAC delimiter, use `show mac-address`.

```
switch(config)# mac-delimiter none

switch(config)# show mac-address
Status and Counters - Port Address Table
MAC Address Port VLAN
---------------- ------------------------------ ----
68b599-a29180 1 1
68b599-a29198 1 1 <<<< Set to default
```

Configuring other MAC-based commands
Enabling/disabling MAC authentication

Syntax:

```
[no] aaa port-access mac-based <port-list>
```

Enables MAC authentication on specified ports.

Use the `no` form of the command to disable MAC authentication on specified ports.

Specifying the maximum number of authenticated MACs allowed on a port

Syntax:

```
aaa port-access mac-based <port-list> [addr-limit <1-256>]
```

Specifies the maximum number of authenticated MACs to allow on the port.

Default: `1`

NOTE | On switches where MAC authenticated and 802.1X operate concurrently, this limit includes the total number of clients authenticated through both methods.

The limit of 256 clients only applies when there are fewer than 16,384 authentication clients on the entire switch. After the limit of 16,384 clients is reached, no additional authentication clients are allowed on any port for any method.

Allowing addresses to move without re-authentication

Syntax:

```
[no] aaa port-access mac-based <port-list> [addr-moves]
```

Allows client moves between the specified ports under MAC authenticated control. When enabled, the switch allows addresses to move without requiring a re-authentication.

When disabled, the switch does not allow moves and when one occurs, the user will be forced to re-authenticate. At least two ports (from ports and to ports) must be specified.

Use the `no` form of the command to disable MAC address moves between ports under MAC authenticated control.

Default: Disabled – no moves allowed

Specifying the VLAN for an authorized client

Syntax:

```
aaa port-access mac-based <port-list> [auth-vid <vid>]
```

```
[no] aaa port-access mac-based <port-list> [auth-vid]
```

Specifies the VLAN to use for an authorized client. The RADIUS server can override the value (accept response includes avid).

If `auth-vid` is `0`, no VLAN changes occur unless the RADIUS server supplies one.

Use the `no` form of the command to set the `auth-vid` to `0`.

Default:`0`

Specifying the time period enforced for implicit logoff

Syntax:

```
[no] aaa port-access mac-based <port-list> [logoff-period] <60-9999999>]
```

Specifies the period, in seconds, that the switch enforces for an implicit logoff. This parameter is equivalent to the MAC age interval in a traditional switch sense. If the switch does not see activity after a logoff-period interval, the client is returned to its pre-authentication state.

Default: `300 seconds`

Specifying how many authentication attempts can time-out before failure

Syntax:

```
[no] aaa port-access mac-based <port-list> [max-requests <1-10>]
```

Specifies the number of authentication attempts that must time-out before authentication fails.

Default: `2`

Specifying how long the switch waits before processing a request from a MAC address that failed authentication

Syntax:

```
[no] aaa port-access mac-based <port-list> [quiet-period <1-65535>]
```

Specifies the time period (in seconds) that the switch waits before processing an authentication request from a MAC address that failed authentication.

Default: `60 seconds`

Specifying time period enforced on a client to re-authenticate

Syntax:

```
[no] aaa port-access mac-based <port-list> [reauth-period <0-9999999>]
```

Specifies the time period (in seconds) that the switch enforces on a client to re-authenticate. The client remains authenticated while the re-authentication occurs.

When set to `0`, re-authentication is disabled.

Default: `300 seconds`

Forcing re-authentication of clients

Syntax:

```
[no] aaa port-access mac-based <port-list> [reauthenticate]
```

Forces a re-authentication of all attached clients on the port.

Specifying how long the switch waits for a server response

Syntax:

```
[no] aaa port-access mac-based <port-list> [server-timeout <1-300>]
```

Specifies the period, in seconds, the switch waits for a server response to an authentication request. Depending on the current `max-requests` value, the switch sends a new attempt or ends the authentication session.

Default: `30 seconds`

Setting the period of time the switch waits before moving the port to the VLAN for unauthenticated clients

Syntax:

```
aaa port-access mac-based <port-list> [unauth-period]
```

Sets the period of time the switch waits before moving the port to the VLAN for unauthenticated clients.

Specifying the VLAN to use when authentication fails

Syntax:

```
aaa port-access mac-based <port-list> [unauth-vid <vid>]
```

```
[no] aaa port-access mac-based <port-list> [unauth-vid]
```

Specifies the VLAN to use for a client that fails authentication. If `unauth-vid` is `0`, no VLAN changes occur. Use the `no` form of the command to set the `unauth-vid` to `0`.

Default: `0`

Configuring custom messages for failed logins

This feature allows administrators to configure custom messages that are displayed when authentication with the RADIUS server fails. The messages are appended to existing internal web pages that display during the authentication process. Messages can be configured using the CLI, or centrally using the RADIUS server, and can provide a description of the reason for a failure as well as possible steps to take to resolve the authentication issue. There is no change to the current web-based authentication functionality.

Syntax:

```
[no] aaa port-access web-based access-denied-message <<access-denied-str>|radius-response>
```

Specifies the text message (ASCII string) shown on the web page after an unsuccessful login attempt. The message must be enclosed in quotes.

The `[no]` form of the command means that no message is displayed upon failure to authenticate.

Default: The internal web page is used. No message will be displayed upon authentication failure.

`access-denied-str`: The text message that is appended to the end of the web page when there is an unsuccessful authentication request. The string can be up to 250 ASCII characters.

`radius-response`: Use the text message provided in the RADIUS server response to the authentication request.

Configuring an access denied message on the switch

```
HP Switch(config)# aaa port-access web-based access-denied-message
"Please contact your system administrator to obtain authentication
privileges."
```

Output showing the custom access denied message

```
HP Switch(config)# show port-access web-based config

Port Access Web-based Configuration

DHCP Base Address     : 192.168.0.0
DHCP Subnet Mask      : 255.255.248.0
DHCP Lease Length     : 10 seconds
Allow RADIUS-assigned dynamic (GVRP) VLANs[No]: Yes
Access Denied Message : Custom:
  Please contact your system administrator to obtain authentication
  privileges.
```

Port	Enabled	Client Limit	Client Moves	Logoff Period	Re-auth Period	Unauth VLAN ID	Auth VLAN ID	Ctrl Dir
1	Yes	1	No	300	60	1	2	both
2	Yes	18	No	999999999	999999999	0	0	both
3	Yes	22	No	999999999	999999999	4096	4096	both

Access denied message when radius-response is configured

```
HP Switch(config)# show port-access web-based config

Port Access Web-based Configuration

DHCP Base Address     : 192.168.0.0
DHCP Subnet Mask      : 255.255.248.0
DHCP Lease Length     : 10 seconds
Allow RADIUS-assigned dynamic (GVRP) VLANs[No]: Yes
Access Denied Message : Retrieved from Radius
```

Port	Enabled	Client Limit	Client Moves	Logoff Period	Re-auth Period	Unauth VLAN ID	Auth VLAN ID	Ctrl Dir
1	Yes	1	No	300	60	1	2	both

```
2    | Yes    18     No     300    999999999 0        0         both
3    | Yes    22     No     300    999999999 4096     4096      both
```

Unauthenticated clients can be assigned to a specific static, untagged VLAN (`unauth-vid`), to provide access to specific (guest) network resources. If no VLAN is assigned to unauthenticated clients, the port is blocked and no network access is available.

Web page display of access denied message

Figure 22: *Example of webpage when `unauth-vid` is configured*

Figure 23: *Example of webpage when `unauth-vid` is not configured*

The `show running-config` command displays the client information, including the configured access denied message.

Running configuration output displaying access denied message

```
HP Switch(config)# show running-config

Running configuration:

; Configuration Editor; Created on release #XX.15.XX

hostname "Switch"
web-management ssl
qos dscp-map 000000 priority 0
vlan 1
   name "DEFAULT_VLAN"
   untagged 1-14,19-24
   ip address dhcp-bootp
   no untagged 15-18
   exit
vlan 100
   name "auth-vid"
```

```
      untagged 15-18
      ip address dhcp-bootp
      exit
radius-server host 10.0.13.118 key 'secret'
aaa authentication port-access eap-radius
snmp-server community "public" Unrestricted
aaa port-access web-based 5
aaa port-access web-based 5 auth-vid 100
aaa port-access web-based 5 unauth-vid 1
aaa port-access web-based dhcp-addr 172.18.0.0 255.255.255.0
aaa port-access web-based access-denied-message "Please contact your system
administrator to obtain authentication privileges."
no autorun
```

Access denied message

Running configuration output when RADIUS response is configured

```
HP Switch(config)# show running-config

Running configuration:

; J9627A Configuration Editor; Created on release #XX.15.XX

hostname "Switch"
web-management ssl
qos dscp-map 000000 priority 0
vlan 1
   name "DEFAULT_VLAN"
   untagged 1-14,19-24
   ip address dhcp-bootp
   no untagged 15-18
   exit
vlan 100
   name "auth-vid"
   untagged 15-18
   ip address dhcp-bootp
   exit
radius-server host 10.0.13.118 key 'secret'
aaa authentication port-access eap-radius
snmp-server community "public" Unrestricted
aaa port-access web-based 5
aaa port-access web-based 5 auth-vid 100
aaa port-access web-based 5 unauth-vid 1
aaa port-access web-based dhcp-addr 172.18.0.0 255.255.255.0
aaa port-access web-based access-denied-message radius-response
```

RADIUS response

Viewing the show commands for MAC authentication

Syntax:

```
show port-access mac-based [port-list]
```

Displays the status of all ports or specified ports that are enabled for MAC authentication. The information displayed for each port includes:

- Number of authorized and unauthorized clients.

- VLAN ID number of the untagged VLAN used. If the switch supports MAC (untagged) VLANs, `MACbased` is displayed to show that multiple untagged VLANs are configured for authentication sessions.

- If tagged VLANs (statically configured or RADIUS-assigned) are used (`Yes` or `No`.)

- If client-specific per-port CoS (Class of Service) values are configured (`Yes` or `No`) or the numerical value of the CoS (802.1p priority) applied to all inbound traffic. For client-specific per-port CoS values, enter the `show port-access web-based clients detailed` command.

- If per-port rate-limiting for inbound traffic is applied (`Yes` or `No`) or the percentage value of the port's available bandwidth applied as a rate-limit value.

- If RADIUS-assigned ACLs are applied.

Information on ports not enabled for MAC authentication is not displayed.

Output for the show port-access mac-based command

```
HP Switch(config)# show port-access mac-based

 Port Access MAC-Based Status

     Auth    Unauth  Untagged Tagged Port     % In   RADIUS Cntl
 Port Clients Clients VLAN     VLANs  COS      Limit  ACL    Dir
 ---- ------- ------- -------- ------ -------- ------ ------ ------
 1    1       1       2003     Yes    70000000 100    Yes    both
 2    2       0       MACbased No     Yes      Yes    Yes    both
 3    4       0       1        Yes    No       No     No     both
```

Viewing session information for MAC authenticated clients on a switch

Syntax:

```
show port-access mac-based clients [port-list]
```

Displays the session status, name, and address for each MAC authenticated client on the switch. The IP address displayed is taken from the DHCP binding table (learned through the DHCP Snooping feature).

If DHCP snooping is not enabled on the switch, `n/a` (not available) is displayed for a client's IP address.

If a MAC-authenticated client uses an IPv6 address, `n/a - IPv6` is displayed.

If DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table, `n/a - no info` is displayed.

Output for the show port-access mac-based clients command

```
HP Switch(config)# show port-access mac-based clients

 Port Access MAC-Based Client Status

Port MAC Address    IP Address                              Session Status
---- ------------- --------------------------------------- -------------
1    001321-eb8063 2001:fecd:ba23:cd1f:dcb1:1010:9234:4088 unauthenticated
```

```
1     000000-111112 192.192.192.192                           authenticated
2     000000-111111 n/a                                       authenticating
```

Output for the show port-access mac-based clients command with a specific port

```
Switch(config)# show port-access mac-based clients 2
Port Access MAC-Based Client Status

Port   MAC Address  IP Address                         Mode Client Status
-----  ------------ ---------------------------------- ---- -------------
2      000000-000010 n/a                                User authenticated
2      000000-000100 n/a                                User authenticated
2      000000-000200 n/a                                User authenticated
2      000000-000300 n/a                                User authenticated
2      000000-000400 n/a                                User authenticated
2      000000-000500 n/a                                User authenticated
2      000000-000600 n/a                                User authenticated
2      000000-000700 n/a                                User authenticated
2      000000-000800 n/a                                User authenticated
2      000000-000900 n/a                                User authenticated
2      000000-000a00 n/a                                User authenticated
2      000000-000b00 n/a                                User authenticated
2      000000-000c00 n/a                                User authenticated
2      000000-000d00 n/a                                User authenticated
2      000000-000e00 n/a                                User authenticated
2      000000-000f00 n/a                                User authenticated
2      000000-001000 n/a                                User authenticated
2      000000-001100 n/a                                User authenticated
2      000000-001200 n/a                                User authenticated
2      000000-001300 n/a                                User authenticated
2      000000-001400 n/a                                User authenticated
2      000000-001500 n/a                                User authenticated
2      000000-001600 n/a                                User authenticated
2      000000-001700 n/a                                User authenticated
2      000000-001800 n/a                                User authenticated
2      000000-001900 n/a                                User authenticated
```

Viewing detail on status of MAC authenticated client sessions

Syntax:

```
show port-access mac-based clients <port-list> detailed
```

Displays detailed information on the status of MAC authenticated client sessions on specified ports.

This syntax shows session status, name, and address for each web-based authenticated client on the switch. The IP address displayed is taken from the DHCP binding table, learned through DHCP snooping. The following can appear if the client's IP address is not available:

`n/a` — DHCP snooping is not enabled on the switch; `n/a` is displayed for a client's IP address.

`n/a-IPv6` — a web-based authenticated client uses an IPv6 address.

`n/a-no info` — DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table.

Output for the show port-access mac-based clients detailed command

```
HP Switch(config)# show port-access mac-based clients 1 detailed

Port Access MAC-Based Client Status Detailed
```

```
 Client Base Details :
  Port          : 1
  Session Status : authenticated   Session Time(sec) : 6
  Username      : client1       MAC Address      : 0010b5-891a9e
  IP            : n/a

 Access Policy Details :
  COS Map      : 12345678              In Limit % : 98
  Untagged VLAN : 4006                 Out Limit % : 100
  Tagged VLANs  : 1, 3, 5, 6, 334, 4001
  RADIUS-ACL List :
    deny in udp from any to 10.2.8.233 CNT
      Hit Count: 0
    permit in udp from any to 10.2.8.233 CNT
      Hit Count: 0
    deny in tcp from any to 10.2.8.233 CNT
      Hit Count: 0
    permit in tcp from any to 10.2.8.233 CNT
      Hit Count: 0
    permit in tcp from
```

Output for the show port-access mac-based clients command with a specific port (detailed)

```
Switch(config)# show port-access mac-based clients 2 detailed

Port Access MAC-Based Client Status Detailed

Client Base Details :
Port : 2
Client Status : authenticated        Session Time : 65 seconds
MAC Address : 000000-000010          Session Timeout : 0 seconds
IP : n/a

Access Policy Details :
COS Map : Not Defined                In Limit Kbps : Not Set
Untagged VLAN : 1                    Out Limit Kbps : Not Set
Tagged VLANs : No Tagged VLANs
Port Mode : 1000FDx                  Auth Mode : User-based
RADIUS ACL List :
permit in ip from any to 10.10.10.20

Client Base Details :
Port : 2
Client Status : authenticated        Session Time : 64 seconds
MAC Address : 000000-000100          Session Timeout : 0 seconds
IP : n/a

Access Policy Details :
COS Map : Not Defined                In Limit Kbps : Not Set
Untagged VLAN : 1                    Out Limit Kbps : Not Set
Tagged VLANs : No Tagged VLANs
Port Mode : 1000FDx                  Auth Mode : User-based
RADIUS ACL List :
permit in ip from any to 10.10.10.20

Client Base Details :
Port : 2
Client Status : authenticated        Session Time : 65 seconds
MAC Address : 000000-000200          Session Timeout : 0 seconds
IP : n/a
```

```
Access Policy Details :
COS Map : Not Defined                    In Limit Kbps : Not Set
Untagged VLAN : 1                        Out Limit Kbps : Not Set
Tagged VLANs : No Tagged VLANs
Port Mode : 1000FDx                      Auth Mode : User-based
RADIUS ACL List :
permit in ip from any to 10.10.10.20
...
```

Error log

Error	Message
RMON_DCA_ILLEGAL_VSA_VALUE	Invalid value of attribute '%s' received on port %s for %s client %s.
	Example:
	Invalid value of attribute HP-Port-Auth-Mode-MA received on port A3 for MacAuth client 001234-567890.
RMON_CONFIG_VSA_CONFLICT	Conflict occurred between configuration and attribute '%s' received on port %s for %s client %s.
	Example:
	Conflict occurred between configuration and attribute HP-Port-Auth-Mode-MA received on port A3 for MacAuth client 001234-567890.
RMON_DCA_ILLEGAL_VSA_COMBINATION	Illegal combination of attributes '%s' and '%s' received on port %s for %s client %s.
	Example:
	Illegal combination of attributes HP-Port-Auth-Mode-MA and HP-Port-Auth-Mode-Dot1x received on port A3 for 802.1X client 001234-567890.

Viewing MAC authentication settings on ports

Syntax:

```
show port-access mac-based config [port-list]
```

Displays the currently configured MAC authentication settings for all switch ports or specified ports, including:

- MAC address format

- Support for RADIUS-assigned dynamic VLANs (`Yes` or `No`)

- Controlled direction setting for transmitting Wake-on-LAN traffic on egress ports

- Authorized and unauthorized VLAN IDs

If the authorized or unauthorized VLAN ID value is `0`, the default VLAN ID is used unless overridden by a RADIUS-assigned value.

Output for the show port-access mac-based Configuration command

```
HP Switch(config)# show port-access mac-based config

Port Access MAC-Based Configuration

 MAC Address Format : no-delimiter
 Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

              Client Client Logoff Re-Auth Unauth   Auth    Cntrl
 Port  Enabled Limit  Moves  Period Period  VLAN ID  VLAN ID Dir
 ----- -------- ------ ------ ------ ------- -------- ------- ----
 1     No       1      No     300    0       0        0       both
 2     Yes      1      No     300    0       0        0       in
```

Viewing details of MAC Authentication settings on ports

Syntax:

```
show port-access mac-based config <port-list> detailed
```

Displays more detailed information on the currently configured MAC authentication settings for specified ports.

Output for the show port-access mac-based config detail command

```
HP Switch(config)# show port-access mac-based config 1 detailed

 Port Access MAC-Based Detailed Configuration

  Port          : 1           Web-based enabled : Yes
  Client Limit  : 1           Client Moves      : No
  Logoff Period : 300         Re-Auth Period    : 0

  Unauth VLAN ID : 0          Auth VLAN ID      : 0

  Max Requests  : 3           Quiet Period      : 60
  Server Timeout : 30
```

Viewing MAC Authentication settings including RADIUS server-specific

Syntax:

```
show port-access mac-based config [port-list] auth-server
```

Displays the currently configured web authentication settings for all switch ports or specified ports and includes RADIUS server-specific settings, such as:

• Timeout waiting period.

• Number of timeouts supported before authentication login fails.

• Length of time (quiet period) supported between authentication login attempts.

Output for the show port-access mac-based config auth-server command

```
HP Switch(config)# show port-access mac-based config auth-server

Port Access MAC-Based Configuration
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

```
                 Client Client Logoff   Re-Auth   Max   Quiet    Server
Port  Enabled    Limit  Moves  Period   Period    Req   Period   Timeout
-----  -------   ------ ------ -------  --------  ----  -------  --------
1      No         1      No     300       0        3     60       30
2      No         1      No     300       0        3     60       30
3      Yes        1      No     300       0        3     60       30
...
```

Client status

The table below shows the possible client status information that can be reported by a web-based or MAC-based `show... clients` command.

Reported status	Available network connection	Possible explanations
authenticated	Authorized VLAN	Client authenticated. Remains connected until logoff-period or reauth-period expires.
authenticating	Switch only	Pending RADIUS request.
rejected-no vlan	No network access	- Invalid credentials supplied. - RADIUS Server difficulties. See log file. - If unauth-vid is specified it cannot be successfully applied to the port. An authorized client on the port has precedence.
rejected-unauth vlan	Unauthorized VLAN only	- Invalid credentials supplied. - RADIUS Server difficulties. See log file.
timed out-no vlan	No network access	RADIUS request timed out. If unauth-vid is specified it cannot be successfully applied to the port. An authorized client on the port has precedence. Credentials resubmitted after quiet-period expires.
timed out-unauth vlan	Unauthorized VLAN only	RADIUS request timed out. After the quiet-period expires credentials are resubmitted when client generates traffic.
unauthenticated	Switch only	Waiting for user credentials.

Configuring MAC pinning

MAC pinning allows administrators to persist authenticated clients by disabling the logoff period associated with the client. The feature is available for clients that use MAC Auth or Local MAC Authentication. During port-flaps and switch reboot, the pinned authenticated client entries will be de-authenticated until those clients reauthenticate.

MAC pinning is disabled by default and can be enabled on a per-port basis or enabled on range of ports. The primary use case for using MAC pinning is for legacy devices such as printers or medical devices that remain silent on the network resulting in de-authentication of those clients.

Restrictions

- This feature is mutually exclusive with port-security learn-mode configurations. Learn-mode can only be set as "continuous" when MAC pinning is enabled on LMA or MAC-based port. If MAC pinning is enabled, port-securities learn mode can be set to" continuous" and" port-access".

- MAC pinning is mutually exclusive with port-security learn-mode configurations. When MAC pinning is enabled, port-security learn-mode configurations must be set as "continuous".

Configuration use cases

- When a client enables LMA with MAC pinning and 802.1x authentication on a port, the MAC address is pinned. If that client tries to authenticate through the 802.1x authentication method, MAC pinning will not function. When MAC pinning does not function, the client must de-authenticated from LMA and reauthenticated through 802.1x which will then take precedence over LMA authentication. The client must check the concurrent auth with the default logoff period of 300 sec.

- When a client enables LMA with MAC-pinning and MAC-based authentication on a port, the MAC-address is pinned through the LMA authentication. If that same client tries to authentication through MAC-based authentication, the LMA authentication takes precedence. No MAC-based authentication clients will be added and MAC-pinning will stay in effect.

- When a client enables LMA with MAC pinning and 802.1x authentication on a port with a logoff period, the client is authenticated through LMA and the MAC address is pinned. The client is then authenticated through both LMA and 802.1x. Once the 802.1x authentication completes, the client must de-authenticate from LMA. The client then configures the logoff period and checks the concurrent Auth between LMA and Dot1x.

- When LMA with MAC pinning has been enabled on a port and the eport is powered down, or power cycles, the client is de-authenticated. When the port is powered up, the client will be re-authenticated when reachable.

- If MAC pinning is disabled on a port, the clients are subjected to log off period behavior when the client is removed from the port.

aaa port-access local-mac <PORT-LIST> mac-pin

Syntax

```
aaa port-access local-mac <PORT-LIST> mac-pin

no aaa port-access local-mac <PORT-LIST> mac-pin
```

Description

Enables MAC pinning configuration on a port or list of ports for LMA-based authentication methods.

The `no` form of this command disables MAC pinning on the port or list or ports.

Command context

```
config
```

Parameters

<PORT-LIST>

Specifies the port or list of ports that will be configured for MAC pinning.

Examples
MAC pinning configuration:

```
switch (config)# aaa port-access local-mac 1
aaa port-access local-mac 1 mac-pin
```

aaa port-access mac-based *<PORT-LIST>* mac-pin

Syntax

```
aaa port-access mac-based <PORT-LIST> mac-pin
```

```
no aaa port-access mac-based <PORT-LIST> mac-pin
```

Description

Enables MAC pinning on a port or list of ports for MAC based authentication methods.

The `no` form of this command disables MAC pinning on a port or list of ports for MAC based authentication methods.

Command context
`config`

Parameters

<PORT-LIST>

Specifies the port or list of ports that will be configured for MAC pinning.

Examples

Configuration of MAC pinning for MAC based authentication.

```
switch (config) # aaa port-access mac-based 1
aaa port-access mac-based 1 mac-pin
aaa port-access mac-based 2-9
aaa port-access mac-based 2-9 mac-pin
```

Overview

Local MAC Authentication (LMA) is a software feature that simplifies deployment for devices such as IP phones and security cameras. In general, it provides dynamic attribute assignment (e.g., VLAN and QoS) through the use of a locally configured authentication repository. The most common use model for LMA is to automatically assign a VLAN to IP phones. In some cases, it can also provide rudimentary access security for the network.

While there are other network technologies that can be used to deploy IP phones (MAC Authentication and IEEE 802.1X), deployment is complex. LMA however is relatively simple to deploy yet offers adequate security for most uses.

Additionally, LMA can be used in environments that deploy a mix of legacy and newer IP phones, even though in the past legacy IP phones did not support newer technologies such as LLDP-MED and IEEE 802.1X.

Concepts

LMA solves dynamic assignment of per client (mac-address) attributes without having to create RADIUS infrastructure. It also allows the user to define authentication polices based on the MAC OUI and MAC/mask, which simplifies management of devices by removing the need to create a policy on a per device basis.

LMA is an addition to existing client authentication methods. Users can configure multiple authentication methods (802.1X, LMA, Mac auth (radius), web-auth (radius)) on a single port concurrently. When multiple authentication methods are configured on a single port the precedence of authentication methods is (right to left): 802.1X -> LMA -> web auth/Mac auth. This means:

- When 802.1X and LMA are enabled on a port, the policy configured for 802.1X takes precedence over LMA.

- When LMA and Mac-auth (radius) are enabled on a port, the policy configured for LMA takes precedence over Mac-auth radius.

- When only LMA is enabled on a port, client access is subjected to the LMA profile configuration.

LMA supports defining configuration profiles called LMA profiles and mac-groups, which significantly reduce the number of configuration entries during Authentication. There are two types of profiles:

- applied – a profile applied to a mac-group

- provisioned – a profile not applied to a group, however the user can use this profile later

LMA mac-groups group different types of mac entities - mac-address, mac-mask and mac-oui.

Possible scenarios for deployment

The following are examples of possible scenarios where LMA can be deployed.

1. In the following scenario multiple clients are connected to a hub that is tagged to vlan "A" and untagged to vlan "B". The hub is attached to a switch port that is tagged to vlan "A" and untagged to vlan "B". LMA authenticates clients and upon authentication places them in appropriate VLANs.
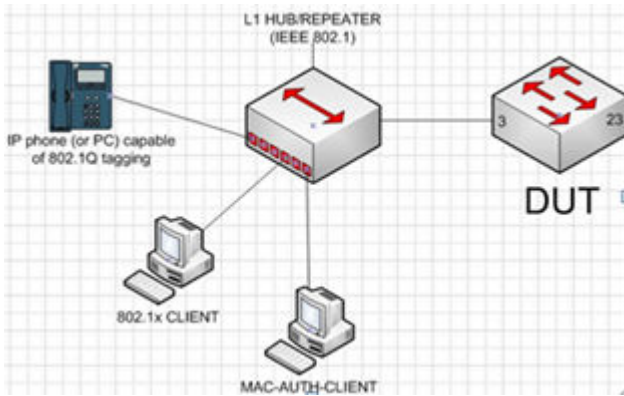
2. In the following scenario, a client is daisy-chained to an IP Phone, which is connected to a tagged vlan port on a switch. The client is authenticating across an untagged vlan. When LMA is enabled on a port and the client connected to it fails to authenticate, the client is assigned attributes configured for the switch's un-auth feature. If the client is authenticated, the switch overrides existing attributes with LMA attributes.
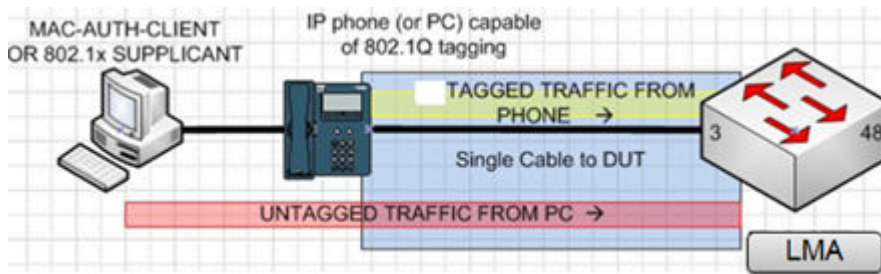
Show commands

LMA supports the following show commands:

- show mac group information

```
HP-Switch# show port-access local-mac mac-group
```

- show default (factory-shipped) mac-group

```
HP-Switch# show port-access local-mac mac-group <mac-group-name>
```

- show profile information

```
HP-Switch# show port-access local-mac profile
```

```
HP-Switch# show port-access local-mac profile provisioned
```

[Note: profiles which are not applied to any mac-group]

- show LMA active clients and applied profiles

```
HP-Switch# show port-access local-mac active
```

- show LMA configuration

```
HP-Switch# show port-access local-mac config

HP-Switch# show port-access local-mac config <port-number> detailed
```
[Note: per port]

- show LMA enabled ports

```
HP-Switch# show port-access local-mac

HP-Switch# show port-access local-mac <port-number>
```
[Note: per port]

- show per port local mac client details

```
HP-Switch# show port-access local-mac <port-number> client [detailed]
```

- show mac-entry and mac-group association

```
HP-Switch# show port-access local-mac association
```

Configuration commands

The configuration strategy below shows the configuration commands that LMA supports. All LMA commands can be prefixed with [no]. For port based commands, a VLAN must be created.

Procedure

1. Enable local mac authentication on switch port '1'

```
HP-Switch(config)#aaa port-access local-mac 1
```

2. Do one of the following:

- Create mac-group, 'ip-phone-grp' for IP phones. The newly created group becomes editable. So, the user can add/delete mac-oui from the mac-group

```
HP-Switch(config)#aaa port-access local-mac mac-group ip-phone-grp
```

- Create mac-group, 'hpphone-grp', from the default (factory-shipped) 'hp-ip-phones' group

```
HP-Switch(config)#aaa port-access local-mac mac-group default hp-ip-phones hpphones-grp
```

Note: To determine the factory-shipped default mac-groups, use `show port-access local-mac mac-group default`

3. Associate mac-address, 005557-9B688B to a mac-group, hpphone-grp

```
HP-Switch(config)#aaa port-access local-mac mac-group hpphones-grp mac-addr005557-9B688B
```

4. Create LMA profile, ip-phone-prof, with attributes, tagged vlan, 2, untagged vlan, 3 and cos 2

```
HP-Switch(config)#aaa port-access local-mac profile ip-phone-prof vlan tagged 2 untagged 3 CoS2
```

5. Associate LMA profile, ip-phone-prof, to a mac-group, hpphone-grp

```
HP-Switch(config)#aaa port-access local-mac apply profile ip-phone-prof mac-group hpphone-grp
```

Per-port attributes

LMA per-port attributes are used to apply attributes for the clients authenticated through LMA profiles. Switches support different per-port values for different authentication methods (802.1X, mac-based and web-based) configured on the same port.

• Configure unauthenticated period

```
HP-Switch(config)#aaa port-access local-mac 1 unauth-period 300
```

• Configure quiet period

```
HP-Switch(config)#aaa port-access local-mac 1 quiet-period 70
```

Configure logoff period

```
HP-Switch(config)#aaa port-access local-mac 1 logoff-period 400
```

Configure AuthVid

```
HP-Switch(config)#aaa port-access local-mac 1 auth-vid 10
```

• Configure UnauthVid

```
HP-Switch(config)#aaa port-access local-mac 1 unauth-vid 12
```

• Configure address limit on a port

```
HP-Switch(config)#aaa port-access local-mac 1  addr-limit 2
```

• Re-authenticate clients on a port

```
HP-Switch(config)#aaa port-access local-mac 1 reauthenticate
```

• Unconfigure LMA on a port

```
HP-Switch(config)#no aaa port-access local-mac 1
```

Configuration examples

Configuration example 1

• In this example, a PC is directly connected to an HPE switch. In addition:

- The corporate PC MAC is 002622bba7ac, and it should end up in VLAN 2 (Notebook of network administrator)

- The rest of the corporate PC series MAC is 00:26:22:bb:* and 00:26:22:bc:*, and it should end up in VLAN 3

- Corporate IP Phones example MAC is 00:80:11:*, and it should end up in VLAN 5 tagged

Configuration example 2

- In this example, PCs are connected to a meeting room HPE 2615 switch series, which is connected to a HPE 3800 switch series (Local MAC authentication happens here). In addition:

  - Authentication of the 2615, example MAC is 00:10:80:* and it should end up in VLAN 15 tagged (management traffic)

  - Corporate PC MAC is: 002622bba7ac, and it should end up in VLAN 2 (Notebook of network administrator)

  - Rest of the corporate PC Series MAC is: 002622bb* and 00:26:22:bc:*, and it should end up in VLAN 3

  - Guest PCs: unknown MAC, and it should end up in Guest VLAN 99

  - Corporate IP Phones, example MAC: 00:80:11:*, and it should end up in VLAN 5 tagged

  - WLAN APs, example MAC: 00:80:12:*, and it should end up in VLAN 10 untagged, 12-14 tagged (10 management, 12-14 SSIDs with local break-out)

For further authentication of any OUIs, predefined in SwitchOS, group default is not allowed.

1. Create 5 LMA profilesThere is no need to create profiles for Guest PCs as you don't know the MACs. Configure unauth-vid (explained in step 3 below) so that such a client fails the authentication and is put into guest VLAN.

```
aaa port-access local-mac profile "corp-switch-prof" vlan tagged 15
```

(for 2615 switches)

```
aaa port-access local-mac profile "corp-pc-prof" vlan untagged 2
```

(for corporate PCs)

```
aaa port-access local-mac profile "rest-pc-prof" vlan untagged 3
```

(for the rest of corporate PCs)

```
aaa port-access local-mac profile "corp-phone-prof" vlan tagged 5
```

 (for corporate ip phones)

```
aaa port-access local-mac profile "wlan-ap-prof" vlan untagged 10 tagged 12-14
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

(for WLAN APs)

2. Associate MACs to these profiles

```
aaa port-ac local-mac apply profile corp-switch-prof mac-oui 001080

aaa port-ac local-mac apply profile corp-pc-prof mac-addr   002622bba7ac

aaa port-ac local-mac apply profile rest-pc-prof mac-mask 002622bb/32  mac-mask  002622bc/32

aaa port-ac local-mac apply profile corp-phone-prof mac-oui  008011

aaa port-ac local-mac apply profile "wlan-ap-prof" mac-oui 008012
```

3. Configure guest VLAN

```
aaa port-ac local-mac <ports> unauth-vid 99
```

4. Enable LMA on ports

```
aaa port-ac local-mac <ports>
```

Configuration using mac-groups

1. Create 3 LMA profiles

```
aaa port-access local-mac profile "corp-pc-prof" vlan untagged 2
```
(for corporate PCs)

```
aaa port-access local-mac profile "rest-pc-prof" vlan untagged 3
```
(for the rest of PCs)

```
aaa port-access local-mac profile "corp-phone-prof" vlan tagged 5
```
(for phones)

2. Create 3 different mac-groups

```
aaa port-ac local-mac mac-group "corp-pc-grp" mac-addr  002622bba7ac
```
(for corporate PCs)

```
aaa port-ac local-mac mac-group "rest-pc-grp" mac-mask  002622bb/32  002622bc/32
```
(for the rest of PCs)

```
aaa port-ac local-mac mac-group "corp-phone-grp" mac-oui  008011
```
(for phones)

3. Associate groups to profiles

```
aaa port-ac local-mac apply profile corp-pc-prof mac-group  corp-pc-grp

aaa port-ac local-mac apply profile rest-pc-prof mac-group  rest-pc-grp

aaa port-ac local-mac apply profile corp-phone-prof mac-group  corp-phone-grp
```

4. Enable LMA on ports

```
aaa port-ac local-mac-auth <ports>
```

Configuration without using mac-groups

1. Create 3 LMA profiles

```
aaa port-access local-mac profile "corp-pc-prof" vlan untagged 2
```
(for corporate PCs)

```
aaa port-access local-mac profile "rest-pc-prof" vlan untagged 3
```
(for the rest of PCs)

```
aaa port-access local-mac profile "corp-phone-prof" vlan tagged 5
```
(for phones)

2. Associate hosts directly to profiles

```
aaa port-ac local-mac apply profile corp-pc-prof   mac-addr   002622bba7ac

aaa port-ac local-mac apply profile rest-pc-prof   mac-mask 002622bb/32

aaa port-ac local-mac apply profile rest-pc-prof  mac-mask  002622bc/32

aaa port-ac local-mac apply profile corp-phone-prof   mac-oui  008011
```

3. Enable LMA on ports

```
aaa port-ac local-mac-auth <ports>
```

Overview

Port-based MAC authentication allows an infrastructure device to be authenticated with a port-based policy that dictates the distribution switch to open the authenticator port to all clients from the authenticated device. This is similar to the existing port-based 802.1X authentication available on HPE switches, except that the port-based 802.1X authentication can also be statically configured on an authenticator port to be persistent over port toggling and switch reboot, while the port-based mode MAC authentication will be dynamic, triggered by the dynamic policy an authenticated client will receive.

Operating notes

1. All existing devices being authenticated with client-based MAC authentication can still be authenticated as usual, until a device with port-based policy is authenticated.

2. A client being authenticated through MAC authentication with a port-based policy will have the connected authenticator port open to all clients without the need of separate authentication.

3. The enhanced MAC authentication will still be able to coexist with web authentication and 802.1X authentication in switch configuration. However it will introduce new dynamic behavior, as defined below:

    a. When a client is authenticated by MAC Authentication with port-based policy.

        I.   If 802.1X authentication is in the client-based mode, all 802.1X clients are removed. With MAC Authentication, only the MAC authenticated client will be accepted and the port will be open.

        II. If 802.1X is in port-based mode, the port is open to all. The ACCESS-ACCEPT (server) for the MAC Authentication client will be rejected with an expectation that the following packets from the client can go through because the port has been opened by 802.1X.

    b. When a client is authenticated as an 802.1X client and the authenticator port has been open by a MAC Authentication client with port-based policy the ACCESS_ACCEPT for the 802.1X client will be rejected appropriately.

    c. When Web authentication and MAC authentication coexist on a switch port, the handling of Web Authentication client will follow the cases of 802.1X client-based mode, for example, 3.a.1. and 3.b.

4. The MAC Authentication Client/Address Limits will be adhered to. If the current number of authenticated clients are at the configured port client/address limit, then additional clients are not accepted (exceptions are clients with client-limit port based policy to increase client-limit). This existing design decision is meant to avoid any DoS of our NAC/AAA solution as any additional authentication requests on this port are pointless as the limit has been hit.

5. A new read-only MIB object will be added to query the authentication mode of a MAC authentication client.

6. When a port is in port-based MAC authentication mode, port-security (port-access mode) blocks traffic from the subsequent users.

7. If a client being MAC authenticated with port-based policy has other assigned policies, (such as NAS-Filter-Rule, Filter-ID, Bandwidth-Max-Ingress/Egress) the port-based policy is then applied only to this client not for any subsequent clients accessing the port in port-mode.

aaa port-access use-lldp-data

Syntax

```
[no] aaa port-access use-lldp-data
```

Description

Authenticate clients using information from their LLDP packets

Usage

[no] aaa port-access use-lldp-data

show port-access config

```
HP-Switch(config)# show port-access config
 Port Access Status Summary
  Port-access authenticator activated [No] : No
  Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No
  Use LLDP data to authenticate [No] : Yes
```

MAC Access Control Lists (ACL)s are an extension of the ACLs feature which include IPv4 Standard, IPv4 Extended ACLs, and IPv6 ACLs. The MAC classes is an extension of Classifier policy feature which includes QoS and Mirror policies.

Classifier Policies and ACLs specify packet attributes on which to match and then take action upon those packets. In the case of ACLs, the actions are permit, deny and log. In the case of Classifier Policies, the actions are specific to the policy type (QoS or Mirror).

The current implementation of ACLs limits packet matching to fields within the IP header of the packet (source IP address, destination IP address, protocol, etc.). MAC ACLs will allow for matching within the Ethernet header of a packet, including source MAC address, destination MAC address and EtherType protocol. MAC ACLs will also allow access to the 802.1q Ethernet frame header values which include the CoS and the VLAN ID. The IP ACLs apply only to Ethernet packets that are of type IP but MAC ACLs will apply to all traffic.

Overview

The MAC ACL and MAC Classes are part of the ACL and Classifier subsystem and they each provide different functionality. Each of the features will be discussed independently to provide the most clarity.

The MAC ACL feature provides a mechanism for the user to permit or deny traffic based on Ethernet frame information. The feature allows for matching traffic based on source MAC address, destination MAC address, Ethernet type, CoS, or VLAN ID. Customers can use this feature to permit or deny specific MAC addresses, block certain types of traffic (for example, appletalk), or block certain CoS/priority packets. The feature extends ACL capabilities down to the Ethernet header and allows matching on most of the fields within the header. This feature's CLI will work very similar to the way IP ACLs are configured but it will need a different context for configuring the match or ignore rules. The context will only allow permit or deny statements with the MAC header fields specified.

The MAC classes feature provides a mechanism for the user to perform actions (for example, remark) on traffic that matches the specified Ethernet header information in the class. The user can create a class that matches the Ethernet header fields: source MAC address, destination MAC address, Ethernet type, VLAN ID or VLAN CoSvalue. After the class is configured the class can be added into a policy and be associated with an action. MAC classes can be included in QoS and Mirror policies and can be applied to those features interfaces (for example, port or VLAN). MAC classes and IPv4/IPv6 classes are mutually exclusive within a policy. A policy that contains both MAC classes and IPv4/IPv6 classes will not be allowed to be configured. Once the policy is applied to an interface any matching traffic will have the specified action applied. This CLI will work very similar to the way classes are defined for IP based traffic.

MAC ACL configuration commands

Mac-access-list creation syntax

This is a new command that needs to be created to allow for the configuration of MAC-based access control lists.

Standard MAC ACL Configuration

Syntax

`mac-access-list standard`

Configure a standard MAC Access Control List.

NAME-STR

The standard MAC ACL name.

200-299

The standard MAC ACL number.

```
mac-access-list standard <200>
```

Description

Configure the standard MAC ACL to filter the packets based on the source MAC address. The standard MAC ACL number ranges from 200 to 299.

Example

```
(config)#mac accss-list standard 200
(config-std-macl)#
```

Extended MAC ACL Configuration

Syntax

```
mac-access-list extended
```

Configure an extended MAC Access Control List.

NAME-STR

The extended MAC ACL name.

300-399

The extended MAC ACL number.

```
mac-access-list extended <300>
```

Description

Configure the extended MAC ACL to filter the packets based on the source MAC address, destination MAC address, ethertype, CoS priority, or VLAN number. The extended MAC ACL number ranges from 300 to 399.

Example

```
(config)#mac accss-list extended 300
(config-ext-macl)#
```

Resequencing MAC ACL

Syntax

```
mac-access-list resequence
```

Renumber the sequence number of the rules in the MAC ACL specified.

<1-2147483647>

The sequence number assigned to the first rule of the specified MAC ACL.

<1-2147483646>

    The increment value that renumbers the subsequent rules in the specified MAC ACL.

```
mac-access-list resequence 200 1 10
```

Description

Re-number the sequence number of the rules in the MAC ACL specified. The first rule receives the sequence number specified in the start-seq-num and the subsequent rule numbers increment per the increment value.

Example

```
(config)# mac-access-list resequence 300 1 10
```

Similar Command

```
ip access-list
```

Mac-access-list standard configuration context

This command is used to configure MAC ACL with a simplified configuration. A simplified configuration provides a way to easily configure MAC ACLs that only require matching on a source MAC address.

Syntax

```
[no] SEQ-NUM < permit | deny > < any | host > SRC-MAC | SRC-MAC-MASK log
```

permit

    Packets matching the specified Ethernet header information.

deny

    Packets matching the specified Ethernet header information.

any

    Match the packets with any source MAC address.

host

    Match the packets with the specified source MAC address.

SRC-MAC

    Match the packets belonging to the specified source MAC address range.

SRC-MAC-MASK

    The MAC address group mask.

log

    Log a debug message when the MAC ACL rule is hit.

Similar Command

```
(config)#ip access-list standard 1
```

Configure standard MAC ACL

```
(config)# mac-access-list standard 200
(config-std-macl)# permit AABB.CCDD.EEFF 0000.0000.FFFF
(config-std-macl)# deny host AABB.CCDD.EEFF log
```

Syntax

[no] *SEQ-NUM* remark

Add a comment for the MAC ACL rule specified. The maximum comment length is 100 characters.

Mac-access-list extended configuration context

Syntax

[no] *SEQ-NUM* < permit | deny > < any | host > *SRC-MAC* | *SRC-MAC-MASK* < any | host > *DST-MAC* | *DST-MAC-MASK* < any | *ETHERTYPE* cos *COS log*

```
[no] [<SEQ-NUM>] permit {any| host <SRC-MAC>|<SRC-MAC> <SRC-MAC-MASK>} {any|host
<DST-MAC> | <DST-MAC> <DST-MAC-MASK>} {any|ETHERTYPE} [cos <priority>|vlan <vlan-
id>] [log]
```

```
[no] [<SEQ-NUM>] deny {any| host <SRC-MAC>|<SRC-MAC> <SRC-MAC-MASK>} {any|host
<DST-MAC> | <DST-MAC> <DST-MAC-MASK>} {any|ETHERTYPE} [cos <priority>|vlan <vlan-
id>] [log]
```

Used to configure an extended MAC ACL. The extended capabilities allow for matching on source MAC address, destination Mac address, EtherType, CoS, and VLAN. The VLAN value is only applicable when the MAC ACL is applied to a port or trunk interface.

permit

Packets matching the specified Ethernet Header information.

deny

Packets matching the specified Ethernet Header information.

any

Match packets with any source/destination MAC address.

host

Match packets with the specified source/destination MAC address.

SRC-MAC

Match packets belonging to the specified source/destination MAC address range.

SRC-MAC-MASK

The source MAC address group mask.

DST-MAC-MASK

The destination MAC address group mask.

<0x600-0xFFFF>

Match a specific EtherType protocol.

aarp

AppleTalk Address Resolution Protocol (AARP)

appletalk

AppleTalk/EtherTalk

arp

Address Resolution Protocol (ARP)

fcoe

Fibre Channel over Ethernet

fcoe-init

Fibre Channel over Ethernet Initialization

lldp

Link Layer Discovery Protocol

ip

Internet Protocol Version 4

ipv6

Internet Protocol Version 6

ipx-arpa

IPX Advanced Research Projects Agency (ARPA)

ipx-non-arpa

IPX non-ARPA

is-is

Intermediate System to Intermediate System

mpls-unicast

MPLS Unicast

mpls-multicast

MPLS Multicast

q-in-q

IEEE 802.1ad encapsulation

rbridge

RBridge Channel Protocol

trill

IETF TRILL protocol

wake-on-lan

Wake on LAN

log

Log a debug message when the MAC ACL rule is hit.

cos

Match packets with a specified 802.1Q Priority Code Point value.

vlan

Match packets with the specified VLAN value.

VLAN-ID

Match packets with the specified VLAN value.

<0-7>

Match packets with a specified 802.1Q Priority Code Point value.

Similar Command

```
(config)#ip access-list extended 100
```

Remark command

The remark command allows for the insertion of a string at the specified sequence number. The remark will consume the sequence number where it is specified and will remain in proper order if the list is resequenced. The remark ability provides a way of tracking notes inside the given ACL but they do not affect the behavior of the ACL.

Syntax

```
[no] SEQ-NUM remark
```

Add a comment for the MAC ACL or MAC ACL rule specified. The maximum comment length is 100 characters.

Mac-access-list application syntax (PACL)

This command is used to apply a MAC ACL to an interface.

Syntax

(Within interface context.)

```
mac-access-group ACL-ID in|out
```

Apply a MAC ACL to traffic on a port. A standard or extended MAC ACL filters packets based on the source MAC address, destination MAC address, ethertype, CoS, or VLAN.

ASCII-STR

The MAC ACL name.

in

Apply MAC ACL on the inbound packets.

Similar command

```
ip access-group name in
```

```
mac-access-group name in
```

Mac-access-list application syntax (VACL)

This command is used to apply a MAC ACL to a VLAN .

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Syntax

(Within interface context.)

```
mac-access-group ACL-ID in|out
```

Apply a MAC ACL to traffic on a VLAN. A standard or extended MAC ACL filters packets based on the source MAC address, destination MAC address, ethertype, CoS, or VLAN.

ASCII-STR

The MAC ACL name.

in

Apply MAC ACL on the inbound packets.

> **NOTE**
>
> **Similar command**
>
> ```
> ip access-group name in|out
> ```
>
> **See: example**

Applying a MAC ACL to VLAN 1

```
(config)#vlan 1
(vlan-1)# mac-access-group name in
```

Show access-list

Syntax

```
show access-list ACL-NAME-STR
 | config | ports | radius | resources | tunnel
<TUNNEL-ID> | vlan <VLAN-ID>
```

Show access control list information. If `no` parameters are specified, a table of ACL information is displayed.

ACL-NAME-STR

Display detailed information about the specified ACL.

config

Show all configured ACLs on the switch using the CLI syntax used to create them.

ports

Show ACLs applied to the specified ports.

radius

Display ACLs applied via RADIUS.

resources

Display ACL resource usage and availability.

tunnel

Show ACLs applied to the specified tunnel.

vlan

   Show ACLs applied to the specified VLAN.

Show access-list by name

This command is used to display the details about a specific ACL.

Syntax

```
show access-list <ACL-ID> config
```

Show access-list 300

```
switch(config)# show access-list 300
Access Control Lists
Name: 300
Type: MAC Extended
Applied: No
SEQ: Entry
-----------------------------------------
10 Action : permit
Src MAC: 1111.2222.3333  Mask: ffff.ffff.0000
Dst MAC: 4444.5555.6666  Mask: ffff.ffff.0000
Ethertype: aarp  CoS: 7  VLAN ID: 1
```

Show access-list 200

```
switch(config)# show access-list 200
Access Control Lists
Name: 200
Type: MAC Standard
Applied: No
SEQ:  Entry
-----------------------------------------------------
10 Action:  permit
Src MAC:    1111.2222.3333 Mask: ffff.ffff.0000
Ethertype : any
```

Show access-list 100

```
switch(config)# show access-list 100
Name: 100
Type: IPv4 Extended
Applied: No
SEQ:  Entry
----------------------------------------------------
10 Action:  deny
   Src IP:    0.0.0.0            Mask: 255.255.255.255   Port(s):
   Dst IP:    0.0.0.0            Mask: 255.255.255.255   Port(s):
   Proto :    TCP
   TOS   :    Precedence:
20 Action:  deny
   Src IP:    0.0.0.0            Mask: 255.255.255.255   Port(s):
   Dst IP:    0.0.0.0            Mask: 255.255.255.255   Port(s):
   Proto :    UDP
   TOS   :    Precedence: -
```

Show access-list v6ACL

```
switch(config)# show access-list v6ACL
Name: 100
Type: IPv6
Applied: No
SEQ  Entry
-----------------------------------------
10  Action:       deny
    Src IP:       Prefix Len: 0
    Dst IP:       Prefix Len: 0
    Src Port(s):  Dst Port(s):
    Proto :       TCP  Option(s):
    Dscp :
```

Show access-list config

Syntax

```
show access-list <ACL-ID> config
```

Used to display a specific ACL as it would be shown in configuration.

mac-access-list

```
(config)# mac-access-list 300 config
10 permit 1111.2222.3333 ffff.ffff.0000 4444.5555.6666 ffff.ffff.0000 aarp
exit

(config)# mac-access-list 200 config
10 permit 1111.2222.3333 4444.5555.6666
exit
```

Show access-list port

Syntax

```
show access-list port <port-list>
```

Used to display the current ACLs that are applied to a specified port.

Show access-list

```
(config)# show access-list port f1
Access Lists for Port F1
IPv4 Inbound : 100   Type: Extended
MAC  Inbound : 300   Type: Extended
```

Show access-list vlan

Syntax

```
show access-list vlan < VLAN-ID | all >
```

Used to display the current ACLs that are applied to a specified VLAN.

VLAN-ID

Show ACLs applied to the specified VLAN.

all

Show ACLs applied to all VLANs.

Show access-list

```
(config)# show access-list vlan 1
Access Lists for VLAN 1
IPv4 Router Inbound           : (None)
IPv4 VLAN Inbound             : (None)
IPv4 Connection Rate Filter   : (None)
IPv6 Router Inbound           : (None)
IPv6 VLAN Inbound             : (None)
MAC  VLAN Inbound             : 300    Type: Extended
```

Show access-list resources

Syntax

```
show access-list resource
```

Used to display current resource usage and availability in the policy enforcement engine.

Show access-list resource

```
Switch# show access-list resource

 Resource usage in Policy Enforcement Engine

 Ingress Policy Enforcement Engine Rules
 Resource usage in Policy Enforcement Engine

            |   Rules   |  Rules Used
      Ports | Available | ACL | QoS |  VT | Mirr |  OF  | Other |
 -------------+----------+-----+-----+-----+------+------+-------|
 1-28        |    4080 |  0 |  0 |  0 |   0 |   0 |    0 |


 Ingress Policy Enforcement Engine Meters

            |   Meters  |  Meters Used
      Ports | Available | ACL | QoS |  VT | Mirr |  OF  | Other |
 -------------+----------+-----+-----+-----+------+------+-------|
 1-28        |    2047 |    |  0 |  0 |     |   0 |    0 |


 Ingress Policy Enforcement Engine Port Ranges

            |Application|
            |Port Ranges|  Application Port Ranges Used
      Ports | Available | ACL | QoS |  VT | Mirr |  OF  | Other |
 -------------+----------+-----+-----+-----+------+------+-------|
 1-28        |      60 |  0 |  0 |     |   0 |   0 |    0 |

 2 of 32 Policy Engine management resources used.


 Egress Policy Enforcement Engine Rules
 Resource usage in Policy Enforcement Engine
```

```
                       |   Rules   |   Rules Used
            Ports | Available | ACL | QoS |  VT | Mirr |  OF  | Other |
       --------------+-----------+-----+-----+-----+------+------+-------|
       1-28          |      2032 |   0 |   0 |   0 |    0 |    0 |     0 |


Egress Policy Enforcement Engine Meters

                       |   Meters   |  Meters Used
            Ports | Available | ACL | QoS |  VT | Mirr |  OF  | Other |
       --------------+-----------+-----+-----+-----+------+------+-------|
       1-28          |      1023 |     |   0 |   0 |      |    0 |     0 |


Egress Policy Enforcement Engine Port Ranges

                   |Application|
                   |Port Ranges|  Application Port Ranges Used
            Ports | Available | ACL | QoS |  VT | Mirr |  OF  | Other |
       --------------+-----------+-----+-----+-----+------+------+-------|
       1-28          |        60 |   0 |   0 |     |    0 |    0 |     0 |

0 of 8 Policy Engine management resources used.
```

The hardware (TCAM) resources used by the ACLs configured on the switch are as follows:

Key	
ACL	Access Control Lists
QoS	Quality of Service
VT	Virus Throttling
Mirror	Mirror Policies, Remote Intelligent Mirror endpoints
Other	Management VLAN, DHCP Snooping, ARP Protection, Jumbo IP-MTU, Transparent Mode.

Resource usage includes resources actually in use, or reserved for future use by the listed feature. Internal dedicated-purpose resources, such as port bandwidth limits or VLAN QoS priority, are not included.

Show statistics

The show statistics command will need to be updated to take a MAC parameter.

Syntax

```
show statistics mac <ACL-NAME-STR> port <PORT-NUM>
```

Used to display hit counts for a given MAC ACL.

mac

   Display the statistics of MAC ACL.

ACL-NAME-STR

   The MAC ACL name.

port

Show statistics for the specified port.

[ethernet] PORT-NUM

The port on which the MAC ACL is applied.

Syntax

```
show statistics mac <ACL-NAME-STR> vlan <VLAN-ID> in|out|vlan
```

vlan

Show statistics for the specified VLAN.

VLAN-ID

The VLAN ID or VLAN name.

in

Show statistics for MAC ACLs that are applied inbound.

out

Show statistics for MAC ACLs that are applied outbound.

---

show statistics mac

```
show statistics mac 300 port 1 in

show statistics mac 300 vlan 10 in

show statistics mac 300 vlan 10 vlan
```

---

show statistics mac superMac vlan 10 in

```
show statistics mac superMac vlan 10 in

HitCounts for ACL superMac
Total
( 540 )     10 permit any 1111.2222.3333 4444.5555.6666
```

clear statistics

The clear statistics command will need to be updated to take a MAC parameter.

Syntax

```
clear statistics mac <ACL-NAME-STR> port <PORT-NUM>
```

Clear all the counters for the ACLs that match the criteria specified.

mac

Clear the statistics for MAC ACL.

ACL-NAME-STR

The MAC ACL name or the MAC ACL number.

port

Clear statistics for the specified port.

[ethernet] PORT-NUM

The port from which the MAC ACL statistics is cleared.

Syntax

```
clear statistics mac <ACL-NAME-STR> port <PORT-NUM> | VLAN <VLAN-ID> in|out|vlan
```

VLAN

Clear statistics for the specified VLAN.

VLAN-ID

The VLAN ID or VLAN name.

in

Clear statistics for inbound packets on the VLAN.

out

Clear statistics for outbound packets on the VLAN.

Clear statistics mac superMac

```
clear statistics mac superMac vlan 10 in
```

CLI command error messages

Event	Message
Attempt to move a port into a trunk with a MAC ACL applied and there are not enough hardware resources available to accommodate the new port.	Operation failed on Port %s: General error.
Attempt to move a port with a MAC ACL applied to it into a trunk.	Unable to apply access control list. Port %s is part of a trunk.
Apply a MAC ACL to a vlan, interface or tunnel and there are not enough hardware resources available.	Unable to apply access control list. Failed to add entry %s.
Add a class to a policy and a non-matching class type is already applied IP/MAC)	Unable to add class. Policy cannot contain both IP and MAC classes.
Apply a MAC ACL on a vlan and it contains a rule that matches vlan-ID	Unable to apply access control list. ACLs applied to VLANs cannot match the VLAN ID parameter.
Apply a MAC Policy on a vlan and it contains a rule that matches vlan-ID	Unable to apply policy. Policies applied to VLANs cannot contain a class that matches the VLAN ID parameter.

Overview

TACACS AAA systems are used as a single point of management to configuring and store user accounts. They are often coupled with directories and management repositories, simplifying the set up maintenance of the end-user accounts.

In the authorization function of the AAA system, network devices with Authentication Services can provide fine-grained control over user capabilities for the duration of the user's session, for example setting access control or session duration. Enforcement of restrictions to a user account can limit available commands and levels of access.

TACACS+ authentication enables you to use a central server to allow or deny access to switches and other TACACS-aware devices in your network. This allows use of a central database to create multiple unique username/password sets with associated privilege levels, for use by individuals accessing the switch from either its console port (local access) or Telnet (remote access).

Figure 24: *Example of TACACS+ operation*
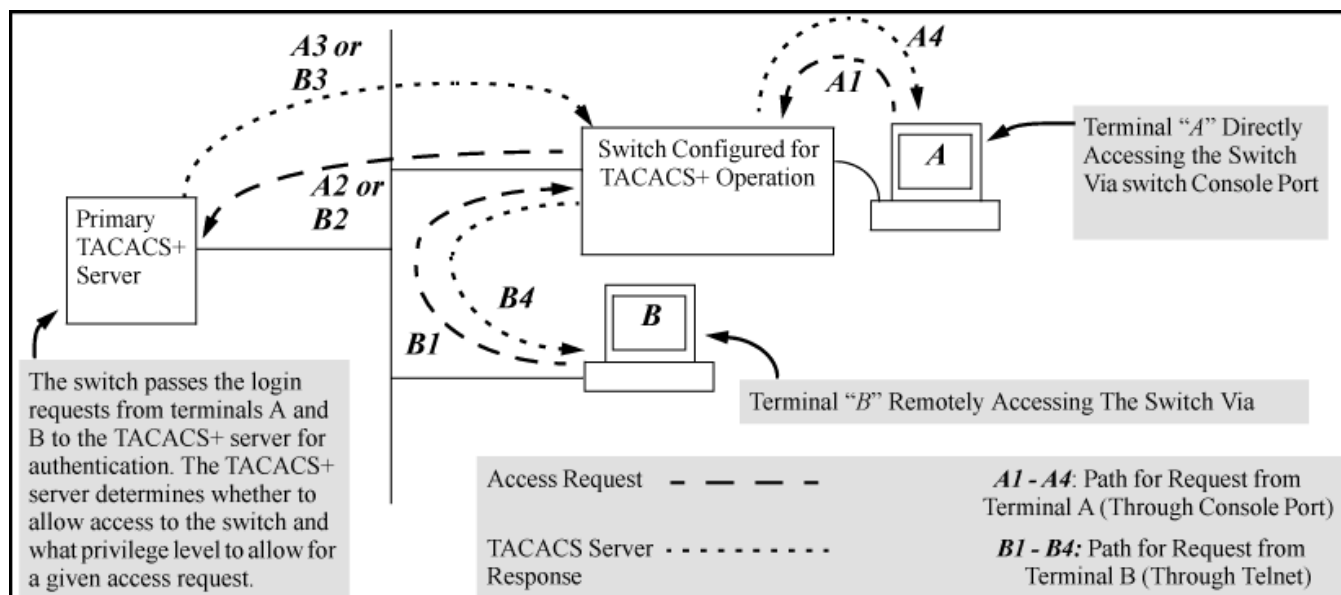
TACACS+ in the switches covered in this guide manages authentication of logon attempts through either the Console port or Telnet. TACACS+ uses an authentication hierarchy consisting of (1) remote passwords assigned in a TACACS+ server and (2) local passwords configured on the switch. That is, with TACACS+ configured, the switch first tries to contact a designated TACACS+ server for authentication services. If the switch fails to connect to any TACACS+ server, it defaults to its own locally assigned passwords for authentication control (if configured to do so). For both Console and Telnet access you can configure a login (read-only) and an enable (read/write) privilege level access.

TACACS+ does not affect WebAgent access. See **Controlling WebAgent access when using TACACS+ authentication** on page 192

General system requirements

To use TACACS+ authentication, you need the following:

- A TACACS+ server application installed and configured on one or more servers or management stations in your network. There are several TACACS+ software packages available.

- A switch configured for TACACS+ authentication, with access to one or more TACACS+ servers.

NOTE

The effectiveness of TACACS+ security depends on correctly using your TACACS+ server application. For this reason, Hewlett Packard Enterprise recommends that you thoroughly test all TACACS+ configurations used in your network.

TACACS-aware switches include the capability of configuring multiple backup TACACS+ servers. Hewlett Packard Enterprise recommends that you use a TACACS+ server application that supports a redundant backup installation. This allows you to configure the switch to use a backup TACACS+ server if it loses access to the first-choice TACACS+ server.

TACACS+ does not affect WebAgent access. See **Controlling WebAgent access when using TACACS+ authentication** on page 192

General authentication setup procedure

It is important to test the TACACS+ service before fully implementing it. Depending on the process and parameter settings you use to set up and test TACACS+ authentication in your network, you could accidentally lock all users, including yourself, out of access to a switch. While recovery is simple, it can pose an inconvenience that can be avoided. To prevent an unintentional lockout on the switch, use a procedure that configures and tests TACACS+ protection for one access type (for example, Telnet access), while keeping the other access type (console, in this case) open in case the Telnet access fails due to a configuration problem. The following procedure outlines a general setup procedure.

NOTE

If a complete access lockout occurs on the switch as a result of a TACACS+ configuration, see "Troubleshooting TACACS+ operation" in the management and configuration guide for your switch.

Procedure

1. Familiarize yourself with the requirements for configuring your TACACS+ server application to respond to requests from the switch. (See the documentation provided with the TACACS+ server software.) This includes knowing whether you need to configure an encryption key, see **Using the encryption key** on page 191.

2. Determine the following:

   - The IP address(es) of the TACACS+ server(s) you want the switch to use for authentication. If you will use more than one server, determine which server is your first-choice for authentication services.

   - The encryption key, if any, for allowing the switch to communicate with the server. You can use either a global key or a server-specific key, depending on the encryption configuration in the TACACS+ server(s).

   - The number of log-in attempts you will allow before closing a log-in session. (Default: 3)

   - The period you want the switch to wait for a reply to an authentication request before trying another server.

   - The username/password pairs you want the TACACS+ server to use for controlling access to the switch.

   - The privilege level you want for each username/password pair administered by the TACACS+ server for controlling access to the switch.

   - The username/password pairs you want to use for local authentication (one pair each for operator and manager levels).

---

3. Plan and enter the TACACS+ server configuration needed to support TACACS+ operation for Telnet access (login and enable) to the switch. This includes the username/password sets for logging in at the operator (read-only) privilege level and the sets for logging in at the manager (read/write) privilege level.

> **NOTE**
> When a TACACS+ server authenticates an access request from a switch, it includes a privilege level code for the switch to use in determining which privilege level to grant to the terminal requesting access. The switch interprets a privilege level code of "15" as authorization for the manager (read/write) privilege level access. Privilege level codes of 14 and lower result in operator (read-only) access. Thus, when configuring the TACACS+ server response to a request that includes a username/password pair that should have manager privileges, you must use a privilege level of 15. For more on this topic, see the documentation you received with your TACACS+ server application.

If you are a first-time user of the TACACS+ service, Hewlett Packard Enterprise recommends that you configure only the minimum feature set required by the TACACS+ application to provide service in your network environment. After you have success with the minimum feature set, you can then want to try additional features that the application offers.

4. Ensure that the switch has the correct local username and password for manager access. (If the switch cannot find any designated TACACS+ servers, the local manager and operator username/password pairs are always used as the secondary access control method.)

> **CAUTION**
> Ensure that the switch has a local manager password. Otherwise, if authentication through a TACACS+ server fails for any reason, unauthorized access will be available through the console port or Telnet.

5. Using a terminal device connected to the switch console port, configure the switch for TACACS+ authentication only for **Telnet login** access and **Telnet enable** access. At this stage, do not configure TACACS + authentication for console access to the switch, as you may need to use the console for access if the configuration for the Telnet method needs debugging.

6. Ensure that the switch is configured to operate on your network and can communicate with your first-choice TACACS+ server. (At a minimum, this requires IP addressing and a successful `ping` test from the switch to the server.)

7. On a remote terminal device, use Telnet to attempt to access the switch. If the attempt fails, use the console access to check the TACACS+ configuration on the switch. If you make changes in the switch configuration, check Telnet access again. If Telnet access still fails, check the configuration in your TACACS+ server application for misconfigurations or missing data that could affect the server's inter-operation with the switch.

8. After your testing shows that Telnet access using the TACACS+ server is working properly, configure your TACACS+ server application for console access. Then test the console access. If access problems occur, check for and correct any problems in the switch configuration, and then test console access again. If problems persist, check your TACACS+ server application for misconfigurations or missing data that could affect the console access.

9. When you are confident that TACACS+ access through both Telnet and the switch console operates properly, use the `write memory` command to save the switch running-config file to flash.

Configuring TACACS+ on the switch

Before you begin:

If you are new to TACACS+ authentication, Hewlett Packard Enterprise recommends that you first read the **General authentication setup procedure** on page 161 and configure your TACACS+ servers before configuring authentication on the switch.

The switch offers three command areas for TACACS+ operation:

- `show authentication`

  and `show tacacs`: Displays the switch TACACS+ configuration and status.

- `aaa authentication:`

  A command for configuring the switch authentication methods.

- `tacacs-server:`

  A command for configuring the switch contact with TACACS+ servers.

show authentication

Syntax

show authentication

Description

List the number of login attempts the switch allows in a single login session, and the primary or secondary access methods configured for each type of access.

```
show authentication


HP-Switch(config)# show authentication
Status and Counters - Authentication Information
  login Attempts : 3
  Lockout Delay : 0
  Respect Privilege : Disabled
  Bypass Username For Operator and Manager Access : Disabled


                | Login          Login         Login
  Access Task   | Primary        Server Group Secondary
  -------------- + -----------    ------------ -----------
  Console       | Local                        None
  Telnet        | Local                        None
  Port-Access   | Local                        None
  Webui         | Local                        None
  SSH           | Two-factor                   None
  Web-Auth      | ChapRadius      radius        None
  MAC-Auth      | ChapRadius      radius        None
  SNMP          | Local                        None
  Local-MAC-Auth | Local                       None

  Enable        | Enable         Enable
  Access Task   | Primary        Server Group Secondary
  -------------- + -----------    ------------ -----------
  Console       | Local                        None
  Telnet        | Local                        None
  Webui         | Local                        None
  SSH           | Two-factor      tacacs        None
```

```
show authentication two-factor


HP-Switch(config)# show authentication two-factor
                | Login          Login
  Access Task   | First          Second
  -------------- + -----------    ------------
```

```
    SSH           | public-key    local

                  | Enable        Enable
    Access Task   | First         Second
    ------------- + ----------    ------------
    SSH           | public-key    local

                  | Login         Login
    Access Task   | First         Second
    ------------- + ----------    ------------
    SSH           | certificate   local

                  | Enable        Enable
    Access Task   | First         Second
    ------------- + ----------    ------------
    SSH           | certificate   local
```

show authentication for RADIUS server dead time

```
switch# show authentication

Status and Counters - Authentication Information
Authorized enabled as backup for secondary login are preceded by *

  Login Attempts : 3
  Lockout Delay : 0
  Respect Privilege : Disabled
  Bypass Username For Operator and Manager Access : Disabled

                  | Login         Login         Login
    Access Task   | Primary       Server Group  Secondary
    ------------- + ----------    ------------  ----------
    Console       | Local                        None
    Telnet        | Local                        None
    Port-Access   | Local                        *Cached
    Webui         | Local                        None
    SSH           | Local                        None
    Web-Auth      | ChapRadius    radius         None
    MAC-Auth      | ChapRadius    radius         *Cached
    SNMP          | Local                        None
    Local-MAC-Auth| Local                        None

                  | Enable        Enable        Enable
    Access Task   | Primary       Server Group  Secondary
    ------------- + ----------    ------------  ----------
    Console       | Local                        None
    Telnet        | Local                        None
    Webui         | Local                        None
    SSH           | Local                        None
```

Viewing the current TACACS+ server contact configuration

Syntax:

```
show tacacs
```

This command lists the timeout period, encryption key, and the IP addresses of the first-choice and backup TACACS+ servers the switch can contact.

Example:

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

If the switch is configured for a first-choice and two backup TACACS+ server addresses, the default timeout period, and `paris-1` for a (global) encryption key, `show tacacs` produces a listing similar to the following:

Figure 25: *Example of the switch TACACS+ configuration listing*

```
                    HP Switch(config)# show tacacs
                     Status and Counters - TACACS Information
First-Choice           Timeout : 5
TACACS+ Server         Encryption Key :
                       Server IP Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx
Second-Choice          --------------- ------ ------ ------ ------ ------- -------
TACACS+ Server         10.30.248.100   0      0      0      0      0       0
                       10.30.248.156   0      0      0      0      0       0
Third-Choice           10.30.248.105   0      0      0      0      0       0
TACACS+ Server
```

Configuring the switch authentication methods

The `aaa authentication` command configures access control for the following access methods:

- Console

- Telnet

- SSH

- Web

- Port-access (802.1X)

However, TACACS+ authentication is only used with the console, Telnet, or SSH access methods. The command specifies whether to use a TACACS+ server or the switch local authentication, or (for some secondary scenarios) no authentication. This means that if the primary method fails, authentication is denied. The command also reconfigures the number of access attempts to allow in a session if the first attempt uses an incorrect username/ password pair.

Using the privilege-mode option for login

When using TACACS+ to control user access to the switch, first login with your username at the operator privilege level using the password for operator privileges, then login again with the same username but using the Manger password to obtain manager privileges. You can avoid this double login process by entering the `privilege-mode` option with the `aaa authentication login` command to enable TACACS+ for a single login. The switch authenticates your username/password, then requests the privilege level (operator or manager) that was configured on the TACACS+ server for this username/password. The TACACS+ server returns the allowed privilege level to the switch. You are placed directly into operator or manager mode, depending on your privilege level.

```
HP Switch(config) aaa authentication login privilege-mode
```

The `no` version of the above command disables TACACS+ single login capability.

Selecting the access method for configuration

Syntax:

```
aaa authentication
```

```
<console|telnet|ssh|web|port-access>
```

Selects the access method for configuration.

```
<enable>
```

The server grants privileges at the manager privilege level.

```
<login[privilege-mode]>
```

The server grants privileges at the operator privilege level. If the `privilege-mode` option is entered, TACACS+ is enabled for a single login. The authorized privilege level (operator or manager) is returned to the switch by the TACACS+ server.

Default: `Single login disabled`.

```
<local|tacas|radius>
```

Selects the type of security access:

`local` – Authenticates with the manager and operator password you configure in the switch.

`tacacs` – Authenticates with a password and other data configured on a TACACS+ server.

`radius` – Authenticates with a password and other data configured on a RADIUS server.

```
[<local|none>]
```

If the primary authentication method fails, determines whether to use the local password as a secondary method or to disallow access.

Syntax:

```
aaa authentication num-attempts <1-10>
```

Specifies the maximum number of login attempts allowed in the current session.

Default: `3`

Authentication parameters

Table 11: *AAA authentication parameters*

Name	Default	Range	Function
console, Telnet, SSH, web or port-access	n/a	n/a	Specifies the access method used when authenticating. TACACS+ authentication only uses the console, Telnet or SSH access methods.
enable	n/a	n/a	Specifies the manager (read/write) privilege level for the access method being configured.

Table Continued

Name	Default	Range	Function
login *<privilege-mode>*	privilege-mode disabled	n/a	login: Specifies the operator (read-only) privilege level for the access method being configured. The privilege-mode option enables TACACS+ for a single login. The authorized privilege level (operator or manager) is returned to the switch by the TACACS+ server.
local - **or** -tacacs	local	n/a	Specifies the primary method of authentication for the access method being configured. local: Use the username/password pair configured locally in the switch for the privilege level being configured tacacs: Use a TACACS+ server.

Table Continued

Name	Default	Range	Function
`local` - **or** -`none`	`none`	n/a	Specifies the secondary (backup) type of authentication being configured. `local`: The username/password pair configured locally in the switch for the privilege level being configured `none`: No secondary type of authentication for the specified method/privilege path. (Available only if the primary method of authentication for the access being configured is local.)
			NOTE If you do not specify this parameter in the command line, the switch automatically assigns the secondary method as follows: • If the primary method is `tacacs`, the only secondary method is `local`. • If the primary method is `local`, the default secondary method is `none`.
`num-attempts`	3	1–10	In a given session, specifies how many tries at entering the correct username/password pair are allowed before access is denied and the session terminated.

Configuring TACACS+ server

Syntax

```
[no] tacacs-server host IP-ADDR key KEY-STR
```

Configures a TACACS+ server for Authentication, Authorization and Accounting. A maximum of 3 TACACS+ servers can be configured.

TACACS-server: Configures a TACACS+ server for Authentication, Authorization and Accounting.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Host: Configures the IP address of a TACACS+ server.

Key: Configures the default access key for all TACACS+ servers.

Timeout: Configures the server response timeout.

Configuring the TACACS+ server for single login

For the single login feature to work correctly you must check some entries in the User Setup on the TACACS+ server:

Procedure

1. In the User Setup, scroll to the Advanced TACACS+ Settings section.

2. Make sure the radio button for "Max Privilege for any AAA Client" is checked and the level is set to 15, as shown in **Figure 26: Advanced TACACS+ settings section of the TACACS+ server user setup** on page 169.

3. Privileges are represented by the numbers 0 through 15, with zero allowing only operator privileges (and requiring two logins) and 15 representing root privileges. The root privilege level is the only level that will allow manager level access on the switch.

Figure 26: *Advanced TACACS+ settings section of the TACACS+ server user setup*

---
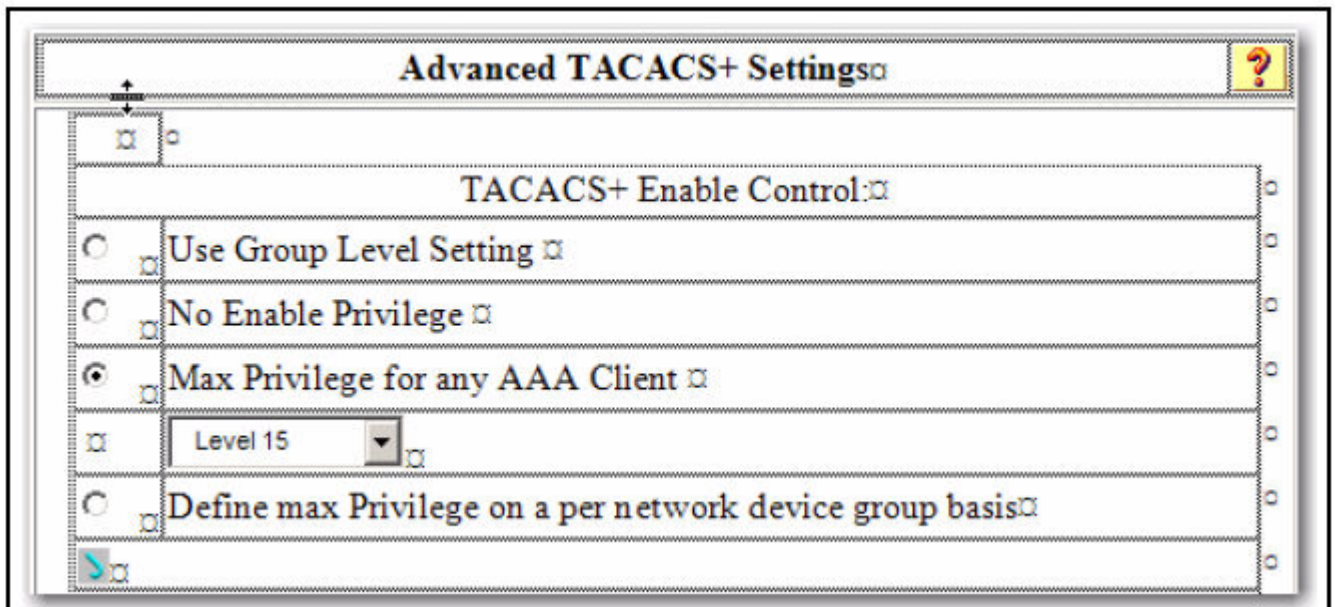
4. Scroll down to the section that begins with "Shell", see **Figure 27: The shell section of the TACACS+ server user setup** on page 170. Check the Shell box.

5. Check the Privilege level box and set the privilege level to 15 to allow "root" privileges. This allows you to use the single login option.

Figure 27: *The shell section of the TACACS+ server user setup*

☑	**Shell (exec)** ¤	
☐	Access control list¤	
☐	Auto command¤	
☐	Callback line¤	
☐	Callback rotary¤	
☐	Idle time¤	
☐	No callback verify¤	☐ Enabled¤
☐	No escape¤	☐ Enabled¤
☐	No hangup¤	☐ Enabled¤
☑	Privilege level¤	15
☐	Timeout¤	123
☐	Custom attributes ¤	

As shown in **Configuring the switch TACACS+ server access** on page 172, login and enable access is always available locally through a direct terminal connection to the switch console port. However, for Telnet access, you can configure TACACS+ to deny access if a TACACS+ server goes down or otherwise becomes unavailable to the switch.

Table 12: *Primary/secondary authentication table*

Access method and privilege level	Authentication options		Effect on access attempts
	Primary	**Secondary**	
Console — Login	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Access method and privilege level	Authentication options		Effect on access attempts
	Primary	Secondary	
Console — Enable	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
Telnet — Login	local	none*	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.
Telnet — Enable	local	none	Local username/password access only.
	tacacs	local	If Tacacs+ server unavailable, uses local username/password access.
	tacacs	none	If Tacacs+ server unavailable, denies access.

⚠ **CAUTION**

Regarding the use of local for login primary access:

During local authentication (which uses passwords configured in the switch instead of in a TACACS+ server), the switch grants read-only access if you enter the operator password, and read-write access if you enter the manager password. For example, if you configure authentication on the switch with Telnet Login Primary as Local and Telnet Enable Primary as Tacacs, when you attempt to Telnet to the switch, you will be prompted for a local password. If you enter the switch local manager password (or, if there is no local manager password configured in the switch) you can bypass the TACACS+ server authentication for Telnet Enable Primary and go directly to read-write (manager) access. Thus, for either the Telnet or console access method, configuring Login Primary for Local authentication while configuring Enable Primary for TACACS+ authentication is not recommended, as it defeats the purpose of using the TACACS+ authentication. If you want Enable Primary log-in attempts to go to a TACACS+ server, then you should configure both Login Primary and Enable Primary for Tacacs authentication instead of configuring Login Primary to Local authentication.

Access options

The following is a set of access options and the corresponding commands to configure them:

- Console login (operator or read-only) access, primary using TACACS+ server and secondary using local

  ```
  HP Switch (config)# aaa authentication console login tacacs local
  ```

- Console enable (manager or read/write) access, primary using TACACS+ server and secondary using local

  ```
  HP Switch (config)# aaa authentication console enable tacacs local
  ```

- Telnet login (operator or read-only) access, primary using TACACS+ server and secondary using local

  ```
  HP Switch (config)# aaa authentication Telnet login tacacs local
  ```

- Telnet enable (manager or read/write) access, primary using TACACS+ server and secondary using local

  ```
  HP Switch (config)# aaa authentication telnet enable tacacs local
  ```

- Deny access and close the session after failure of two consecutive username/password pairs

  ```
  HP Switch (config)# aaa authentication num-attempts 2
  ```

Configuring the switch TACACS+ server access

The `tacacs-server` command configures these parameters:

- **The host IP address(es)**

  for up to three TACACS+ servers; one first-choice and up to two backups. Designating backup servers provides for a continuation of authentication services in case the switch is unable to contact the first-choice server.

- **An optional encryption key.**

  This key helps to improve security, and must match the encryption key used in your TACACS+ server application. In some applications, the term "secret key" or "secret" may be used instead of "encryption key". If you need only one encryption key for the switch to use in all attempts to authenticate through a TACACS+ server, configure a global key. However, if the switch is configured to access multiple TACACS+ servers having different encryption keys, you can configure the switch to use different encryption keys for different TACACS+ servers.

- **The timeout value**

  in seconds for attempts to contact a TACACS+ server. If the switch sends an authentication request, but does not receive a response within the period specified by the timeout value, the switch resends the request to the next server in its Server IP Addr list, if any. If the switch still fails to receive a response from any TACACS+ server, it reverts to whatever secondary authentication method was configured using the `aaa authentication` command (local or none), see **Configuring the switch authentication methods** on page 165.

> **NOTE**
>
> As described in **General authentication setup procedure** on page 161, Hewlett Packard Enterprise recommends that youconfigure, test, and troubleshoot authentication via Telnet access before you configure authentication via console port access. This helps to prevent accidentally locking yourself out of switch access due to errors or problems in setting up authentication in either the switch or your TACACS+ server.

Syntax:

```
tacacs-server host <ip-addr> [key <key-string>]
```

Adds a TACACS+ server and optionally assigns a server-specific encryption key. If the switch is configured to access multiple TACACS+ servers having different encryption keys, you can configure the switch to use different encryption keys for different TACACS+ servers.

```
[no] tacacs-server host <ip-addr>
```

Removes a TACACS+ server assignment (including its server-specific encryption key, if any).

```
tacacs-server [key <key-string>]
```

Configures an optional global encryption key. Keys configured in the switch must exactly match the encryption keys configured in the TACACS+ servers that the switch will attempt to use for authentication.

```
[no] tacacs-server key
```

Removes the optional global encryption key. This does not affect any server-specific encryption key assignments.

```
tacacs-server timeout <1-255>
```

Changes the wait period for a TACACS server response.

Default: 5 seconds.

TACACS+ authorization and accounting commands
Command to enable authorization

Syntax:

```
[no] aaa authorization commands <radius|tacacs|local|auto|none>
```

```
[no] aaa authorization commands access-level <manager|all>
```

These commands are used to enable TACACS+ authorization.

Brief description of TACACS authorization options:

Local

Locally authenticated clients goes through local authorization. No authentication is performed for RADIUS/TACACS+ authenticate clients.

RADIUS

Locally authenticated clients go through local authorization. RADIUS authenticated clients go through RADIUS authorization. No authorization is performed for TACACS+ authenticated clients.

TACACS

TACACS authenticated clients go through TACACS authorization. No authorization is performed for RADIUS/locally authenticated users.

Auto

Uses the same method as Authentication and Authorization. For example local/radius/tacacs authenticated clients will go through local/radius/tacacs authorization respectively.

Authorization method:

Manager

Enables authorization for manager level commands.

All

Enables authorization for all commands. This is the default option.

Command to configure dead time

Syntax:

```
[no] tacacs-server dead-time <0-1440>
```

Configures the dead time for unavailable TACACS+ servers. When a server stops responding, the switch ignores this for a given amount of time and proceeds immediately to the next backup. Configuring the dead time improves server response time as the switch no longer has to wait for connections to time out before contacting the next backup server. The default value of zero disables skipping unavailable servers.

dead-time

Configure the dead time for unavailable TACACS+ servers.

0-1440

The server unavailability time in minutes (default is 0).

Command to enable authorization

Syntax:

```
[no] aaa authorization commands < radius | local | tacacs | auto | none >
```

```
[no] aaa authorization commands access-level < manager | all >
```

Configure command authorization. For each command issued by the user, an authorization request is sent to the server. Command authorization can be applied to all commands or only manager-level commands

AAA

Configure the switch Authentication, Authorization, and Accounting features.

Commands

Configure command authorization.

Group

Create or remove an authorization rule.

Local

Authorize commands using local groups.

RADIUS

Authorize commands using RADIUS.

None

Do not require authorization for command access.

TACACS

Authorize commands using TACACS+.

Auto

Authorize commands with the same protocol used for authentication.

Access-level

Configure command authorization level.

Manager

Allow authorization only for 'manager' level commands.

All

Allow authorization for all commands.

Command to enable accounting

Syntax:

```
[no] aaa accounting < exec | network | system | commands | start-stop | stop-only | intermim-update radius | syslog | tacacs>
```

```
[no] aaa accounting network < start-stop | stop-only | interim-update radius >
```

Configures the accounting service on the device. Accounting can be configured for EXEC sessions, network connection, commands and system. The accounting data is collected by a RADIUS, SYSLOG, or TACACS+ server.

> **NOTE**
>
> Network accounting is not supported through TACACS+ and SYSLOG. `session-id` accounting is not supported for TACACS+.

Periodic

Configures how periodic accounting updates are sent.

Null-username

Suppresses accounting for users with no username

Accounting

Configures the accounting service on the device.

Commands

Configures `command` type of accounting.

Exec

Configures Exec type of accounting.

Network

Configures network type of accounting.

Session-id

Configures accounting sessions identification scheme.

Suppress

Does not generate accounting records for a specific type of user.

System

Configures system type of accounting.

Update

Configures update accounting records mechanism.

RADIUS

Uses RADIUS for accounting.

TACACS

Uses TACACS+ for accounting.

Show all authorization configurations

Syntax:

```
show authorization group <group-name>
```

Show authorization configuration.

Authorization

Show authorization configuration.

---

show authorization

```
HP Switch(config)# show authorization
Status and Counters - Authorization Information
Access Level Requiring Authorization: Manager

 Type      | Method
 -------- + -------
 Commands |  tacacs
```

Show all accounting configurations

Syntax:

```
show accounting sessions
```

Shows accounting configuration parameters. If `sessions` is specified, the command will show accounting data for all active sessions.

Accounting

Shows accounting configuration parameters.

Show authorization

```
HP Switch(config)# show accounting
Status and Counters - Accounting Information
Interval(min) : 0
Suppress Empty User : No
Sessions Identification : Common

 Type      | Method Mode            Server Group
 -------- + ------ -------------- ------------
 Network  | None
 Exec     | None
 System   | tacacs Start-Stop     tacacs
 Commands | None
```

Show TACACS+

Syntax:

```
show tacacs host <IP-addr>
```

Show TACACS+ status and statistics information.

TACACS:

Shows TACACS+ status and statistics.

Example:

```
HP Switch# show tacacs
TACACS+ Information
Timeout : 5
Source IP Selection : Outgoing Interface
Encryption Key :

Server Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx
------------ ------ ------ ------ ------ ------- -------
```

Show TACACS+ host details

Syntax:

```
show tacacs host <IP-addr>
```

Show TACACS+ status and statistics information.

Host:

Show information for the specified TACACS+ server.

IP-addr:

The IP address of the TACACS+ server.

---

show tacacs

```
HP Switch# show tacacs

TACACS+ Information
Timeout : 5
Source IP Selection : Outgoing Interface
Encryption Key :
Server Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx OOBM
------------ ------ ------ ------ ------ ------- ------- ----
```

OOBM information is displayed only for switches that support OOBM.

---

show tacacs host

```
HP Switch(config)# show tacacs host <IP>
 TACACS+ Server Information

Server Addr      : 10.0.0.3

OOBM             : Enabled
Sessions Opened  :
```

```
Sessions Closed   :
Sessions Aborted  :
Sessions Error    :

Authentication :
Packets Tx        : 0       Packets Rx      : 0
Timeouts          : 0

Authorization :
Packets Tx        : 0       Packets Rx      : 0
Timeouts          : 0

Accounting :
Packets Tx        : 0       Packets Rx      : 0
Timeouts          : 0
```

OOBM information is displayed only for switches that support OOBM.

Show accounting sessions

Syntax:

```
show accounting sessions
```

Show accounting data for all active sessions.

Sessions:

Show accounting data for all active sessions.

Example:

```
HP Switch(config)#
Active Accounted actions on SWITCH, User (n/a) Priv (n/a),
Acct-Session-Id 0x013E00000006, System Accounting record, 1:45:34 Elapsed,
system event 'Accounting On', method 'radius'
Active Accounted actions on SWITCH, User (n/a) Priv (n/a),
Task-id 0x013E00000006, Command Accounting record, 1:45:34 Elapsed,
method 'tacacs'.
```

Device running a TACACS+ server application

Syntax:

```
host <ip-addr> [key <key-string>]
```

Specifies the IP address of a device running a TACACS+ server application. Optionally, can also specify the unique, per-server encryption key to use when each assigned server has its own, unique key. For more on the encryption key, see **Using the encryption key** on page 191 and the documentation provided with your TACACS + server application.

You can enter up to three IP addresses; one first-choice and two (optional) backups (one second-choice and one third-choice).

Use `show tacacs` to view the current IP address list.

If the first-choice TACACS+ server fails to respond to a request, the switch tries the second address, if any, in the show tacacs list. If the second address also fails, then the switch tries the third address, if any.

See **Figure 25: Example of the switch TACACS+ configuration listing** on page 165, the priority (first-choice, second-choice, and third-choice) of a TACACS+ server in the switch TACACS+ configuration depends on the order in which you enter the server IP addresses:

1. When there are no TACACS+ servers configured, entering a server IP address makes that server the first-choice TACACS+ server.

2. When there is one TACACS+ serves already configured, entering another server IP address makes that server the second-choice (backup) TACACS+ server.

3. When there are two TACACS+ servers already configured, entering another server IP address makes that server the third-choice (backup) TACACS+ server.

The above position assignments are fixed. If you remove one server and replace it with another, the new server assumes the priority position that the removed server had. For example, suppose you configured three servers, A, B, and C, configured in order:

First-Choice: A

Second-Choice: B

Third-Choice: C

If you removed server B and then entered server X, the TACACS+ server order of priority would be:

First-Choice: A

Second-Choice: X

Third-Choice: C

If there are two or more vacant slots in the TACACS+ server priority list and you enter a new IP address, the new address will take the vacant slot with the highest priority. Thus, if A, B, and C are configured as above and you (1) remove A and B, and (2) enter X and Y (in that order), then the new TACACS+ server priority list would be X, Y, and C.

The easiest way to change the order of the TACACS+ servers in the priority list is to remove all server addresses in the list and then re-enter them in order, with the new first-choice server address first, and so on.

To add a new address to the list when there are already three addresses present, you must first remove one of the currently listed addresses.

See also **General authentication process using a TACACS+ server** on page 190.

Default: None

Optional, global "encryption key"

Syntax:

```
key <key-string>
```

Specifies the optional, global "encryption key" that is also assigned in the TACACS+ servers that the switch will access for authentication. This option is subordinate to any "per-server" encryption keys you assign, and applies only to accessing TACACS+ servers for which you have not given the switch a "per-server" key. (See the `host <ip-addr> [key <key-string>]` entry above.)

You can configure a TACACS+ encryption key that includes a tilde (~) as part of the key, for example, "hp~switch".

For more on the encryption key, see **Using the encryption key** on page 191 and the documentation provided with your TACACS+ server application.

Configuring a host-specific key

```
HP Switch(config)# tacacs-server host 10.10.10.2 key hp~networking
```

Use the `show running-config` command to display the key information.

The running config file showing the host-specific key for TACACS+ with the "~" included

```
HP Switch(config)# show running-config

Running configuration:

; J9627A Configuration Editor; Created on release #XX.15.XX

hostname "Switch"
vlan 1
   name "DEFAULT_VLAN"
   untagged 1-24
   ip address dhcp-bootp
   exit
banner motd "good morning
tacacs-server host 10.10.10.2 key "hp~networking"
snmp-server community "public" unrestricted
```

Specifying how long the switch waits for a TACACS+ server to respond to an authentication request

Syntax:

`timeout <1-255>`

Specifies how long the switch waits for a TACACS+ server to respond to an authentication request. If the switch does not detect a response within the timeout period, it initiates a new request to the next TACACS+ server in the list. If all TACACS+ servers in the list fail to respond within the timeout period, the switch uses either local authentication (if configured) or denies access (if `none` configured for local authentication).

Default: 5 seconds

Adding, removing, or changing the priority of a TACACS+ server

Suppose the switch is configured to use TACACS+ servers at 10.28.227.10 and 10.28.227.15. 10.28.227.15 was entered first and so is listed as the first-choice server:

Figure 28: *Example of the switch with two TACACS+ server addresses configured*

```
HP Switch(config)# show tacacs
  Status and Counters - TACACS Information
    Timeout : 5
    Encryption Key :        First-Choice TACACS+ Server

    Server IP Addr   Opens   Closes Aborts Errors Pkts Rx Pkts Tx
    --------------   ------  ------ ------ ------ ------- -------
    10.28.227.15     0       0      0      0      0       0
    10.28.227.10     0       0      0      0      0       0
```

To move the "first-choice" status from the "15" server to the "10" server, use the `no tacacs-server host <ip-addr>` command to delete both servers, then use `tacacs-server host <ip-addr>` to re-enter the "10" server first, then the "15" server.

The servers would then be listed with the new "first-choice" server, that is:

Figure 29: *Example of the switch after assigning a different "first-choice" server*

```
HP Switch(config)# show tacacs
 Status and Counters - TACACS Information
   Timeout : 5
   Encryption Key          The "10" server is now the "first-choice" TACACS+ authentication device.

   Server IP Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx
   -------------- ------ ------ ------ ------ ------- -------
   10.28.227.10    0       0      0      0      0       0
   10.28.227.15    0       0      0      0      0       0
```

To remove the 10.28.227.15 device as a TACACS+ server, you would use this command:

```
HP Switch(config)# no tacacs-server host 10.28.227.15
```

Configuring an encryption key

Use an encryption key in the switch if the switch will be requesting authentication from a TACACS+ server that also uses an encryption key. (If the server expects a key, but the switch either does not provide one, or provides an incorrect key, then the authentication attempt will fail.)

- Use a global encryption key if the same key applies to all TACACS+ servers the switch may use for authentication attempts.

- Use a per-server encryption key if different servers the switch may use will have different keys.(For more details on encryption keys, see **Using the encryption key** on page 191.

Configuring a global encryption key

To configure **north01** as a global encryption key:

```
HP Switch(config) tacacs-server key north01
```

Configuring a per-server encryption key

To configure **north01** as a per-server encryption key:

```
HP Switch(config)# tacacs-server host 10.28.227.63 key north01
```

An encryption key can contain up to 100 characters, without spaces, and is likely to be case-sensitive in most TACACS+ server applications.

Deleting a global encryption key

To delete a global encryption key from the switch, use this command:

```
HP Switch(config)# no tacacs-server key
```

Deleting a per-server encryption key

To delete a per-server encryption key in the switch, re-enter the tacacs-server host command without the key parameter. For example, if you have **north01** configured as the encryption key for a TACACS+ server with an IP address of 10.28.227.104 and you want to eliminate the key, you would use this command:

```
HP Switch(config)# tacacs-server host 10.28.227.104
```

NOTE You can save the encryption key in a configuration file by entering this command:

```
HP Switch(config)# tacacs-server key <keystring>
```

The <*keystring* > parameter is the encryption key in clear text.

NOTE The `show tacacs` command lists the global encryption key, if configured. However, to view any configured per-server encryption keys, you must use `show config` or `show config running` (if you have made TACACS+ configuration changes without executing `write mem`).

Configuring the timeout period

The timeout period specifies how long the switch waits for a response to an authentication request from a TACACS+ server before either sending a new request to the next server in the switch Server IP Address list or using the local authentication option. For example, to change the timeout period from 5 seconds (the default) to 3 seconds:

```
HP Switch(config)# tacacs-server timeout 3
```

Configuring cipher text for TACACS+ key

To improve security, when entering a TACACS+ key, the key is displayed in CIPHER text. When the feature is active, a masking of plaintext keys is supported using the CIPHER key (*) for configuration commands where sensitive information is entered. Show commands are shown in plaintext.

Enabling the CLI command `hide-sensitive-data` supports the masking of the sensitive information. This command is disabled by default.

Process of configuring TACACS+ key with encrypt-credentials and hide-sensitive-data

Procedure

1. Enable the ciphertext feature at the switch.

    See **hide-sensitive-data**.

2. Enable the TACACS+ server key to accept the input of sensitive information in an enhanced secure mode.

    See **tacacs-server key**.

3. Encrypt passwords and authentication keys in show commands.

    See **encrypt-credentials**.

hide-sensitive-data

Syntax

```
hide-sensitive-data

no hide-sensitive-data
```

Description

Enables ciphertext feature in standard secure mode for the hiding of sensitive data. In the standard secure mode, the command is unavailable by default.

The `no` form of this command disables the ciphertext feature.

Command context

`config`

Restrictions

This command cannot be configured in enhanced-secure mode.

Example

Enable the ciphertext feature at the switch.

```
switch(config)# hide-sensitive-data
```

Example

After enabling ciphertext at the switch, use the command `include-credentials`.

```
Switch(config)# include-credentials

                        **** CAUTION ****

 You have invoked the command 'include-credentials'.  This action will make
 changes to the password and SSH public-key storage.

 It will affect *all* stored configurations, which might need to be updated.
 Those credentials will no longer be readable by older software revisions.
 It also may break some of your existing user scripts. Continue?[y/n] y

                        **** CAUTION ****

 This will insert possibly sensitive information in switch configuration files,
 and as a part of some CLI commands output. It is strongly recommended that you
 use SFTP rather than TFTP for transfer of the configuration over the network,
 and that you use the web configuration interface only with SSL enabled.

 Erasing configurations with 'include-credentials' enabled will erase stored
 passwords and security credentials.  The system will reboot with the factory
 default configuration.

 Proceed?[y/n] y
```

tacacs-server key

Syntax

`tacacs-server key`

`no tacacs-server key`

Description

The command `tacacs-server key` turns on the enhanced secure mode which uses the ciphertext for sensitive information during input.

After entering the command `hide-sensitive-data`, enable the enhanced secure mode for TACACS+ with the command `tacacs-server-key` to ensure enhanced security for sensitive information during input.

The `no` form of this command disables the enhanced secure mode of input for TACACS+.

Command context

config

Restrictions

This command is not allowed in enhanced secure mode.

Examples

Enabling `tacacs-server key` will hide sensitive information.

```
Switch(config)# tacacs-server key
Enter key-str: ********
Re-enter key-str:********
```

TACACS+ key configuration with `include-credentials`.

```
Switch(config)# tacacs-server key
    Enter key-str: ********
    Re-enter key-str:  ********

Switch(config)# tacacs-server host 10.0.0.10 key
    Enter key-str:  ********
    Re-enter key-str:  ********

Switch(config)# show include-credentials
    Stored in Configuration         : Yes
    Enabled in Active Configuration : Yes

Switch(config)# show encrypt-credentials
    Encryption    : Disabled
    Pre-shared Key: none

Switch(config)# show running-config

Running configuration:
    ; J9850A Configuration Editor; Created on release #KB.16.03.0000x
    ; Ver #0f:7f.ff.bb.ff.7c.59.fc.7b.ff.ff.fc.ff.ff.3f.ef:45
    hostname "HP-Switch-5406Rzl2"
    module A type j9989a
    module F type j9534a
    hide-sensitive-data
    include-credentials
    tacacs-server host 10.0.0.10 key "procurve"
    tacacs-server key "procurve"
    snmp-server community "public" unrestricted
    snmpv3 engineid "00:00:00:0b:00:00:a0:48:1c:f7:ee:00"
    oobm
   ip address dhcp-bootp
   exit
    vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24,F1-F24
   ip address dhcp-bootp
   exit

Switch(config)# show tacacs

 Status and Counters - TACACS Information
  Deadtime(min) : 0
  Timeout : 5
```

```
    Source IP Selection : Outgoing Interface
    Encryption Key : procurve

    Server IP Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx OOBM
    --------------- ------ ------ ------ ------ ------- ------- ----
    10.0.0.10       0      0      0      0      0       0       No
```

TACACS+ key configuration with `encrypt-credentials`.

```
Switch(config)# show encrypt-credentials

    Encryption    : Enabled
    Pre-shared Key: none

Switch(config)# show include-credentials
    Stored in Configuration         : Yes
    Enabled in Active Configuration : Yes

Switch(config)# tacacs-server key
    Enter key-str: ********
    Re-enter key-str:  ********

Switch(config)# tacacs-server host 10.0.0.10 key
    Enter key-str:  ********
    Re-enter key-str:  ********

Switch(config)# show running-config

Running configuration:

    ; J9850A Configuration Editor; Created on release #KB.16.03.0000x
    ; Ver #0f:7f.ff.bb.ff.7c.59.fc.7b.ff.ff.fc.ff.ff.3f.ef:45
    ; encrypt-cred 38qcQq/OETUfXNO7/eGOb5TgG3IBzILkhHOspcJkM2Y/5JvgL27NSkoQGjVEPz5a
    hostname "HP-Switch-5406Rzl2"
    module A type j9989a
    module F type j9534a
    encrypt-credentials
    hide-sensitive-data
    include-credentials
    tacacs-server host 10.0.0.10 encrypted-key
    "6T8PEZYO7uz4gIaWdWUg23gEZAjO33D21I6V2KOTECk="
    tacacs-server encrypted-key "HHa0HOmjKae6yzZ9Fn9JqZBuQhkGJV898+DHtb/3r9E="
    snmp-server community "public" unrestricted
    snmpv3 engineid "00:00:00:0b:00:00:a0:48:1c:f7:ee:00"
    oobm
   ip address dhcp-bootp
   exit
    vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24,F1-F24
   ip address dhcp-bootp
   exit

Switch(config)# show tacacs

 Status and Counters - TACACS Information

  Deadtime(min) : 0
  Timeout : 5
  Source IP Selection : Outgoing Interface
  Encryption Key : gJ5AeXfDFHJqjOOgOaa+NAmzneHDqs/aMqQuWsW01Qs=


  Server IP Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx OOBM
```

```
    --------------- ------ ------ ------ ------ ------- ------- ----
    10.0.0.10          0      0      0      0       0       0     No
```

TACACS+ key configuration without `include-credentials`.

```
Switch(config)# hide-sensitive-data

Switch(config)# tacacs-server key
    Enter key-str: ********
 Re-enter key-str: ********

Switch(config)# tacacs-server host 10.0.0.10 key
    Enter key-str: ********
    Re-enter key-str: ********

Switch(config)# show running-config

Running configuration:

 ; J9850A Configuration Editor; Created on release #KB.16.03.0000x
 ; Ver #0f:7f.ff.bb.ff.7c.59.fc.7b.ff.ff.fc.ff.ff.3f.ef:45
    hostname "HP-Switch-5406Rzl2"
    module A type j9989a
    module F type j9534a
    hide-sensitive-data
    tacacs-server host 10.0.0.10 key "test1"
    tacacs-server key "test"
    snmp-server community "public" unrestricted
    oobm
    ip address dhcp-bootp
    exit
    vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24,F1-F24
    ip address dhcp-bootp
    exit

Switch(config)#show include-credentials
 Stored in Configuration          : No
    Enabled in Active Configuration : N/A

Switch(config)# show encrypt-credentials
    Encryption   : Disabled
    Pre-shared Key: none

Switch(config)# show tacacs
 Status and Counters - TACACS Information
        Deadtime(min) : 0
  Timeout : 5
  Source IP Selection : Outgoing Interface
  Encryption Key : procurve

  Server IP Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx OOBM
  --------------- ------ ------ ------ ------ ------- ------- ----
  10.0.0.10          0      0      0      0       0       0     No
```

TACACS+ key configuration without `hide-sensitive-data`.

```
Switch(config)# tacacs-server key procurve
Switch(config)# tacacs-server host 10.0.0.10 key procurve

Switch(config)# show encrypt-credentials
    Encryption   : Enabled
```

```
      Pre-shared Key: none

Switch(config)# show running-config
    Running configuration:
    ; J9850A Configuration Editor; Created on release #KB.16.03.0000x
    ; Ver #0f:7f.ff.bb.ff.7c.59.fc.7b.ff.ff.fc.ff.ff.3f.ef:45
    ; encrypt-cred 38qcQq/OETUfXNO7/eGOb5TgG3IBzILkhHOspcJkM2Y/5JvgL27NSkoQGjVEPz5a
    hostname "HP-Switch-5406Rzl2"
    module A type j9989a
    module F type j9534a
    encrypt-credentials
    tacacs-server host 10.0.0.10 encrypted-key
    "GU3k9AV3u4eKyxBERotdYG87TbHLyv1RxVBnP3KhDhs="
    tacacs-server encrypted-key "7ViIcKdWMqJzWKDn/bT6AiAAehx3ASz+nldMZ9TI5eg="
    snmp-server community "public" unrestricted
    oobm
  ip address dhcp-bootp
  exit
   vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24,F1-F24
  ip address dhcp-bootp
  exit

Switch(config)# show tacacs

 Status and Counters - TACACS Information
  Deadtime(min) : 0
  Timeout : 5
  Source IP Selection : Outgoing Interface
  Encryption Key : gJ5AeXfDFHJqjOOgOaa+NAmzneHDqs/aMqQuWsW01Qs=

  Server IP Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx OOBM
  --------------- ------ ------ ------ ------ ------- ------- ----
  10.0.0.10       0      0      0      0      0       0       No
```

encrypt-credentials

Syntax

```
encrypt-credentials

no enrypt-credentials
```

Description

Encrypts all passwords and authentication keys in show commands.

The `no` form of this command removes encryption so that passwords and authentication keys are shown in plain text.

Command context

```
config
```

Examples

This example encrypts all credentials in show commands.

```
switch(config)# encrypt-credentials

                            **** CAUTION ****
    This will encrypt all passwords and authentication keys.
```

```
      The encrypted credentials will not be understood by older software versions.
      The resulting config file cannot be used by older software versions.
      It also may break some of your existing user scripts.

      Before proceeding, please save a copy of your current config file, and
      associate the current config file with the older software version saved in
      flash memory. See "Best Practices for Software Updates" in the Release Notes.

      A config file with 'encrypt-credentials' may prevent previous software
      versions from booting. It may be necessary to reset the switch to factory
      defaults. To prevent this, remove the encrypt-credentials command or use
      an older config file.
Save config and continue (y/n)? y

Switch(config)# tacacs-server key procurve

Switch(config)# show running-config

      Running configuration:

      ; J9850A Configuration Editor; Created on release #KB.16.03.0000x
      ; Ver #0f:7f.ff.bb.ff.7c.59.fc.7b.ff.ff.fc.ff.ff.3f.ef:45
      ; encrypt-cred 38qcQq/OETUfXNO7/eGOb5TgG3IBzILkhHOspcJkM2Y
   /5JvgL27NSkoQGjVEPz5a
      hostname "HP-Switch-5406Rzl2"
      module A type j9989a
      module F type j9534a
      encrypt-credentials
      tacacs-server encrypted-key "7ViIcKdWMqJzWKDn
      /bT6AiAAehx3ASz+nldMZ9TI5eg="
      snmp-server community "public" unrestricted
      oobm
   ip address dhcp-bootp
   exit
      vlan 1
   name "DEFAULT_VLAN"
   untagged A1-A24,F1-F24
   ip address dhcp-bootp
   exit

HP-Switch-5412Rzl2(config)# show tacacs

    Status and Counters - TACACS Information

     Deadtime(min) : 0
     Timeout : 5
     Source IP Selection : Outgoing Interface
     Encryption Key : 82qT9SBeCEs7iUtT7jSp
          /Jb2Xr5VMZPaB2YTveaq+F0=


     Server IP Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx OOBM
     --------------- ------ ------ ------ ------ ------- ------- ----
```

Configured the TACACS+ key with `encrypt-credentials`.

```
Switch(config)# encrypt-credentials

                         **** CAUTION ****
 This will encrypt all passwords and authentication keys.

 The encrypted credentials will not be understood by older software versions.
 The resulting config file cannot be used by older software versions.
```

```
   It also may break some of your existing user scripts.

   Before proceeding, please save a copy of your current config file, and
   associate the current config file with the older software version saved in
   flash memory. See "Best Practices for Software Updates" in the Release Notes.

   A config file with 'encrypt-credentials' may prevent previous software
   versions from booting. It may be necessary to reset the switch to factory
   defaults. To prevent this, remove the encrypt-credentials command or use
   an older config file.

       Save config and continue (y/n)? y

Switch(config)# hide-sensitive-data

Switch(config)# tacacs-server key
       Enter key-str:   ********
       Re-enter key-str:  ********

Switch(config)# tacacs-server host 10.0.0.10 key
       Enter key-str:   ********
       Re-enter key-str:  ********


Switch(config)# show include-credentials
       Stored in Configuration        : No
       Enabled in Active Configuration : N/A

Switch(config)# show encrypt-credentials
       Encryption    : Enabled
       Pre-shared Key: none

Switch(config)# show running-config

   Running configuration:

   ; J9850A Configuration Editor; Created on release #KB.16.03.0000x
   ; Ver #0f:7f.ff.bb.ff.7c.59.fc.7b.ff.ff.fc.ff.ff.3f.ef:45
   ; encrypt-cred 38qcQq/OETUfXNO7/eGOb5TgG3IBzILkhHOspcJkM2Y/
5JvgL27NSkoQGjVEPz5a
   hostname "HP-Switch-5406Rzl2"
   module A type j9989a
   module F type j9534a
   encrypt-credentials
   hide-sensitive-data
   tacacs-server host 10.0.0.10 encrypted-key
"6T8PEZYO7uz4gIaWdWUg23gEZAjO33D21I6V2KOTECk="
   tacacs-server encrypted-key "SV4/HLQCyOUoEspTiIEhsKPW21e6zfMDkJ1mdG8CrQc="
   snmp-server community "public" unrestricted
   oobm
  ip address dhcp-bootp
  exit
   vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A24,F1-F24
  ip address dhcp-bootp
 exit

Switch(config)# show tacacs

 Status and Counters - TACACS Information

  Deadtime(min) : 0
  Timeout : 5
```

```
Source IP Selection : Outgoing Interface
Encryption Key : gJ5AeXfDFHJqjOOgOaa+NAmzneHDqs/aMqQuWsW01Qs=


Server IP Addr  Opens  Closes Aborts Errors Pkts Rx Pkts Tx OOBM
--------------- ------ ------ ------ ------ ------- ------- ----
10.0.0.10         0       0      0      0      0       0      No
```

How authentication operates

General authentication process using a TACACS+ server

Authentication through a TACACS+ server operates generally as described below. For specific operating details, see the documentation you received with your TACACS+ server application.

Figure 30: *Using a TACACS+ server for authentication*
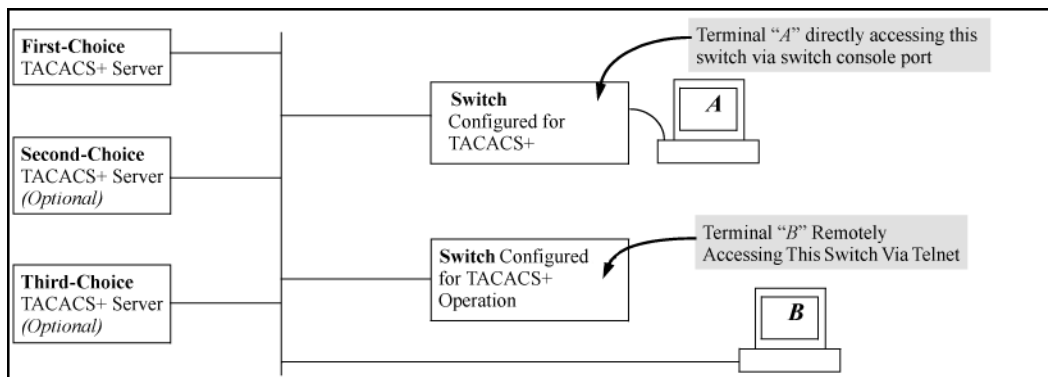
Using **Figure 30: Using a TACACS+ server for authentication** on page 190, after either switch detects an operator's logon request from a remote or directly connected terminal, the following events occur:

Procedure

1. The switch queries the first-choice TACACS+ server for authentication of the request.

    •   If the switch does not receive a response from the first-choice TACACS+ server, it attempts to query a secondary server. If the switch does not receive a response from any TACACS+ server, then it uses its own local username/password pairs to authenticate the logon request, see **Local authentication process (TACACS+)** on page 191.

    •   If a TACACS+ server recognizes the switch, it forwards a username prompt to the requesting terminal via the switch.

2. When the requesting terminal responds to the prompt with a username, the switch forwards it to the TACACS+ server.

3. After the server receives the username input, the requesting terminal receives a password prompt from the server via the switch.

4. When the requesting terminal responds to the prompt with a password, the switch forwards it to the TACACS+ server and one of the following actions occurs:

    a. If the username/password pair received from the requesting terminal matches a username/password pair previously stored in the server, then the server passes access permission through the switch to the terminal.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

b. If the username/password pair entered at the requesting terminal does not match a username/password pair previously stored in the server, access is denied. In this case, the terminal is again prompted to enter a username and repeat steps 2 through 4 In the default configuration, the switch allows up to three attempts to authenticate a login session. If the requesting terminal exhausts the attempt limit without a successful TACACS+ authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Local authentication process (TACACS+)

When the switch is configured to use TACACS+, it reverts to local authentication only if one of these two conditions exists:

- "Local" is the authentication option for the access method being used.

- TACACS+ is the primary authentication mode for the access method being used. Local is the secondary authentication method if the switch is unable to connect to any TACACS+ servers.

For a listing of authentication options, see **Configuring the switch TACACS+ server access** on page 172.

For local authentication, the switch uses the operator-level and manager-level username/password set(s) previously configured locally on the switch. (These are the usernames and passwords you can configure using the CLI password command, the WebAgent, or the menu interface–which enables only local password configuration).

- If the operator at the requesting terminal correctly enters the username/password pair for either access level, access is granted.

- If the username/password pair entered at the requesting terminal does not match either username/password pair previously configured locally in the switch, access is denied. In this case, the terminal is again prompted to enter a username/password pair. In the default configuration, the switch allows up to three attempts. If the requesting terminal exhausts the attempt limit without a successful authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

> **NOTE**
> The switch menu allows you to configure only the local operator and manager passwords, and not any usernames. In this case, all prompts for local authentication will request only a local password. However, if you use the CLI or the WebAgent to configure usernames for local access, you will see a prompt for both a local username and a local password during local authentication.

Using the encryption key

General operation

When used, the encryption key (sometimes termed "key", "secret key", or "secret") helps to prevent unauthorized intruders on the network from reading username and password information in TACACS+ packets moving between the switch and a TACACS+ server. At the TACACS+ server, a key may include both of the following:

- **Global key:**

  A general key assignment in the TACACS+ server application that applies to all TACACS-aware devices for which an individual key has not been configured.

- **Server-Specific key:**

  A unique key assignment in the TACACS+ server application that applies to a specific TACACS-aware device.

Thus, on the TACACS+ server side, you have a choice as to how to implement a key. On the switch side, it is necessary only to enter the key parameter so that it exactly matches its counterpart in the server. For information on how to configure a general or individual key in the TACACS+ server, see the documentation you received with the application.

Encryption options in the switch

When configured, the encryption key causes the switch to encrypt the TACACS+ packets it sends to the server. When left at "null", the TACACS+ packets are sent in clear text. The encryption key (or just "key") you configure in the switch must be identical to the encryption key configured in the corresponding TACACS+ server. If the key is the same for all TACACS+ servers the switch will use for authentication, then configure a global key in the switch. If the key is different for one or more of these servers, use "server-specific" keys in the switch. (If you configure both a global key and one or more per-server keys, the per-server keys will override the global key for the specified servers.)

For example, you would use the next command to configure a global encryption key in the switch to match a key entered as **north40campus** in two target TACACS+ servers. (That is, both servers use the same key for your switch.) Note that you do not need the server IP addresses to configure a global key in the switch:

```
HP Switch(config)# tacacs-server key north40campus
```

Suppose that you subsequently add a third TACACS+ server (with an IP address of 10.28.227.87) that has **south10campus** for an encryption key. Because this key is different than the one used for the two servers in the previous example, you will need to assign a server-specific key in the switch that applies only to the designated server:

```
HP Switch(config)# tacacs-server host 10.28.227.87 key south10campus
```

With both of the above keys configured in the switch, the **south10campus** key overrides the **north40campus** key only when the switch tries to access the TACACS+ server having the 10.28.227.87 address.

Controlling WebAgent access when using TACACS+ authentication

Configuring the switch for TACACS+ authentication does not affect WebAgent access. To prevent unauthorized access through the WebAgent, do one or more of the following:

- Configure local authentication (a manager user name and password and, optionally, an operator user name and password) on the switch.

- Configure the switch Authorized IP manager feature to allow WebAgent access only from authorized management stations. (The Authorized IP manager feature does not interfere with TACACS+ operation.)

- Disable WebAgent access to the switch by going to the System Information screen in the Menu interface and configuring the `Web Agent Enabled` parameter to `No`.

Messages related to TACACS+ operation

Table 13: *Messages related to TACACS+ operation*

Message	Meaning
Connecting to TACACS server	The switch is attempting to contact the TACACS+ server identified in the switch's `tacacs-server` configuration as the first-choice (or only) TACACS+ server.
Connecting to secondary TACACS server	The switch was not able to contact the first-choice TACACS+ server, and is now attempting to contact the next (secondary) TACACS+ server identified in the switch's `tacacs-server` configuration.
Invalid password	The system does not recognize the username or the password or both. Depending on the authentication method (tacacs or local), either the TACACS+ server application did not recognize the username/password pair or the username/password pair did not match the username/password pair configured in the switch.
No TACACS server(s) responding.	The switch has not been able to contact any designated TACACS+ servers. If this message is followed by the Username prompt, the switch is attempting local authentication.
Not legal combination of authentication methods.	For console access, if you select `tacacs` as the primary authentication method, you must select `local` as the secondary authentication method. This prevents you from being locked out of the switch if all designated TACACS+ servers are inaccessible to the switch.
Record already exists	When resulting from a `tacacs-server host <ip addr>` command, indicates an attempt to enter a duplicate TACACS+ server IP address.

Operating notes

- If you configure Authorized IP managers on the switch, it is not necessary to include any devices used as TACACS+ servers in the authorized manager list. That is, authentication traffic between a TACACS+ server and the switch is not subject to Authorized IP manager controls configured on the switch. Also, the switch does not attempt TACACS+ authentication for a management station that the Authorized IP manager list excludes because, independent of TACACS+, the switch already denies access to such stations.

- When TACACS+ is not enabled on the switch-or when the switch only designated TACACS+ servers are not accessible-setting a local operator password without also setting a local manager password does not protect the switch from manager-level access by unauthorized persons.

- When using the `copy` command to transfer a configuration to a TFTP server, any optional, server-specific and global encryption keys in the TACACS configuration will not be included in the transferred file. Otherwise, a security breach could occur, allowing access to the TACACS+ username/password information.

Overview

RADIUS (Remote Authentication Dial-In User Service) enables you to use up to 15 servers and maintain separate authentication and accounting for each RADIUS server employed.

Authentication Services

Authentication with RADIUS allows for a unique password for each user, instead of the need to maintain and distribute switch-specific passwords to all users.

RADIUS verifies identity for the following types of primary password access to the switch:

- Serial port (console)

- Telnet

- SSH

- SFTP/SCP

- WebAgent

- Port-Access (802.1X)

The switch also supports RADIUS accounting for Web Authentication and MAC authentication sessions.

NOTE	The switch does not support RADIUS security for SNMP (network management) access. For information on blocking access through the WebAgent, see **Controlling WebAgent access** on page 213.

Accounting services

RADIUS accounting on the switch collects resource consumption data and forwards it to the RADIUS server. This data can be used for trend analysis, capacity planning, billing, auditing, and cost analysis.

RADIUS-administered CoS and rate-limiting

The switches covered in this guide take advantage of vendor-specific attributes (VSAs) applied in a RADIUS server to support these optional, RADIUS-assigned attributes:

- 802.1p (CoS) priority assignment to inbound traffic on the specified port(s) (port-access authentication only)

- Per-Port Rate-Limiting on a port with an active link to an authenticated client (port-access authentication only)

RADIUS-administered commands authorization

This feature enables RADIUS server control of an authenticated client's access to CLI commands on the switch. See **Commands authorization** on page 214.

SNMP access to the switch authentication configuration MIB

The switch default configuration allows SNMP access to the `hpSwitchAuth` MIB (Management Information Base). A management station running an SNMP networked device management application can access the switch MIB for read access to the switch status and read/write access to the switch configuration. For more information, including the CLI command to use for disabling this feature, see **Using SNMP to view and configure switch authentication features** on page 210.

Switch operating rules for RADIUS

- You must have at least one RADIUS server accessible to the switch.

- The switch supports authentication and accounting using up to 15 RADIUS servers. The switch accesses the servers in the order in which they are listed by `show radius`. If the first server does not respond, the switch tries the next one, and so-on. To change the order in which the switch accesses RADIUS servers, see **Changing RADIUS-server access order** on page 237.

- You can select RADIUS as the primary authentication method for each type of access. (Only one primary and one secondary access method is allowed for each access type.)

- In the switch, EAP RADIUS uses MD5 and TLS to encrypt a response to a challenge from a RADIUS server.

- When primary/secondary authentication is set to Radius/Local (for either Login or Enable) and the RADIUS server fails to respond to a client attempt to authenticate, the failure is noted in the Event Log with the message `radius: Can't reach RADIUS server <server-ip-addr>` . When this type of failure occurs, the switch prompts the client again to enter a username and password. In this case, use the local username (if any) and password configured on the switch itself.

- Zero-length usernames or passwords are not allowed for RADIUS authentication, even though allowed by some RADIUS servers.

General RADIUS setup procedure

Preparation:

Procedure

1. Configure one to 15 RADIUS servers to support the switch. See the documentation provided with the RADIUS server application.

2. Before configuring the switch, collect the following information:

   a. Determine the access methods (console, Telnet, Port-Access (802.1X), WebAgent and/or SSH) for which you want RADIUS as the primary authentication method. Consider both operator (login) and manager

(enable) levels, and what secondary authentication methods to use (local or none) if the RADIUS authentication fails or does not respond.

Figure 31: *Possible RADIUS access assignments*

```
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled

             | Login        Login         Login
  Access Task | Primary      Server Group  Secondary
  ----------- + ----------   ------------  ----------
  Console     | Local                      Local
  Telnet      | Local                      None
  Port-Access | Local                      None
  Webui       | Local                      None
  SSH         | Local                      None
  Web-Auth    | ChapRadius radius          None
  MAC-Auth    | ChapRadius radius          None
  SNMP        | Local                      None


             | Enable       Enable        Enable
  Access Task | Primary      Server Group  Secondary
  ----------- + ----------   ------------  ----------
  Console     | Local                      None
  Telnet      | Local                      None
  Webui       | Local                      None
  SSH         | Local                      None
```

Note: The **WebAgent** access task shown in this figure is available only on the switches covered in this guide.

Console access requires Local as secondary method to prevent lockout if the primary RADIUS access fails due to loss of RADIUS server access or other problems with the server.

b. Determine the IP addresses of the RADIUS servers to support the switch. You can configure the switch for up to 15 RADIUS servers. See the documentation provided with the RADIUS server application for more information.

c. If you need to replace the default UDP destination port (1812) the switch uses for authentication requests to a specific RADIUS server, select it before beginning the configuration process.

d. If you need to replace the default UDP destination port (1813) the switch uses for accounting requests to a specific RADIUS server, select it before beginning the configuration process.

e. Determine whether to use one global encryption key for all RADIUS servers or if unique keys will be required for specific servers. With multiple RADIUS servers, if one key applies to two or more of these servers, then you can configure this key as the global encryption key. For any server whose key differs from the global key you are using, you must configure that key in the same command that you use to designate that server IP address to the switch.

f. Determine an acceptable timeout period for the switch to wait for a server to respond to a request. Hewlett Packard Enterprise recommends that you begin with the default (five seconds).

g. Determine how many times the switch can contact a RADIUS server before trying another RADIUS server or quitting. This depends on how many RADIUS servers you have configured the switch to access.

h. Determine whether you want to bypass a RADIUS server that fails to respond to requests for service. To shorten authentication time, you can set a bypass period in the range of 1 to 1440 minutes for nonresponsive servers. This requires that you have multiple RADIUS servers accessible for service requests.

i. Optional: Determine whether the switch access level (manager or operator) for authenticated clients can be set by a Service Type value the RADIUS server includes in its authentication message to the switch, see **Enabling manager access privilege (optional)** on page 200.

j. Configure RADIUS on servers used to support authentication on the switch.

Configuring the switch for RADIUS authentication

Configure RADIUS authentication for controlling access through one or more of the following:

- Serial port
- Telnet
- SSH
- Port-Access (802.1X)
- WebAgent

Procedure

1. RADIUS authentication on the switch must be enabled to override the default authentication operation which is to automatically assign an authenticated client to the operator privilege level. This applies the privilege level specified by the service type value received from the RADIUS server.

2. Configure the switch for accessing one or more RADIUS servers (one primary server and up to two backup servers):

   - Server IP address
   - (Optional) UDP destination port for authentication requests (default: 1812; recommended)
   - (Optional) UDP destination port for accounting requests (default: 1813; recommended)
   - (Optional) Encryption key for use during authentication sessions with a RADIUS server. This key overrides the global encryption key you can also configure on the switch, and must match the encryption key used on the specified RADIUS server. (default: null)

   > **NOTE:** Step 2 assumes you have already configured the RADIUS servers to support the switch. See your RADIUS server documentation for details.

3. Configure the global RADIUS parameters.

   a. Server key: This key must match the encryption key used on the RADIUS servers the switch contacts for authentication and accounting services unless you configure one or more per-server keys. (default: null)

   b. Timeout period: The timeout period the switch waits for a RADIUS server to reply. (default: 5 seconds; range: 1 to 15 seconds)

c. **Retransmit attempts:** The number of retries when there is no server response to a RADIUS authentication request. (default: 3; range of 1 to 5)

d. **Server dead-time:** The period during which the switch will not send new authentication requests to a RADIUS server that has failed to respond to a previous request. This avoids a wait for a request to time out on a server that is unavailable. If you want to use this feature, select a dead-time period of 1 to 1440 minutes. (default: disabled; range: 1-1440 minutes.) If your first-choice server was initially unavailable, but then becomes available before the dead-time expires, you can nullify the dead-time by resetting it to zero and then trying to log on again. As an alternative, you can reboot the switch, (thus resetting the dead-time counter to assume the server is available) and then try to log on again.

e. **Number of login attempts:** This is actually an `aaa authentication` command. It controls how many times per session a RADIUS client (and clients using other forms of access) can try to log in with the correct username and password. (default: Three times per session)

Configuring authentication for the access methods that RADIUS protects

Configure the switch for RADIUS authentication through the following access methods:

- **Console:** Either direct serial-port connection or modem connection.

- **Telnet:** Inbound Telnet must be enabled (the default).

- **SSH:** To use RADIUS for SSH access, first configure the switch for SSH operation.

- **WebAgent:** You can enable RADIUS authentication for WebAgent access to the switch.

You can configure RADIUS as the primary password authentication method for the above access methods. You also need to select either `local`, `none`, or `authorized` as a secondary, or backup, method. Note that for console access, if you configure `radius` (or `tacacs`) for primary authentication, you must configure `local` for the secondary method. This prevents the possibility of being completely locked out of the switch in the event that all primary access methods fail.

Syntax:

```
aaa authentication <console|telnet|ssh|web|<enable|login <local|radius>> web-based|mac-based <chap-radius|peap-radius>>
```

Configures RADIUS as the primary password authentication method for console, Telnet, SSH, and/or the WebAgent.

The default primary `<enable|login>` authentication is `local`.

```
<console|telnet|ssh|web>
```

```
[<local|none|authorized>]
```

Provides options for secondary authentication. For console access, secondary authentication must be `local` if primary access is not `local`. This prevents you from being locked out of the switch in the event of a failure in other access methods.

```
<<web-based|mac-based> login> <chap-radius|peap-mschapv2>
```

Password authentication for web-based or MAC-based port access to the switch. Use `peap-mschapv2` when you want password verification without requiring access to a plain text password; it is more secure.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Default: `chap-radius`

`[none|authorized]`

Provides options for secondary authentication. The `none` option specifies that a backup authentication method is not used. The `authorized` option allows access without authentication.

Default: `none`.

In certain situations, RADIUS servers can become isolated from the network. Users are not able to access the network resources configured with RADIUS access protection and are rejected. To address this situation, configuring the `authorized` secondary authentication method allows users unconditional access to the network when the primary authentication method fails because the RADIUS servers are unreachable.

> ⚠️ **CAUTION**
>
> Configuring `authorized` as the secondary authentication method used when there is a failure accessing the RADIUS servers allows clients to access the network unconditionally. Use this method with care.

Here is an example of the `show authentication` command displaying `authorized` as the secondary authentication method for port-access, web-based authentication access, and MAC authentication access. Since the configuration of `authorized` means no authentication will be performed and the client has unconditional access to the network, the "Enable Primary" and "Enable Secondary" fields are not applicable (N/A).

Figure 32: *Example of AAA authentication using authorized for the secondary authentication method*

```
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled

             | Login       Login         Login
 Access Task | Primary     Server Group  Secondary
 ----------- + ---------- ------------ ----------
 Console     | Local                     None
 Telnet      | Local                     None
 Port-Access | Local                     Authorized  ◄
 Webui       | Local                     None
 SSH         | Local                     None
 Web-Auth    | ChapRadius  radius        Authorized  ◄
 MAC-Auth    | ChapRadius  radius        None

             | Enable      Enable        Enable
 Access Task | Primary     Server Group  Secondary
 ----------- + ---------- ------------ ----------
 Console     | Local                     None
 Telnet      | Local                     None
 Webui       | Local                     None
 SSH         | Local                     None
```

The access methods with secondary authentication configured as **authorized** allows the client access to the network even if the RADIUS server is unreachable.

Suppose you already configured local passwords on the switch, but want RADIUS to protect primary Telnet and SSH access without allowing a secondary Telnet or SSH access option (the switch local passwords):

Figure 33: *Example configuration for RADIUS authentication*

```
HP Switch(config)# aaa authentication telnet login radius none
HP Switch(config)# aaa authentication telnet enable radius none
HP Switch(config)# aaa authentication ssh login radius none
HP Switch(config)# aaa authentication ssh enable radius none
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled

              | Login        Login         Login
  Access Task | Primary      Server Group  Secondary
  ----------- + ----------   ------------  ----------
  Console     | Local                      None
  Telnet      | Radius                     None
  Port-Access | Local                      None
  Webui       | Local                      None
  SSH         | Radius                     None
  Web-Auth    | ChapRadius   radius        None
  MAC-Auth    | ChapRadius   radius        None

              | Enable       Enable        Enable
  Access Task | Primary      Server Group  Secondary
  ----------- + ----------   ------------  ----------
  Console     | Local                      None
  Telnet      | Radius                     None
  Webui       | Local                      None
  SSH         | Radius                     None
```

The switch now allows Telnet and SSH authentication only through RADIUS.

NOTE

If you configure the **Login Primary** method as `local` instead of `radius` (and local passwords are configured on the switch), then clients connected to your network can gain access to either the operator or manager level without encountering the RADIUS authentication specified for **Enable Primary**. See **Local authentication process (RADIUS)** on page 213.

Enabling manager access privilege (optional)

In the default RADIUS operation, the switch automatically admits any authenticated client to the login (operator) privilege level, even if the RADIUS server specifies enable (manager) access for that client. Thus, an authenticated user authorized for the manager privilege level must authenticate again to change privilege levels. Using the optional `login privilege-mode` command overrides this default behavior for clients with enable access. That is, with `privilege-mode` enabled, the switch immediately allows enable (manager) access to a client for whom the RADIUS server specifies this access level.

Syntax:

```
[no] aaa authentication login privilege-mode
```

When enabled, the switch reads the Service-Type field in the client authentication received from a RADIUS server. The following table describes the applicableService-Type values and corresponding client access levels the switch allows upon authentication by the server.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Table 14: *Service-type value*

Service-type	Value	Client access level
Administrative-user	6	manager
NAS-prompt-user	7	operator
Any other type	Any value **except** 6 or 7	Access Denied

This feature applies to console (serial port), Telnet, SSH, and WebAgent access to the switch. It does not apply to 802.1X port-access.

> **NOTE**
> While this option is enabled, a Service-Type value other than 6 or 7, or an unconfigured (null) Service-Type causes the switch to deny access to the requesting client.

The `no` form of the command returns the switch to the default RADIUS authentication operation. The default behavior for most interfaces is that a client authorized by the RADIUS server for Enable (manager) access will be prompted twice, once for Login (operator) access and once for Enable access. In the default RADIUS authentication operation, the WebAgent requires only one successful authentication request. For more information on configuring the Service Type in your RADIUS application, see the documentation provided with the application.

Configuring the switch to access a RADIUS server

This section describes how to configure the switch to interact with a RADIUS server for both authentication and accounting services.

> **NOTE**
> If you want to configure RADIUS accounting on the switch, see **Accounting services** on page 222.

Syntax:

```
[no] radius-server host <ip-address>
```

Adds a server to the RADIUS configuration or (with `no`) deletes a server from the configuration. You can configure up to three RADIUS servers, and up to 15 RADIUS server addresses. See **Using multiple RADIUS server groups** on page 206 for information about grouping multiple RADIUS servers.

The switch uses the first server it successfully accesses, see **Changing RADIUS-server access order** on page 237.

```
[auth-port <port-number>]
```

Optional. Changes the UDP destination port for authentication requests to the specified RADIUS server (host). If you do not use this option with the `radius-server host` command, the switch automatically assigns the default authentication port number. The `auth-port` number must match its server counterpart.

Default: **1812**

```
[acct-port <port-number>]
```

Optional. Changes the UDP destination port for accounting requests to the specified RADIUS server. If you do not use this option with the `radius-server host` command, the switch automatically assigns the default accounting port number. The `acct-port` number must match its server counterpart.

Default: **1813**

```
[dyn-authorization]
```

Enables or disables the processing of Disconnect and Change of Authorization messages from this host. When enabled, the RADIUS server can dynamically terminate or change the authorization parameters (such as VLAN assignment) used in an active client session on the switch. The UDP port specified in the `radius-server dyn-autz-port` command (defaults to 3799) is the port used to listen for Change of Authorization messages (CoA) or Disconnect messages (DM).

Default: **Disabled**

```
[key <key-string>]
```

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the specified server. This key must match the encryption key used on the RADIUS server. Use this command only if the specified server requires a different encryption key than configured for the global encryption key.

> **NOTE**
> Formerly, when you saved the configuration file using Xmodem or TFTP, the RADIUS encryption key information was not saved in the file. This caused RADIUS authentication to break when the startup configuration file was loaded back onto the switch. You now can save the configured RADIUS shared secret (encryption) key to a configuration file by entering the following commands:
>
> ```
> include-credentials
> ```
>
> ```
> write memory
> ```

```
[time-window <0-65535>]
```

The time window in seconds within which the received dynamic authorization requests are considered to be current and accepted for processing. A zero value means there is no time limit. A non-zero value indicates that the even-timestamp attribute is expected as part of all Change of Authorization and Disconnect request messages. If the timestamp attribute is not present the message is dropped.

Default: **300 seconds.**

```
no radius-server host <ip-address> key
```

Use the `no` form of the command to remove the key for a specified server.

Suppose you have configured the switch as shown in the first figure and you now need to make the following changes:

Procedure

1. Change the encryption key for the server at 10.33.18.127 to "source0127".

2. Add a RADIUS server with an IP address of 10.33.18.119 and a server-specific encryption key of "source0119".

Figure 34: *Configuration for RADIUS server before changing the key and adding another server*

```
HP Switch(config)# radius-server host 10.22.18.127 key source0127
HP Switch(config)# radius-server host 10.22.18.119 key source0119

HP Switch# show radius

 Status and Counters - General RADIUS Information

  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key :
  Dynamic Authorization UDP Port : 3799
  Source IP Selection : Outgoing Interface


                  Auth Acct DM/ Time
  Server IP Addr  Port Port CoA Window Encryption Key          OOBM
  --------------- ---- ---- --- ------ -------------------------- ----
  10.33.18.127    1812 1813 No  300    TempKey01               No
```

OOBM information is displayed only for switches that support OOBM.

To make these changes, perform the following:

Figure 35: *Configuration for RADIUS server after changing the key and adding another server*

```
HP Switch(config)# radius-server host 10.33.18.127 key source0127 *
HP Switch(config)# radius-server host 10.33.18.119 key source0119 **
HP Switch(config)# show radius

 Status and Counters - General RADIUS Information

  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key :myg10balkey
  Dynamic Authorization UDP Port : 3799


                  Auth Acct DM/ Time
  Server IP Addr  Port Port CoA Window Encryption Key
  --------------- ---- ---- --- ------ ----------------
  10.33.18.127    1812 1813 No  10     source0127
  10.33.18.119    1812 1813 No  10     source0119
```

* Changes the key for the existing server to "source012 7" (step 1, above).

** Adds the new RADIUS server with its required "source0119" key.

To change the order in which the switch accesses RADIUS servers, see **Changing RADIUS-server access order** on page 237.

Configuring the switch global RADIUS parameters

Configure the switch for the following global RADIUS parameters:

- **Number of login attempts:** In a given session, this specifies how many tries at entering the correct username and password pair are allowed before access is denied and the session terminated. This is a general `aaa authentication` parameter and is not specific to RADIUS.

- **Global server key:** The server key the switch uses for contacts with all RADIUS servers for which there is not a server-specific key configured by `radius-server host <ip-address> key <key-string>` . This key is optional if you configure a server-specific key for each RADIUS server entered in the switch.

- **Server timeout:** Defines the time period in seconds for authentication attempts. If the timeout period expires before a response is received, the attempt fails.

- **Server dead time:** Specifies the time in minutes during which the switch avoids requesting authentication from a server that has not responded to previous requests.

- **Retransmit attempts:** If the first attempt to contact a RADIUS server fails, this specifies how many retries to allow the switch to attempt on that server.

- **Change of Authorization port:** The `dyn-autz-port` parameter specifies the UDP port number that listens for the Change of Authorization and Disconnect messages.The UDP port range is 1024-49151. The default port is 3799.

Syntax:

```
aaa authentication num-attempts <1-10>
```

Specifies how many tries for entering the correct username and password are allowed before shutting down the session due to input errors.

Default: **3**; Range: 1–10.

```
[no] radius-server
```

```
key <global-key-string>
```

Specifies the global encryption key the switch uses with servers for which the switch does not have a server-specific key assignment. This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key.

Default: Null.

```
[encrypted-key <global-key-string>]
```

Global encryption key, specified using a base64–encoded aes-256 encrypted string.

```
dead-time <1-1440>
```

Optional. Specifies the time in minutes during which the switch will not attempt to use a RADIUS server that has not responded to an earlier authentication attempt.

Default: **0**; Range: 1–1440 minute

```
dyn-autz-port <1024-49151>
```

Specifies the UDP port number that listens for Change of Authorization or Disconnect messages. The range of ports is 1024–49151.

Default: **3799**

```
radius-server timeout <1-15>
```

Specifies the maximum time the switch waits for a response to an authentication request before counting the attempt as a failure.

Default: **5** seconds; Range: 1–15 seconds

```
radius-server retransmit <1-5>
```

If a RADIUS server fails to respond to an authentication request, specifies how many retries to attempt before closing the session.

Default: **3**; Range: 1–5

> **NOTE** Where the switch has multiple RADIUS servers configured to support authentication requests, if the first server fails to respond, then the switch tries the next server in the list, and so-on. If none of the servers respond, then the switch attempts to use the secondary authentication method configured for the type of access being attempted (console, Telnet, or SSH). If this occurs, see the "Troubleshooting" chapter of the management and configuration guide for your switch.

Example:

Suppose your switch is configured to use three RADIUS servers for authenticating access through Telnet and SSH. Two of these servers use the same encryption key. In this case the plan is to configure the switch with the following global authentication parameters:

- Allow only two tries to correctly enter username and password.

- Use the global encryption key to support the two servers that use the same key. (For this example, assume that you did not configure these two servers with a server-specific key.)

- Use a dead-time of five minutes for a server that fails to respond to an authentication request.

- Allow three seconds for request timeouts.

- Allow two retries following a request that did not receive a response.

Global configuration exercise for RADIUS authentication

```
HP Switch(config)# aaa authentication num-attempts 2
HP Switch(config)# radius-server key My-Global-KEY-1099
HP Switch(config)# radius-server dead-time 5
HP Switch(config)# radius-server timeout 3
```

```
HP Switch(config)# radius-server retransmit 2
HP Switch(config)# write mem
```

Figure 36: *Global RADIUS parameters configured in Global configuration exercise for RADIUS authentication*

```
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

   Login Attempts : 2                           After two attempts failing due to
   Respect Privilege : Disabled                 username or password entry
                                                 errors, the switch will terminate
                                                 the session.
              | Login      Login        Login
   Access Task | Primary    Server Group Secondary
   ----------- + ---------- ------------ ----------
   Console     | Local                   None
   Telnet      | Radius                  None
   Port-Access | Local                   Authorized
   Webui       | Local                   None
   SSH         | Radius                  None
   Web-Auth    | ChapRadius radius       Authorized
   MAC-Auth    | ChapRadius radius       None


              | Enable     Enable       Enable
   Access Task | Primary    Server Group Secondary
   ----------- + ---------- ------------ ----------
   Console     | Local                   None
   Telnet      | Radius                  None
   Webui       | Local                   None
   SSH         | Radius                  None

HP Switch(config)# show radius              Global RADIUS parameters
                                            from figure 6-6.
 Status and Counters - General RADIUS Information
                                            Server-specific encryption key
   Deadtime(min) : 5                        for the RADIUS server that will
   Timeout(secs) : 3                        not use the global encryption
   Retransmit Attempts : 2                  key.
   Global Encryption Key : My-Global-Key-1099
   Dynamic Authorization UDP Port : 3799

                 Auth Acct DM/ Time
   Server IP Addr  Port Port CoA Window Encryption Key                        OOBM
   --------------- ---- ---- --- ------ -------------------------------- -----
   10.33.18.127    1812 1813 No  10     source0127                       No
   10.33.18.119    1812 1813 No  10                                      No
   10.33.18.151    1812 1813 No  10                These two servers will use the
                                                   global encryption key.
```

Using multiple RADIUS server groups

The authentication and accounting features on the switch can use up to fifteen RADIUS servers and these servers can be put into groups. Up to 5 groups of 3 RADIUS servers each can be configured. The authentication and accounting features can choose which RADIUS server group to communicate with. End-user authentication methods (802.1X, MAC-based and web-based) can authenticate with different RADIUS servers from the management interface authentication methods (console, telnet, ssh, web).

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Several commands are used to support the RADIUS server group option. The RADIUS server must be configured before it can be added to a group. See **Configuring the switch for RADIUS authentication** on page 197 for more information.

Figure 37: *RADIUS server group command output*

```
HP Switch(config)# radius-server host 10.33.18.151 acct-port 1750 key
source0151
HP Switch(config)# write mem

HP Switch(config)# show radius

 Status and Counters - General RADIUS Information
                                    Because the radius-server command includes an acct-
                                    port keyword with a non-default UDP port number of
  Deadtime(min) : 0                 1750, the switch assigns this value as the UDP
  Timeout(secs) : 5                 accounting port.
  Retransmit Attempts : 3
  Global Encryption Key :
  Dynamic Authorization UDP Port : 3799

                  Auth Acct DM/ Time
  Server IP Addr  Port Port CoA Window Encryption Key                  OOBM
  --------------- ---- ---- --- ------ ------------------------------- -----
  10.33.18.151    1812 1750 No  10     source0151                      No
```

Connecting a RADIUS server with a server group

Syntax:

[no] radius-server host *<ip-address>*

Adds a server to the RADIUS configuration or (with no) deletes a server from the configuration. You can configure up to fifteen RADIUS server addresses. The switch uses the first server it successfully accesses.

Syntax:

aaa server-group radius *<group-name>* host *<ip-addr>*

no aaa server-group radius *<group-name>* host *<ip-addr>*

Associates a RADIUS server with a server group. Each group can contain up to 3 RADIUS servers. The default group (called 'radius'), can only contain the first three RADIUS servers. The default group cannot be edited.

The no form of the command removes the RADIUS server with the indicated IP address from the server group. If that server was the last entry in the group, the group is removed.

radius *<group-name>* : The group name of the RADIUS server group. The name has a maximum length of 12 characters. Up to five groups can be configured with a a maximum of three RADIUS servers in each group. The first group slot is used by the default group.

host *<ip-addr>* : The IP address of the RADIUS server used.

Configuring the primary password authentication method for console, Telnet, SSH and WebAgent

The following commands have the server-group option. If no server-group is specified, the default RADIUS group is used. The server group must already be configured.

> **NOTE**
> The last RADIUS server in a server group cannot be deleted if any authentication or accounting method is using the server group.

Syntax:

```
aaa authentication <console|telnet|ssh|web> <enable|[login <local]|radius [server-group <group-name>|local|none|authorized]
```

Configures the primary password authentication method for console, Telnet, SSH, and the WebAgent.

```
{<local | radius>}
```

Primary authentication method.

Default: **local**

```
<local|radius>
```

Use either the local switch user/password database or a RADIUS server for authentication.

```
<server-group <group-name>>
```

Specifies the server group to use.

```
[local|none|authorized]
```

Provides options for secondary authentication.

Default: **none**

Note that for console access, secondary authentication must be `local` if primary access is not `local`. This prevents being locked out of the switch in the event of a failure in other access methods.

Configuring the primary password authentication method for port-access, MAC-based, and web-based access

Syntax:

```
aaa authentication <port-access <local|eap-radius|chap-radius>|<mac-based|web-
based|<chap-radius|peap-mschapv2> [none|authorized|server-group|<group-name>]>>
```

Configures the primary authentication method for port-access, MAC-based, or web-based access.

```
mac-based|web-based <chap-radius|peap-mschapv2>
```

Password authentication for web-based or MAC-based port access to the switch. Use `peap-mschapv2` for password verification without requiring access to a plain text password; it is more secure.

Default: `chap-radius`

```
port-access <local|eap-radius|chap-radius>
```

Configures `local`, `chap-radius (MD5)`, or `eap-radius` as the primary password authentication method for port-access.

Default primary authentication: `local`.

```
[none|authorized|server-group <group-name>]
```

`none`: No backup authentication method is used.

`authorized`: Allow access without authentication

`server-group <group-name>`: Specifies the server group to use with RADIUS.

Syntax:

```
aaa accounting <exec|network|system|commands| <start-stop|stop-only> radius [server-group <group-name>]
```

Configures accounting type and how data will be sent to the RADIUS server.

radius

Uses RADIUS protocol as accounting method.

server-group <group-name>

Specifies the server group to use with RADIUS.

Viewing RADIUS server group information

Syntax:

```
show server-group radius
```

Displays the same information as the `show radius` command, but displays the servers in their server groups.

Output from the show server-group radius command

```
HP Switch(config)# show server-group radius

 Status and Counters - AAA Server Groups

  Group Name: radius

               Auth  Acct  DM/  Time
  Server IP Addr Port  Port  CoA  Window  Encryption Key
  -------------- ----- ----- ---- ------- --------------------
  192.168.1.3    1812  1813  No   300     default_key
  192.168.3.3    1812  1813  No   300     grp2_key
  192.172.4.5    1812  1813  No   300     grp2_key
  192.173.6.7    1812  1813  No   300     grp2_key

Group Name: group2
               Auth  Acct  DM/  Time
  Server IP Addr Port  Port  CoA  Window  Encryption Key
  -------------- ----- ----- ---- ------- --------------------
  192.168.3.3    1812  1813  No   300     grp2_key
  192.172.4.5    1812  1813  No   300     grp2_key
  192.173.6.7    1812  1813  No   300     grp2_key

Group Name: group3
               Auth  Acct  DM/  Time
  Server IP Addr Port  Port  CoA  Window  Encryption Key
  -------------- ----- ----- ---- ------- --------------------
  192.168.30.3   1812  1813  No   300     grp3_key
  192.172.40.5   1812  1813  No   300     grp3_key
  192.173.60.7   1812  1813  No   300     grp3_key
```

Output from the show authentication command

```
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled
```

```
                | Login       Login         Login
Access Task  | Primary     Server Group  Secondary
-----------  + ---------- ------------ ----------
Console      | Local       radius        None
Telnet       | Local       radius        None
Port-Access  | Local                     None
Webui        | Local                     None
SSH          | Local                     None
Web-Auth     | ChapRadius  group3        None
MAC-Auth     | ChapRadius  group3        None


                | Enable      Enable        Enable
Access Task  | Primary     Server Group  Secondary
-----------  + ---------- ------------ ----------
Console      | Local       radius        None
Telnet       | Radius      group2        None
Webui        | Local                     None
SSH          | Local                     None
```

Output from the show accounting command

```
HP Switch(config)# show accounting

 Status and Counters - Accounting Information

  Interval(min) : 0
  Suppress Empty User : No
  Sessions Identification : Unique

  Type      | Method Mode        Server Group
  --------  + ------ ---------- ------------
  Network   | None
  Exec      | Radius Start-Stop group2
  System    | Radius Stop-Only  group2
  Commands  | Radius Start-Stop radius
```

Using SNMP to view and configure switch authentication features

SNMP MIB object access is available for switch authentication configuration (`hpSwitchAuth`) features. This means that the switches covered by this guide allow, by default, manager-only SNMP read/write access to a subset of the authentication MIB objects for the following features:

- number of primary and secondary login and enable attempts

- TACACS+ server configuration and status

- RADIUS server configuration

- selected 802.1X settings

- key management subsystem chain configuration

- key management subsystem key configuration

- local switch operator and manager usernames and passwords

With SNMP access to the `hpSwitchAuth` MIB enabled, a device with management access to the switch can view the configuration for the authentication features listed above (excluding usernames, passwords, and keys).

Using SNMP sets, a management device can change the authentication configuration (including changes to usernames, passwords and keys). operator read/write access to the authentication MIB is always denied.

> **NOTE**
>
> Regarding Security: All usernames, passwords, and keys configured in the `hpSwitchAuth` MIB are not returned through SNMP, and the response to SNMP queries for such information is a null string. However, SNMP sets can be used to configure username, password, and key MIB objects.
>
> To help prevent unauthorized access to the switch authentication MIB, Hewlett Packard Enterprise recommends following the reviewing **Viewing and changing the SNMP access configuration** on page 211.
>
> If you do not want to use SNMP access to the switch authentication configuration MIB, then use the `snmp-server mib hpswitchauthmib excluded` command to disable this access, as described in the next section.
>
> If you choose to leave SNMP access to the security MIB open (the default setting), Hewlett Packard Enterprise recommends that you configure the switch with the SNMP version 3 management and access security feature, and disable SNMP version 2c access. See "**SNMP access to the authentication configuration MIB** on page 41."

Viewing and changing the SNMP access configuration

Syntax:

```
snmp-server mib hpswitchauthmib <excluded|included>
```

```
included
```

Enables manager-level SNMP read/write access to the switch authentication configuration (`hpSwitchAuth`) MIB.

```
excluded
```

Disables manager-level SNMP read/write access to the switch authentication configuration (`hpSwitchAuth`) MIB.

Default: included

Syntax:

```
show snmp-server
```

The output for this command has been enhanced to display the current access status of the switch authentication configuration MIB in the `Excluded MIBs` field.

Example:

To disable SNMP access to the switch authentication MIB and then display the result in the Excluded MIB field, execute the following two commands.

Figure 38: *Disabling SNMP access to the authentication MIB and displaying the result*

```
HP Switch(config)# snmp-server mib hpswitchauthmib excluded
HP Switch(config)# show snmp-server
                                              Indicates that SNMP security MIB access
                                              is disabled, which is the nondefault
 SNMP Communities                             setting.

  Community Name       MIB View Write Access
  -------------------- -------- ------------
  public               Manager  Unrestricted
  test                 Manager  Unrestricted

 Trap Receivers

  Link-Change Traps Enabled on Ports [All] : All

   Traps Category                  Current Status
  _____          _____
  SNMP Authentication           : Extended
  Password change               : Enabled
  Login failures                : Enabled
  Port-Security                 : Enabled
  Authorization Server Contact  : Enabled
  DHCP-Snooping                 : Enabled
  Dynamic ARP Protection        : Enabled
  Dynamic IP Lockdown           : Enabled

  Address                Community               Events    Type    Retry   Timeout
  ---------------------- ----------------------- --------- ------- ------- -------
  15.255.131.57          public                  None      trap    3       15
  192.169.1.106          public                  None      trap    3       15
  15.255.135.68          public                  None      trap    3       15
  15.255.135.235         public                  None      trap    3       15
  15.255.135.222         public                  None      trap    3       15


 Excluded MIBs              This command disables
                            SNMP security MIB
  hpSwitchAuthenticationMIB  access.


 Snmp Response Pdu Source-IP Information

  Selection Policy   : rfc1517

 Trap Pdu Source-IP Information

  Selection Policy   : rfc1517
```

An alternate method of determining the current Authentication MIB access state is to use the `show run` command.

Figure 39: *Using the show run command to view the current authentication MIB access state*

```
HP Switch(config)# show run

Running configuration:

; J9091A Configuration Editor; Created on release #K.15.01.000x

hostname "HP Switch"
ip default-gateway 10.10.24.55
vlan 1
    name "DEFAULT_VLAN"
    untagged A1-A24,B1-B4
    ip address 10.10.24.100 255.255.255.0
    exit
snmp-server community "public" Operator
snmp-server mib hpSwitchAuthMIB excluded          ◄────  Indicates that SNMP access to the
password manager                                          authentication configuration MIB
                                                          (hpSwitchAuth) is disabled.
```

Local authentication process (RADIUS)

When the switch is configured to use RADIUS, it reverts to local authentication only if one of these two conditions exists:

- `Local`

  is the authentication option for the access method being used.

- The switch has been configured to query one or more RADIUS servers for a primary authentication request, but has not received a response, and `Local` is the configured secondary option.

For local authentication, the switch uses the operator-level and manager-level username/password set(s) previously configured locally on the switch. These are the usernames and passwords you configure using the CLI password command, the WebAgent, or the menu interface which enables only local password configuration.

- If the operator at the requesting terminal correctly enters the username/password pair for either access level (operator or manager), access is granted on the basis of which username/password pair was used. For example, suppose you configure Telnet primary access for RADIUS and Telnet secondary access for local. If a RADIUS access attempt fails, then you can still get access to either the operator or manager level of the switch by entering the correct username/password pair for the level you want to enter.

- If the username/password pair entered at the requesting terminal does not match either local username/ password pair previously configured in the switch, access is denied. In this case, the terminal is again prompted to enter a username/password pair. In the default configuration, the switch allows up to three attempts. If the requesting terminal exhausts the attempt limit without a successful authentication, the login session is terminated and the operator at the requesting terminal must initiate a new session before trying again.

Controlling WebAgent access

To help prevent unauthorized access through the WebAgent, do one or more of the following:

- Configure the switch to support RADIUS authentication for WebAgent access.

◦ Configure local authentication (a manager user name and password and, optionally, an operator user name and password) on the switch.

◦ Configure the switch's Authorized IP manager feature to allow WebAgent access only from authorized management stations. (The Authorized IP manager feature does not interfere with TACACS+ operation.)

◦ Use one of the following methods to disable WebAgent access to the switch via http (Port 80):

    – CLI: `no web-management`

    – Menu Interface—From the Main menu, select the following:

        2: Switch Configuration

        1: System Information

        WebAgent Enabled: No

Commands authorization

The RADIUS protocol combines user authentication and authorization steps into one phase. The user must be successfully authenticated before the RADIUS server will send authorization information from the user's profile to the Network Access Server (NAS). After user authentication has occurred, the authorization information provided by the RADIUS server is stored on the NAS for the duration of the user's session. Changes in the user's authorization profile during this time will not be effective until after the next authentication occurs.

You can limit the services for a user by enabling AAA RADIUS authorization. The NAS uses the information set up on the RADIUS server to control the user's access to CLI commands.

The authorization type implemented on the switches is the "commands" method. This method explicitly specifies on the RADIUS server which commands are allowed on the client device for authenticated users. This is done on a per-user or per-group basis.

NOTE The commands authorization will only be executed for commands entered from Telnet, SSH, or console sessions. The Web management interface is not supported.

By default, all users may execute a minimal set of commands regardless of their authorization status, for example, "exit" and "logout". This minimal set of commands can prevent deadlock on the switch due to an error in the user's authorization profile on the RADIUS server.

Enabling authorization

To control access to the CLI commands, enter this command at the CLI.

Syntax:

```
[no] aaa authorization <commands> <local|radius|none>
```

Configures authorization for controlling access to CLI commands. When enabled, the switch checks the list of commands supplied by the RADIUS server during user authentication to determine if a command entered by the user can be executed.

`local`: Use local groups as the authorization method.

`radius`: The NAS requests authorization information from the RADIUS server. Authorization rights are assigned by user or group.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

`none`: The NAS does not request authorization information.

For example, to enable the RADIUS protocol as the authorization method:

`HP Switch(config)# aaa authorization commands radius`

When the NAS sends the RADIUS server a valid username and password, the RADIUS server sends an Access-Accept packet that contains two attributes the command list and the command exception flag. When an authenticated user enters a command on the switch, the switch examines the list of commands delivered in the RADIUS Access-Accept packet as well as the command exception flag, which indicates whether the user has permission to execute the commands in the list. See **Configuring commands authorization on a RADIUS server** on page 215.

After the Access-Accept packet is delivered, the command list resides on the switch. Any changes to the user's command list on the RADIUS server are not seen until the user is authenticated again.

Viewing authorization information

Syntax:

`show authorization`

Configures authorization for controlling access to CLI commands. When enabled, the switch checks the list of commands supplied by the RADIUS server during user authentication to determine if a command entered by the user can be executed.

The show authorization command

```
HP Switch(config)# show authorization

 Status and Counters - Authorization Information

  Type     | Method
  -------- + ------
  Commands | RADIUS
```

Configuring commands authorization on a RADIUS server

Using vendor specific attributes (VSAs)

Some RADIUS-based features implemented on switches use VSAs for information exchange with the RADIUS server. RADIUS Access-Accept packets sent to the switch may contain the vendor-specific information.

The attributes supported with `commands` authorization are:

- HP-Command-String: List of commands (regular expressions) that are permitted (or denied) execution by the user. The commands are delimited by semi-colons and must be between 1 and 249 characters in length. Multiple instances of this attribute may be present in Access-Accept packets. (A single instance may be present in Accounting-Request packets.)

- HP-Command-Exception: A flag that specifies whether the commands indicated by the HP-Command-String attribute are permitted or denied to the user. A zero (0) means permit all listed commands and deny all others; a one (1) means deny all listed commands and permit all others.

The results of using the HP-Command-String and HP-Command-Exception attributes in various combinations are shown below.

Table 15: *HP command string and exception*

HP-command-string	HP-command-exception	Description
Not present	Not present	If command authorization is enabled and the RADIUS server does not provide any authorization attributes in an Access-Accept packet, the user is denied access to the server. This message appears: "Access denied: no user's authorization info supplied by the RADIUS server."
Not present	DenyList-PermitOthers(1)	Authenticated user is allowed to execute all commands available on the switch.
Not present	PermitList-DenyOthers(0)	Authenticated user can only execute a minimal set of commands (those that are available by default to any user).
Commands List	DenyList-PermitOthers(1)	Authenticated user may execute all commands except those in the Commands list.
Commands List	PermitList-DenyOthers(0)	Authenticated user can execute only those commands provided in the Commands List, plus the default commands.
Commands List	Not present	Authenticated user can only execute commands from the Commands List, plus the default commands.
Empty Commands List	Not present	Authenticate user can only execute a minimal set of commands (those that are available by default to any user).
Empty Commands List	DenyList-PermitOthers(1)	Authenticated user is allowed to execute all commands available on the switch.
Empty Commands List	PermitList-DenyOthers(0)	Authenticate user can only execute a minimal set of commands (those that are available by default to any user).

You must configure the RADIUS server to provide support for the VSAs. There are multiple RADIUS server applications; the two examples below show how a dictionary file can be created to define the VSAs for that RADIUS server application.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Example configuration using FreeRADIUS

Procedure

1. Create a dictionary file (for example, dictionary.hp) containing VSA definitions. An example file is:

   Figure 40: *Example dictionary file*

```
#
#   dictionary.hp
#
#   As posted to the list by User <user_email>
#
#   Version: $Id: dictionary.hp, v 1.0 2006/02/23 17:07:07
#

VENDOR          Hp      11

# HP Extensions


ATTRIBUTE       Hp-Command-String       2     string   Hp
ATTRIBUTE       Hp-Command-Exception    3     integer  Hp

# Hp-Command-Exception Attribute Values

VALUE           Hp-Command-Exception       Permit-List       0
VALUE           Hp-Command-Exception       Deny-List         1
```

2. Find the location of the dictionary files used by FreeRADIUS (try `/usr/local/share/freeradius`).

3. Copy dictionary.hp to that location. Open the existing dictionary file and add this entry:

   ```
   $ INCLUDE dictionary.hp
   ```

4. You can now use VSAs with other attributes when configuring user entries.

Dynamic port access auth via RADIUS

Overview

This feature allows a common port policy to be configured on all access ports by creating new RADIUS VSAs that will dynamically override the authentication limits. The changes are always applied to the port on the authenticator switch associated with the supplicant being authenticated.

> **NOTE**
> All the changes requested by the VSAs must be valid for the switch configuration. For example, if either MAC-based or Web-based port access is configured while 802.1X port access is in client mode, a RADIUS client with a VSA to change the 802.1X port access to port-based mode is not allowed. 802.1X in port-based mode is not allowed with MAC-based or web-based port access types. However, if the authenticating client has VSAs to disable MAC-based and Web-based authentication in conjunction with changing 802.1X to port-based mode, then client authentication is allowed.

Configuring the RADIUS VSAs

Only RADIUS-authenticated port-access clients will be able to dynamically change the port access settings using the new proprietary RADIUS VSAs. The settings that can be overridden are:

- Client limit (address limit with mac-based port access)

- Disabling the port-access types

- Setting the port mode in which 802.1X is operating

If the VSA client limit decreases the switch configured client limit, all clients except the client that is overriding the settings is deauthenticated. Only one client session at a time can override the port-access settings on a port. When the client session is deauthenticated, the port resets itself to the configured settings. This port reset causes the deauthentication of all clients for the port-access authentication types that had their settings changed dynamically.

The new VSAs are:

- **HP-Port-Client-Limit-Dot1x:** This VSA temporarily alters the 802.1X authentication client limit to the value contained in the VSA. Values range from 0 to 32 clients. A zero client limit means this VSA is disabled. This is an HPE proprietary VSA with a value of 10.

- **HP-Port-Client-Limit-MA:** This VSA temporarily alters the MAC authentication client limit to the value contained in the VSA. Values range from 0 to 256 clients. A zero client limit means this VSA is disabled. This is an HPE proprietary VSA with a value of 11.

- **HP-Port-Client-Limit-WA:** This VSA temporarily alters the web-based authentication client limit to the value contained in the VSA. Values range from 0 to 256 clients. A zero client limit means this VSA is disabled. This is an HPE proprietary VSA with a value of 12.

- **HP-Port-Auth-Mode-Dot1x:** This VSA temporarily alters the 802.1X authentication mode to be either port-based or user-based depending on the value in the VSA. A port-based VSA is set with a value of 1; a user-based VSA is set with a value of 2. This is an HPE proprietary VSA with a value of 13. If an 802.1X port is operating in port-based mode, it is invalid to set the 802.1X client limit using the HP-Port-Client-Limit VSA.

NOTE: The changing of the client limits for a port using VSAs is temporary. The running configuration file is not changed and still displays the client limit and address limit settings.

Each authentication type may have a unique value for the client limit. If the value of the VSA is zero, the authentication type corresponding to that VSA will be disabled.

Settings for these VSAs are in effect for the duration of the authenticated session of the downstream supplicant switch. If for any reason there is a loss of the session (link loss between authenticator switch and supplicant switch, or authentication failure during reauthentication), the originally configured 802.1X and MAC authentication limits are restored.

Viewing port-access information

The `show port-access summary` command displays the dynamically changed client limit settings.

Syntax:

```
show port-access summary [radius-overridden]
```

Displays summary configuration information for all ports, including the ports that have client limits set by RADIUS VSAs.

`radius-overridden`: Displays only the ports with client limits that are overridden by RADIUS attributes.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

If the command `no aaa port-access authentication <port-list> client-limit` is executed, the port access is in port-mode.

If the 802.1X client-limit is configured with a value from 1-32, the port access is in user-mode.

Summary configuration information

```
Port Access Status Summary

  Port-access authenticator activated [No] : No
  Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No
  Use LLDP data to authenticate [No] : No

       |    Authenticator      |   Web Auth   |      MAC Auth          |   Local MAC
  Port | Enable  Mode   Limit | Enable Limit | Enable Mode   Limit | Enable Limit
  ----- - ----- ----- ------ - ------ ----- - ------ ------ ----- - ------ ---
  1    | No      Port   0     | No      1    | No      Port  1     | No      1
  2    | No      Port   0     | No      1    | No      User  1     | No      1
  3    | No      Port   0     | No      1    | No      User  1     | No      1
  4    | No      Port   0     | No      1    | No      User  1     | No      1
  5    | No      Port   0     | No      1    | No      User  1     | No      1
  6    | No      Port   0     | No      1    | No      User  1     | No      1
```

To display the configuration information for just those ports that are dynamically overridden by RADIUS attributes, use the `show port-access summary radius-overridden` command.

Output for client-limit values that are RADIUS overridden

```
HP Switch(config)# show port-access summary radius-overridden

 Port Access Status Summary

  Port-access authenticator activated [No} : No
  Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

  Note: * indicates values dynamically overridden by RADIUS

          Authenticator        Web Auth        MAC Auth
  Port    Enabled Mode Limit    Enabled Limit    Enabled Limit
  ---- + ------- ---- ----- + ------- ----- + ------- -----
  1    | Yes     user* 1*    | Yes      1    | Yes      1
  2    | Yes     user  32    | Yes      32*  | Yes      32
  4    | No      port  1     | No       1    | No*      1
```

Operating notes

• Only RADIUS authentication supports the new VSAs. Other authentication types, such as TACACS, are not supported.

• If the RADIUS server delivers a new VSA to an authenticator switch that does not understand it, the Access-Accept message is accepted and the new VSA is ignored by the switch.

• If the RADIUS server delivers a new VSA to an authenticator switch that does not understand it, the Access-Accept message is rejected.

VLAN assignment in an authentication session

A switch supports concurrent 802.1X and either web-based or MAC authentication sessions on a port (with up to 32 clients allowed). If you have configured RADIUS as the primary authentication method for a type of access, when a client authenticates on a port, the RADIUS server assigns an untagged VLAN that is statically configured on the switch for use in the authentication session. See the documentation provided with the RADIUS server application.)

If a switch port is configured to accept multiple 802.1X and/or web-based or MAC authentication client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session. On a port where one or more authenticated client sessions are already running, all clients are on the same untagged VLAN.

If the RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail.

Tagged and untagged VLAN attributes

When you configure a user profile on a RADIUS server to assign a VLAN to an authenticated client, you can use either the VLAN name or VLAN ID (VID) number. For example, if a VLAN configured in the switch has a VID of 100 and is named `vlan100`, you could configure the RADIUS server to use either "100" or "vlan100" to specify the VLAN.

After the RADIUS server validates a client username and password, the RADIUS server returns an Access-Accept packet that contains the VLAN assignment and the following attributes for use in the authentication session:

- Egress-VLANID: Configures an optional, egress VLAN ID for either tagged or untagged packets (RFC 4675).

- Egress-VLAN-Name: Configures an optional, egress VLAN for either tagged or untagged packets when the VLAN ID is not known (RFC 4675).

- Tunnel-Type, Tunnel-Medium-Type, and Tunnel-Private-Group-ID: Tunnel attributes that specify an untagged VLAN assignment (RFC 3580).Tunnel (untagged VLAN) attributes may be included in the same RADIUS packet as the Egress-VLANID and Egress-VLAN-Name attributes. These attributes are not mutually exclusive.

The switch processes the VLAN information returned from the remote RADIUS server for each successfully 802.1X-, web-based, and MAC authenticated client (user). The VLAN information is part of the user profile stored in the RADIUS server database and is applied if the VLANs exist on the switch.

Additional RADIUS attributes

These attributes are included in Access-Request and Access-Accounting packets sent from the switch to the RADIUS server, advertising switch capabilities, reporting authentication session information, and dynamically reconfiguring authentication parameters:

- MS-RAS-Vendor (RFC 2548): Allows switches to inform a Microsoft RADIUS server that the switches are from Hewlett Packard Enterprise Networking. This feature assists the RADIUS server in its network configuration.

- HP-capability-advert: An HPE proprietary RADIUS attribute that allows a switch to advertise its current capabilities to the RADIUS server for port-based (MAC, Web, or 802.1X) authentication; for example, VSAs for port QoS, ingress rate-limiting, RFC 4675 QoS and VLAN attributes, and RFC 3580 VLAN-related attributes. The RADIUS server uses this information to make a more intelligent policy decision on the configuration settings to return to the switch for a client session.

- HP-acct-terminate-cause: An HPE proprietary RADIUS accounting attribute that allows a switch to report to the RADIUS server why an authentication session was terminated. This information allows customers to diagnose

network operational problems and generate reports on terminated sessions. This attribute provides extended information on the statistics provided by the acct-terminate-cause attribute.

- Change-of-Authorization (CoA) (RFC 3576: Dynamic Authorization Extensions to RADIUS): A mechanism that allows a RADIUS server to dynamically disconnect messages (DM) or change the authorization parameters (such as VLAN assignment) used in an active client session on the switch. The switch (NAS) does not have to initiate the exchange. For example, for security reasons you may want to limit the network services granted to an authenticated user. In this case, you can change the user profile on the RADIUS server and have the new authorization settings take effect immediately in the active client session. The Change-of-Authorization attribute provides the mechanism to dynamically update an active client session with a new user policy that is sent in RADIUS packets.

Output for dynamic authorization configuration

```
HP Switch(config)# show radius dyn-authorization

 Status and Counters - RADIUS Dynamic Authorization Information

  NAS Identifier : LAB-8212
  Invalid Client Addresses (CoA-Reqs) : 0
  Invalid Client Addresses (Disc-Reqs) : 0

                 Disc      Disc      Disc      CoA       CoA       CoA
  Client IP Addr Reqs      ACKs      NAKs      Reqs      ACKs      NAKs
  -------------- --------  --------  --------  --------  --------  --------
  154.34.23.106  1         1         0         2         2         0
  154.45.234.12  2         1         1         3         3         0
```

Output showing dynamic authorization statistics

```
HP Switch(config)# show radius host 154.23.45.111 dyn-authorization
Status and Counters - RADIUS Dynamic Authorization Information

  Authorization Client IP Address : 154.23.45.111
  Unknown PKT Types Received : 0

  Disc-Reqs                  : 2      CoA-Reqs                  : 1
  Disc-Reqs Authorize Only : 0        CoA-Reqs Authorize Only : 0
  Disc-ACKs                  : 2      CoA-ACKs                  : 1
  Disc-NAKs                  : 0      CoA-NAKs                  : 0
  Disc-NAKs Authorize Only : 0        CoA-NAKs Authorize Only : 0
  Disc-NAKs No Ses. Found  : 0        CoA-NAKs No Ses. Found  : 0
  Disc-Reqs Ses. Removed   : 0        CoA-Reqs Ses. Changed   : 0
  Disc-Reqs Malformed      : 0        CoA-Reqs Malformed      : 0
  Disc-Reqs Bad Authentic. : 0        CoA-Reqs Bad Authentic. : 0
  Disc-Reqs Dropped        : 0        CoA-Reqs Dropped        : 0
```

MAC-based VLANs

MAC-Based VLANs (MBVs) allow multiple clients on a single switch port to receive different untagged VLAN assignments. VLAN assignment of untagged traffic is based on the source MAC address rather than the port. Clients receive their untagged VLAN assignment from the RADIUS server. This feature adheres to the requirement that if all known RADIUS attributes for a given client cannot be applied, the authentication request for that client must be rejected.

Both authenticated and unauthenticated clients can reside on the same port on different VLANs, but only if the mixed-mode configuration is enabled. This is not the default behavior. The normal operating behavior is not to allow unauthenticated clients on the port when at least one authenticated client is present on the port. If an

unauthenticated client is present on the unauthorized VLAN and another client successfully authenticates on that port, the unauthenticated client is kicked off the port.

When an MBV cannot be applied due to a conflict with another client on that port, a message indicating VID arbitration error is logged.

When an MBV cannot be applied due to lack of resources, a message indicating lack of resources is logged.

There is no command line support for this feature. The decision to use an MBV is made automatically if the hardware is capable and if the situation necessitates. If multiple clients authenticate on different untagged VLANs on hardware that does not support MBVs, the switch will reject all clients authorized on a VLAN different from the first client's VLAN - the first authenticated client sets the Port VID (PVID).

This feature has the side effect of allowing egress traffic from one client's VLAN to be accepted by all untagged clients on that port. For example, suppose that clients A and B are both on the same switch port, but on two different VLANs. If client A is subscribing to a multicast stream, then client B also receives that multicast traffic.

Accounting services

RADIUS accounting collects data about user activity and system events and sends it to a RADIUS server when specified events occur on the switch, such as a logoff or a reboot.

Accounting service types

The switch supports four types of accounting services:

- **Network accounting:**

  Provides records containing the information listed below on clients directly connected to the switch and operating under Port-Based Access Control (802.1X):

◦ Acct-Session-Id	◦ Acct-Output-Packets	◦ Service-Type
◦ Acct-Status-Type	◦ Acct-Input-Octets	◦ NAS-IP-Address
◦ Acct-Terminate-Cause	◦ Nas-Port	◦ NAS-Identifier
◦ Acct-Authentic	◦ Acct-Output-Octets	◦ Calling-Station-Id
◦ Acct-Delay-Time	◦ Acct-Session-Time	◦ HP-acct-terminate-cause
◦ Acct-Input-Packets	◦ User-Name	◦ MS-RAS-Vendor

- **Exec accounting:**

  Provides records holding the information listed below about login sessions (console, Telnet, and SSH) on the switch:

◦ Acct-Session-Id	◦ Acct-Delay-Time	◦ NAS-IP-Address
◦ Acct-Status-Type	◦ Acct-Session-Time	◦ NAS-Identifier
◦ Acct-Terminate-Cause	◦ User-Name	◦ Calling-Station-Id
◦ Acct-Authentic	◦ Service-Type	◦ MS-RAS-Vendor

- **System accounting:**

  Provides records containing the information listed below when system events occur on the switch, including system reset, system boot, and enabling or disabling of system accounting.

◦ Acct-Session-Id	◦ Acct-Delay-Time	◦ NAS-Identifier
◦ Acct-Status-Type	◦ Username	◦ Calling-Station-Id
◦ Acct-Terminate-Cause	◦ Service-Type	◦ Acct-Session-Time
◦ Acct-Authentic	◦ NAS-IP-Address	◦ MS-RAS-Vendor

- **Commands accounting:**

   Provides records containing information on CLI command execution during user sessions.

◦ Acct-Session-Id	◦ User-Name	◦ Calling-Station-Id
◦ Acct-Status-Type	◦ NAS-IP-Address	◦ HP-Command-String
◦ Service-Type	◦ NAS-Identifier	◦ Acct-Delay-Time
◦ Acct-Authentic	◦ NAS-Port-Type	

- **RADIUS accounting with IP attribute:**

   The RADIUS Attribute 8 (Framed-IP-Address) feature provides the RADIUS server with information about the client's IP address after the client is authenticated. DHCP snooping is queried for the IP address of the client, so DHCP snooping must be enabled for the VLAN of which the client is a member.

   When the switch begins communications with the RADIUS server it sends the IP address of the client requesting access to the RADIUS server as RADIUS Attribute 8 (Framed-IP-Address) in the RADIUS accounting request. The RADIUS server can use this information to build a map of usernames and addresses.

   It may take a minute or longer for the switch to learn the IP address and then send the accounting packet with the Framed-IP-Address attribute to the RADIUS server. If the switch does not learn the IP address after a minute, it sends the accounting request packet to the RADIUS server without the Framed-IP-Address attribute. If the IP address is learned at a later time, it will be included in the next accounting request packet sent.

The switch forwards the accounting information it collects to the designated RADIUS server, where the information is formatted, stored, and managed by the server. For more information on this aspect of RADIUS accounting, see the documentation provided with your RADIUS server.

Operating rules for RADIUS accounting

- You can configure up to four types of accounting to run simultaneously: exec, system, network, and command.

- RADIUS servers used for accounting are also used for authentication.

- The switch must be configured to access at least one RADIUS server.

- RADIUS servers are accessed in the order in which their IP addresses were configured in the switch. Use `show radius` to view the order. As long as the first server is accessible and responding to authentication requests from the switch, a second or third server will not be accessed. For more on this topic, see **Changing RADIUS-server access order** on page 237.

- If access to a RADIUS server fails during a session, but after the client has been authenticated the switch continues to assume the server is available to receive accounting data. Thus, if server access fails during a session, it will not receive accounting data transmitted from the switch.

Acct-Session-ID options in a management session

The switch can be configured to support either of the following options for the accounting service types used in a management session. (See **Accounting service types** on page 222.)

- Unique Acct-Session-ID for each accounting service type used in the same management session (the default)

- Same Acct-Session-ID for all accounting service types used in the same management session

Unique Acct-Session-ID operation

In the Unique mode (the default), the various service types running in a management session operate as parallel, independent processes. Thus, during a specific management session, a given service type has the same Acct-Session-ID for all accounting actions for that service type. However, the Acct- Session-ID for each service type differs from the ID for the other types.

NOTE

In Unique Acct-Session-ID operation, the Command service type is a special case in which the Acct-Session-ID for each executed CLI command in the session is different from the IDs for other service types used in the session and also different for each CLI command executed during the session. That is, the ID for each successive CLI command in the session is sequentially incremented from the ID value assigned to the immediately preceding CLI command in that session.

Figure 41: *Example of accounting in the (default) unique mode*

```
User "fred" starts Exec        Acct-Session-Id = "003300000008"
Accounting session              Acct-Status-Type = Start
"003300000008".                 Service-Type = NAS-Prompt-User
                                Acct-Authentic = RADIUS
                                NAS-IP-Address = 10.1.242.15
                                NAS-Identifier = "gsf_dosx_15"
                                User-Name = "fred"
                                Calling-Station-Id = "172.22.17.101"
                                Acct-Delay-Time = 0

                                Acct-Session-Id = "003300000009"
                                Acct-Status-Type = Stop
User "fred" then executes show   Service-Type = NAS-Prompt-User
ip, which results in this        Acct-Authentic = RADIUS
accounting entry. Notice the     User-Name = "fred"
session ID (003300000009)        NAS-IP-Address = 10.1.242.15
assigned to this accounting      NAS-Identifier = "gsf_dosx_15"
entry incrementally follows the  NAS-Port-Type = Virtual
preceeding Acct-Session-Id.      Calling-Station-Id = "172.22.17.101"
This incrementing of the         HP-Command-String = "show ip"
session ID is normal operation   Acct-Delay-Time = 0
for command accounting in the
(default) Unique mode.           Acct-Session-Id = "00330000000A"
                                Acct-Status-Type = Stop
                                Service-Type = NAS-Prompt-User
User "fred" executes the logout  Acct-Authentic = RADIUS
command. The session ID          User-Name = "fred"
(00330000000A) assigned to       NAS-IP-Address = 10.1.242.15
this accounting entry            NAS-Identifier = "gsf_dosx_15"
incrementally follows the        NAS-Port-Type = Virtual
preceeding Acct-Session-Id.      Calling-Station-Id = "172.22.17.101"
This is another instance of      HP-Command-String = "logout"
normal Command accounting        Acct-Delay-Time = 0
operation in the Unique mode.
                                Acct-Session-Id = "003300000008"
                                Acct-Status-Type = Stop
                                Service-Type = NAS-Prompt-User
                                Acct-Authentic = RADIUS
                                NAS-IP-Address = 10.1.242.15
Terminate Exec Accounting        NAS-Identifier = "gsf_dosx_15"
Session "003300000008"           User-Name = "fred"
                                Calling-Station-Id = "172.22.17.101"
                                Acct-Terminate-Cause = User-Request
                                Acct-Session-Time = 29
                                Acct-Delay-Time = 0
```

Common Acct-Session-ID operation

In this case, all service types running in a given management session operate as subprocesses of the same parent process, and the same Acct-Session-ID is used for accounting of all service types, including successive CLI commands.

Figure 42: *Example of accounting in common mode (with same session ID throughout)*

```
User "fred" starts Exec          Acct-Session-Id = "00330000000B"
Accounting session               Acct-Status-Type = Start
"00330000000B".                  Service-Type = NAS-Prompt-User
                                 Acct-Authentic = RADIUS
                                 NAS-IP-Address = 10.1.242.15
                                 NAS-Identifier = "gsf_dosx_15"
                                 User-Name = "fred"
                                 Calling-Station-Id = "172.22.17.101"
                                 Acct-Delay-Time = 0

User "fred" then executes        Acct-Session-Id = "00330000000B"
show ip, which results in this   Acct-Status-Type = Stop
command accounting entry.        Service-Type = NAS-Prompt-User
Because this example assumes     Acct-Authentic = RADIUS
Common Mode configuration,       User-Name = "fred"
the session ID (00330000000B)    NAS-IP-Address = 10.1.242.15
assigned to this accounting      NAS-Identifier = "gsf_dosx_15"
entry is identical to the session NAS-Port-Type = Virtual
ID assigned when the session     Calling-Station-Id = "172.22.17.101"
was opened. No incrementing      HP-Command-String = "show ip"
of the session ID is done for    Acct-Delay-Time = 0
individual commands.

User "fred" executes the logout  Acct-Session-Id = "00330000000B"
command. The session ID          Acct-Status-Type = Stop
(00330000000B) used for the      Service-Type = NAS-Prompt-User
earlier Exec and Command         Acct-Authentic = RADIUS
accounting entries continues to  User-Name = "fred"
be the same as was originally    NAS-IP-Address = 10.1.242.15
assigned to the session.         NAS-Identifier = "gsf_dosx_15"
                                 NAS-Port-Type = Virtual
                                 Calling-Station-Id = "172.22.17.101"
                                 HP-Command-String = "logout"
                                 Acct-Delay-Time = 0

Terminate Exec Accounting        Acct-Session-Id = "00330000000B"
Session "00330000000B"           Acct-Status-Type = Stop
                                 Service-Type = NAS-Prompt-User
                                 Acct-Authentic = RADIUS
                                 NAS-IP-Address = 10.1.242.15
                                 NAS-Identifier = "gsf_dosx_15"
                                 User-Name = "fred"
                                 Calling-Station-Id = "172.22.17.101"
                                 Acct-Terminate-Cause = User-Request
                                 Acct-Session-Time = 29
                                 Acct-Delay-Time = 0
```

Configuring RADIUS accounting

NOTE

This procedure assumes:

- RADIUS authentication is configured on the switch for one or more access methods

- One or more RADIUS servers is configured to support the switch

Steps for configuring RADIUS accounting

Procedure

1. Configure the switch for accessing a RADIUS server.

You can configure up to three RADIUS servers (one primary, two backup). The switch operates on the assumption that a server can operate in both accounting and authentication mode. See the documentation for your RADIUS server application for additional information.

- Use the same `radius-server host` command that you would use to configure RADIUS authentication.

- Provide the following:

  ◦ A RADIUS server IP address.

  ◦ Optional — UDP destination port for authentication requests. Otherwise the switch assigns the default UDP port (1812; recommended).

  ◦ Optional — if you are also configuring the switch for RADIUS authentication, and need a unique encryption key for use during authentication sessions with the RADIUS server you are designating, configure a server-specific key. This key overrides the global encryption key you can also configure on the switch, and must match the encryption key used on the specified RADIUS server. Default: null

2. (Optional) Reconfigure the desired Acct-Session-ID operation.

   a. **Unique (the default setting):** Establishes a different Acct-Session-ID value for each service type, and incrementing of this ID per CLI command for the Command service type.

   b. **Common:** Establishes the same Acct-Session-ID value for all service types, including successive CLI commands in the same management session.

3. Configure accounting types and the controls for sending reports to the RADIUS server.

   a. Accounting types:

      - exec

      - network

      - system

      - commands

   b. Trigger for sending accounting reports to a RADIUS server: At session start and stop or only at session stop.

4. (Optional) Configure session blocking and interim updating options.

   a. **Updating:** Periodically update the accounting data for sessions-in-progress.

   b. **Suppress accounting:** Block the accounting session for any unknown user with no username trying to access to the switch.

Configuring a switch to access a RADIUS server

Before you configure the actual accounting parameters, configure the switch to use a RADIUS server. This process is outlined in **Configuring the switch to access a RADIUS server** on page 201. Repeat this now only if one of the following applies:

- The switch is not yet configured to use a RADIUS server

- Your server data has changed

- You need to specify a non-default UDP destination port for accounting requests

NOTE   Switch operation expects a RADIUS server to accommodate both authentication and accounting.

Syntax:

`[no] radius-server host <ip-address>`

Adds a server to the RADIUS configuration or (with no) deletes a server from the configuration.

`[acct-port < port-number >]`

Optional. Changes the UDP destination port for accounting requests to the specified RADIUS server. If you do not use this option, the switch automatically assigns the default accounting port number. (Default: 1813)

`[key < key-string >]`

Optional. Specifies an encryption key for use during accounting or authentication sessions with the specified server. This key must match the encryption key used on the RADIUS server. Use this command only if the specified server requires a different encryption key than configured for the global encryption key.

NOTE   If you save the config file using Xmodem or TFTP, the key information is not saved in the file. This causes RADIUS authentication to fail when the config file is loaded back onto the switch.

`[encrypted-key < key-string >]`

Encryption key to use with the RADIUS server, specified using a base64-encoded aes-256 encrypted string.

Example:

Suppose you want the switch to use the RADIUS server described below for both authentication and accounting purposes.

- IP address: 10.33.18.151

- A non-default UDP port number of 1750 for accounting.

For this example, assume that all other RADIUS authentication parameters for accessing this server are acceptable at their default settings, and RADIUS is already configured as an authentication method for one or more types of access to the switch (Telnet, Console, etc.).

Because the `radius-server` command includes an `acct-port` keyword with a non-default UDP port number of 1750, the switch assigns this value as the UDP accounting port.

Configuring for a RADIUS server with a non-default accounting UDP port number

The radius-server command configures the switch to use a RADIUS server at IP address 10.33.18.151, with a non-default UDP accounting port of 1750, and a server-specific key of "source0151".

```
HP Switch(config)# radius-server host 10.33.18.151
 acct-port 1750 key source0151
HP Switch(config)# write mem

HP Switch(config)# show radius

 Status and Counters - General RADIUS Information

  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key :
  Dynamic Authorization UDP Port : 3799


                Auth Acct DM/ Time
  Server IP Addr  Port Port CoA Window Encryption Key
  --------------- ---- ---- --- ------ --------------------
  10.33.18.151    1812 1750 No  10     source0151
```

Reconfiguring the Acct-Session-ID operation (optional)

Syntax:

```
aaa accounting session-id <unique|common>
```

Optional command to reconfigure the Acct-Session-ID mode to apply to the accounting service type records for a given management session.

`unique`: Configures the switch to use a different Acct-Session-ID for each accounting service type. (Default setting)

`common`: Configures the switch to apply the same Acct-Session-ID to all accounting service types in the same management session.

Figure 43: *Accounting configured for the common option*

```
HP Switch(config)# aaa accounting session-id common
HP Switch(config)# show accounting

 Status and Counters - Accounting Information

  Interval(min) : 0
  Suppress Empty User : No
  Sessions Identification : Common
                                          Example of common
                                          Session ID Configuration
  Type      | Method Mode             Server
  --------- + ------ -------------- ------------
  Network   | None
  Exec      | None
  System    | None
  Commands  | None
```

Configure accounting types and controls for sending reports to the RADIUS server
Accounting service types to track

Configure one or more accounting service types to track:

- **Exec:** Use `exec` if you want to collect accounting information on login sessions on the switch via the console, Telnet, or SSH.

- **System:** Use `system` if you want to collect accounting data when:

◦ A system boot or reload occurs

◦ System accounting is turned on or off

> **NOTE**
> There is no time span associated with using the `system` option. It simply causes the switch to transmit whatever accounting data it currently has when one of the above events occurs.

- **Network:** Use `network` if you want to collect accounting information on 802.1X port-based-access to the network by users connected to the physical ports on the switch.

- **Commands:** When commands accounting is enabled, an accounting notice record is sent after the execution of each command.

Accounting controls

These options are enabled separately, and define how the switch will send accounting data to a RADIUS server:

- **Start-Stop:** Applies to the `exec`, `network`, and `system` accounting service types:

  ◦ Send a "start record accounting" notice at the beginning of the accounting session and a "stop record notice" at the end of the session. Both notices include the latest data the switch has collected for the requested accounting type.

  ◦ Do not wait for an acknowledgment.

- **Stop-Only:** Applies to the `network`, `exec`, `system`, and `command` service types, as described below:

  ◦ Send a stop record accounting notice at the end of the accounting session. The notice includes the latest data the switch has collected for the requested accounting type (`network`, `exec`, or `system` service types). For the `commands` service type, sends the "Stop" accounting notice after execution of each CLI command.

  ◦ Do not wait for an acknowledgment.The system option always delivers `stop-only` operation because the switch sends the accumulated data only when there is a reboot, reload, or accounting on/off event.

- **Interim-Update:** Applies only to the `command` service type, and is intended for use when the optional `common` session ID is configured. Enabling `interim-update` in this case results in the command accounting records appearing as enclosed sub-parts of the `exec` service type record for a given management session. Using interim-update when the `unique` session ID is configured has no effect because in this case, the different service types appear as separate accounting processes with separate Acct-Session-ID values.

> **NOTE**
> Configuring `interim-update` for Command accounting results in all commands being reported as "update" records, regardless of whether common or unique is configured for the accounting session ID.

Syntax:

```
[no] aaa accounting <exec|network|system> <start-stop|stop-only> <radius|syslog>
```

```
[no] aaa accounting command <stop-only|interim-only> <radius|syslog>
```

Configures RADIUS accounting service type and how data will be sent to the RADIUS server.

```
<exec|network|system|command>
```

Specifies an accounting service type to configure. See **Accounting service types** on page 222.

`start-stop`: Applies to exec, network, and system accounting service types.

`stop-only`: Applies to all accounting service types.

`radius`: Uses RADIUS as the accounting period.

`syslog`: Uses syslog as the accounting protocol.

`interim-update`: Applies to the commands accounting service type.

Example:

To configure RADIUS accounting on the switch with `start-stop` for Exec functions, `stop-only` for system functions, and `interim-update` for `commands` functions.

Figure 44: *Example of configuring accounting types and controls*

```
HP Switch(config)# aaa accounting exec start-stop radius
HP Switch(config)# aaa accounting system stop-only radius
HP Switch(config)# aaa accounting commands interim-update radius
HP Switch(config)# show accounting

 Status and Counters - Accounting Information

  Interval(min) : 0
  Suppress Empty User : No
  Sessions Identification : Common

  Type      | Method Mode             Server Group
  -------- + ------ --------------- -------------
  Network   | None
  Exec      | Radius Start-Stop
  System    | Radius Stop-Only
  Commands  | Radius Interim-Update
```

Common is configured to apply the same Acct-Session-ID to all accounting records for a given switch management session.

Exec, System, and Commands accounting are active. (Assumes the switch is configured to access a reachable RADIUS server.)

Example:

If the switch is configured with RADIUS accounting on the switch to use `start-stop` for Exec, System, and `Command` functions, there will be an "Accounting-On" record when the switch boots up and an "Accounting-Off" record when the switch reboots or reloads. (Assume that Acct-Session-Id is configured for `common`.)

Figure 45: *Example of accounting session operation with "start-stop" enabled*

```
Record of Switch Bootup              Acct-Session-Id = "003600000001"
                                     Acct-Status-Type = Accounting-
                            On
                                     NAS-IP-Address = 1.1.1.15
                                     NAS-Identifier = "gsf_dosx_15"
                                     Acct-Delay-Time = 5

Record of User Session Start         Acct-Session-Id =
                            "003600000002"
                                     Acct-Status-Type = Start
                                     Service-Type = NAS-Prompt-User
                                     Acct-Authentic = Local
                                     NAS-IP-Address = 10.1.242.15
                                     NAS-Identifier = "gsf_dosx_15"
                                     Calling-Station-Id = "0.0.0.0"
                                     Acct-Delay-Time = 0

Record of reload Command             Acct-Session-Id =
                            "003600000002"
                                     Acct-Status-Type = Interim-
                            Update
                                     Service-Type = NAS-Prompt-User
                                     Acct-Authentic = Local
                                     NAS-IP-Address = 10.1.242.15
                                     NAS-Identifier = "gsf_dosx_15"
                                     NAS-Port-Type = Virtual
                                     Calling-Station-Id = "0.0.0.0"
                                     HP-Command-String = "reload"
                                     Acct-Delay-Time = 0

Record of System Accounting          Acct-Session-Id =
Off When Switch Reboots     "003600000001"
                                     Acct-Status-Type = Accounting-
                            Off
```

Configuring session blocking and interim updating options (optional)

These optional parameters give you additional control over accounting data.

- **Updates:** In addition to using a `Start-Stop` or `Stop-Only` trigger, you can optionally configure the switch to send periodic accounting record updates to a RADIUS server.

- **Suppress:** The switch can suppress accounting for an unknown user having no user name.

Syntax:

`[no] aaa accounting update periodic <1-525600>`

Sets the accounting update period for all accounting sessions on the switch.

The `no` form disables the update function and resets the value to zero.

Default: zero; disabled

Syntax:

`[no] aaa accounting suppress null-username`

Disables accounting for unknown users having no username.

Default: suppression disabled

To continue the example in **Accounting controls**, suppose you want the switch to:

- Send updates every 10 minutes on in-progress accounting sessions.

- Block accounting for unknown users (no username).

Figure 46: *Example of optional accounting update period and accounting suppression on unknown user*

```
HP Switch(config)# aaa accounting update periodic 10
HP Switch(config)# aaa accounting suppress null-username
HP Switch(config)# show accounting
 Status and Counters - Accounting Information

  Interval(min) : 10
  Suppress Empty User : Yes              ────────────  Update Period
  Sessions Identification : Common       ────────────

  Type      | Method Mode            Server Group       Suppress Unknown User
  -------- + ------ -------------- ------------
  Network   | None
  Exec      | Radius Start-Stop
  System    | Radius Stop-Only
  Commands | Radius Interim-Update
```

Viewing RADIUS statistics

General RADIUS statistics

Syntax:

`show radius [host <ip-addr>]`

Shows general RADIUS configuration, including the server IP addresses. Optional form shows data for a specific RADIUS host. To use `show radius`, the server's IP address must be configured in the switch, which. requires prior use of the `radius-server host` command. See **Accounting services** on page 222 for more information.

General RADIUS information from show radius command

```
HP Switch# show radius

 Status and Counters - General RADIUS Information

  Deadtime(min) : 5
  Timeout(secs) : 10
  Retransmit Attempts : 2
  Global Encryption Key : myg10balkey
  Dynamic Authorization UDP Port : 3799
  Source IP Selection : Outgoing Interface


                Auth Acct DM/ Time
  Server IP Addr  Port Port CoA Window Encryption Key  OOBM
  --------------- ---- ---- --- ------ -------------- ----
  192.33.12.65    1812 1813 No  300    my65key        No
```

OOBM information is displayed only for switches that support OOBM.

RADIUS server information from the show radius host command

```
HP Switch(config)# show radius host 192.33.12.65

 Status and Counters - RADIUS Server Information

  Server IP Addr : 192.33.12.65
  Authentication UDP Port : 1812    Accounting UDP Port  : 1813
  Round Trip Time         : 2       Round Trip Time      : 7
  Pending Requests        : 0       Pending Requests     : 0
  Retransmissions         : 0       Retransmissions      : 0
  Timeouts                : 0       Timeouts             : 0
  Malformed Responses     : 0       Malformed Responses  : 0
  Bad Authenticators      : 0       Bad Authenticators   : 0
  Unknown Types           : 0       Unknown Types        : 0
  Packets Dropped         : 0       Packets Dropped      : 0
  Access Requests         : 2       Accounting Requests  : 2
  Access Challenges       : 0       Accounting Responses : 2
  Access Accepts          : 0
  Access Rejects          : 0
```

Table 16: *Values for show radius host output*

Term	Definition
Round Trip Time	The time interval between the most recent Accounting-Response and the Accounting-Request that matched it from this RADIUS accounting server.
Pending Requests	The number of RADIUS Accounting-Request packets sent to this server that have not yet timed out or received a response. This variable is incremented when an accounting-Request is sent and decremented due to receipt of an Accounting-Response, a timeout or a retransmission.
Retransmissions	The number of RADIUS Accounting-Request packets retransmitted to this RADIUS accounting server. Retransmissions include retries where the Identifier and Acct-Delay have been updated, as well as those in which they remain the same.
Timeouts	The number of accounting timeouts to this server. After a timeout the client may retry to the same server, send to a different server, or give up. A retry to the same server is counted as a retransmit as well as a timeout. A send to a different server is counted as an Accounting-Request as well as a timeout.
Malformed Responses	The number of malformed RADIUS Accounting-Response packets received from this server. Malformed packets include packets with an invalid length. Bad authenticators and unknown types are not included as malformed accounting responses.
Bad Authenticators	The number of RADIUS Accounting-Response packets which contained invalid authenticators received from this server.
Unknown Types	The number of RADIUS packets of unknown type which were received from this server on the accounting port.

Table Continued

Term	Definition
Packets Dropped	The number of RADIUS packets which were received from this server on the accounting port and dropped for some other reason.
Access Requests	The number of RADIUS Access-Requests the switch has sent since it was last rebooted. (Does not include retransmissions.)
Accounting Requests	The number of RADIUS Accounting-Request packets sent. This does not include retransmissions.
Access Challenges	The number of RADIUS Access-Challenge packets (valid or invalid) received from this server.
Access Accepts	The number of RADIUS Access-Accept packets (valid or invalid) received from this server.
Access Rejects	The number of RADIUS Access-Reject packets (valid or invalid) received from this server.
Responses	The number of RADIUS packets received on the accounting port from this server.

RADIUS authentication statistics

Syntax:

```
show authentication
```

Displays the primary and secondary authentication methods configured for the Console, Telnet, Port-Access (802.1X), and SSH methods of accessing the switch. Also displays the number of access attempts currently allowed in a session.

```
show radius authentication
```

Displays NAS identifier and data on the configured RADIUS server and the switch interactions with this server. Requires prior use of the `radius-server host` command to configure a RADIUS server IP address in the switch, see **Accounting services** on page 222.

Login attempt and primary/secondary authentication information from the show authentication command

```
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled

            | Login        Login         Login
  Access Task | Primary    Server Group  Secondary
  ----------- + ---------- ------------ ----------
  Console    | Local                     None
  Telnet     | Radius                    None
  Port-Access | Local                    None
  Webui      | Local                     None
  SSH        | Radius                    None
  Web-Auth   | ChapRadius  radius        None
```

```
    MAC-Auth      | ChapRadius   radius        None

                  | Enable       Enable        Enable
    Access Task   | Primary      Server Group  Secondary
    ----------- + ----------  ------------  ----------
    Console       | Local                     None
    Telnet        | Radius                    None
    Webui         | Local                     None
    SSH           | Radius                    None
```

RADIUS authentication information from a specific server

```
HP Switch(config)# show radius authentication

 Status and Counters - RADIUS Authentication Information

   NAS Identifier : Networking
   Invalid Server Addresses : 0

                 UDP
   Server IP Addr Port Timeouts  Requests  Challenges Accepts  Rejects
   -------------- ---- --------- --------- ---------- -------- -------
   192.33.12.65   1812 0         2         0          2        0
```

RADIUS accounting statistics

Syntax:

```
show accounting
```

Lists configured accounting interval, "Empty User" suppression status, session ID, accounting types, methods, and modes.

```
show radius accounting
```

Lists accounting statistics for one or more RADIUS servers configured in the switch (using the `radius-server host` command).

```
show accounting sessions
```

Lists the accounting sessions currently active on the switch.

Listing the accounting configuration in the switch

```
HP Switch(config)# show accounting

 Status and Counters - Accounting Information

   Interval(min) : 5
   Suppress Empty User : No
   Sessions Identification : Common

   Type      | Method Mode            Server Group
   -------- + ------ --------------  ------------
   Network   | None
   Exec      | Radius Start-Stop
   System    | Radius Stop-Only
   Commands  | Radius Interim-Update
```

RADIUS accounting information for a specific server

```
HP Switch(config)# show radius accounting

 Status and Counters - RADIUS Accounting Information

  NAS Identifier : Networking
  Invalid Server Addresses : 0

                 UDP
  Server IP Addr Port  Timeouts   Requests   Responses
  -------------- ----- ---------- ---------- ----------
  192.33.12.65   1813  0          1          1
```

Listing of active RADIUS accounting sessions on the switch

```
HP Switch(config)# show accounting sessions

Active Accounted actions on SWITCH, User (n/a) Priv (n/a),
  Acct-Session-Id 0x013E00000006, System Accounting record,
  1:45:34 Elapsed
  system event 'Accounting On
```

Changing RADIUS-server access order

The switch tries to access RADIUS servers according to the order in which their IP addresses are listed by the `show radius` command.

NOTE

When you add a server IP address, it is placed in the highest empty position in the list.

Adding or deleting a RADIUS server IP address leaves an empty position, but does not change the position of any other server addresses in the list. For example if you initially configure three server addresses, they are listed in the order in which you entered them. However, if you subsequently remove the second server address in the list and add a server address, the new address will be placed second in the list.

Thus, to move up a server address in the list, you must delete it from the list, ensure that the position to which you want to move it is vacant, and then re-enter it. For example, suppose that you have already configured the following three RADIUS server IP addresses in the switch:

Figure 47: *Search order for accessing a RADIUS server*

To exchange the positions of the addresses so that the server at 10.10.10.3 is the first choice and the server at 10.10.10.1 is the last, perform the following:

Procedure

1. Delete 10.10.10.3 from the list. This opens the third (lowest) position in the list.

2. Delete 10.10.10.1 from the list. This opens the first (highest) position in the list.

3. Re-enter 10.10.10.3. Because the switch places a newly entered address in the highest-available position, this address becomes first in the list.

4. Re-enter 10.10.10.1. Because the only position open is the third position, this address becomes last in the list.

Figure 48: *Example of new RADIUS server search order*

```
HP Switch(config)# no radius host 10.10.10.1
HP Switch(config)# radius host 10.10.10.3        Removes the "3" and "1" addresses from the
HP Switch(config)# radius host 10.10.10.1        RADIUS server list.
HP Switch# show radius

 Status and Counters - General RADIUS Information

  Deadtime(min) : 0                              Inserts the "3" address in the first position in the
  Timeout(secs) : 5                              RADIUS server list, and inserts the "1" address
  Retransmit Attempts : 3                        in the last position in the list.
  Global Encryption Key :
  Dynamic Authorization UDP Port : 3799
  Source IP Selection : Outgoing Interface       Shows the new order in which the switch
                                                 searches for a RADIUS server.

                 Auth Acct DM/ Time
  Server IP Addr  Port Port CoA Window Encryption Key                          OOBM
  --------------- ---- ---- --- ------ -------------------------------- ----
  10.10.10.3      1812 1813 No  300                                         No
  10.10.10.2      1812 1813 No  300                                         No
  10.10.10.1      1812 1813 No  300                                         No
```

Creating local privilege levels

This feature allows more granular localized control over user access when accessing the switch through the console or by telnet or SSH. Instead of allowing access to all commands with the "manager" command, or very restricted access with the "operator" command, the local access can be customized to allow the commands that the local account is authorized to execute. The new local accounts are in addition to and independent of the existing manager and operator accounts, with the exception that if a username is set for a manager or operator account, that name cannot be the same as any of the local user account names.

To do this, groups are created that contain up to 16 user accounts. The group has a list of match commands that determine if that user is authorized to execute that command. Up to 100 local user accounts are supported. The local user accounts are stored in the configuration as an SHA1 hash, which is only displayed if "include-credentials" is enabled. A password is required for the local user accounts, but nothing else.

There is one default group—operator. Users assigned to the operator group have only operator privileges.

Applying the authorization group to a local user account only occurs if the user logs in using `local` as the primary authentication method and the `aaa authorization commands local` command has been executed. Authorization groups are not supported when the login method is set as secondary local authentication.

These commands are authorized at all access levels:

- exit

- logout

- page

- redo

- repeat

- end

Configuring groups for local authorization

You must create a group for local authorization before you can assign local users to it. When creating the group, at least one command is created as part of that group. Typically, multiple commands are assigned to a group.

> **NOTE:** You must enable local authorization by executing `aaa authorization` commands `local` to use this feature.

To create a group, enter this command:

Syntax:

```
[no] aaa authorization group <group-name> <1-2147483647> match-command <command-string> <permit|deny> [log]
```

Create a local authorization group with the specified name. The name is case-sensitive and may not contain spaces. Duplicate names are not allowed. You can create a maximum of 16 groups. The name of the group can have a maximum of 16 characters.

```
<1-2147483647>
```

The evaluation order for the match commands.

```
match-command <command-string>
```

The `<command-string>` is the CLI command. It must be surrounded in double quotes of it contains any spaces, for example, "`vlan*`".

The `<command-string>` is a POSIX regular expression and follows POSIX matching rules. For example, the "*" character means match the preceding character zero or more times, so ab*c will match "ac", "abc", "abbc", etc. The "." character means match any character, so ".*" would match anything, while the command string "aaa.*" would match commands that have "aaa" followed by zero or more characters. The "^" character means match to the beginning of the string, so "^aaa.*" would mean the string must start with "aaa" and can have anything after that.

```
<permit|deny>
```

Either permit or deny execution of the command.

```
[log]
```

Optional. Indicates the matching of such commands will generate an event log entry for either permitted or denied.

Typically multiple commands are assigned to a group. Each command is entered on a separate line. Commands are evaluated in numerical order of the sequence number until a match is found, then the `permit` or `deny` action for that command is executed.

Commands are expanded before the comparison is done, for example, `sh ver` would be expanded to `show version` and then this command is compared against the command strings of the authorization group.**See: example**

When a command must be preceded by the execution of another command, then both commands need to be permitted for the command authorization group. For example, you must execute the `configure` command before you can enter the `vlan` context, so both commands must be permitted.

Configuring authorized commands for a group in the correct order

```
HP Switch(config)# aaa authorization group Redgroup 100
match-command configure permit
HP Switch(config)# aaa authorization group Redgroup 200
match-command "vlan *" permit
```

Some commands cause the switch CLI to enter a special context, such as test mode, and the input is not processed by the normal CLI. Keyboard input is not checked against the command authorization group. If these special contexts are permitted, the user can proceed outside the control and logging of the command group configuration.

Creating a local authorization group and assigning the commands authorized

```
HP Switch(config)# aaa authorization group Bluegroup 100
match-command configure permit
HP Switch(config)# aaa authorization group Bluegroup 200
match-command telnet permit
HP Switch(config)# aaa authorization group Bluegroup 300
match-command menu permit
```

Configuring a local user for a group

Local manager user logins and authorized command configuration are mutually exclusive with RADIUS or TACACS authentication and with RADIUS authorization and accounting.

To create a local user enter this command for the group with the appropriate authorizations.

Syntax:

```
[no] aaa authorization local-user <username> group <group-name> password <plaintext|sha1> <password>
```

Defines a local user for a defined group.

```
local-user <password>
```

The local user being added to the authorization group. The username can have a maximum of 16 characters. It must not contain spaces and is case-sensitive.

```
group <group-name>
```

The authorization group the local user belongs to. The group must have been created already.

```
password <plaintext|sha1 <password>
```

The plaintext password string can have a maximum of 16 characters. It must not contain spaces and is case-sensitive.

You are not allowed to actually enter the plaintext password in-line as part of the command; you will be prompted for it. The password is obscured when you enter it. The password is obscured when you enter it. This is similar to entering the password for the manager or operator.

If `include-credentials` is enabled, displaying the configuration shows the user passwords as SHA1 hash. If `include-credentials` is not enabled, then no password information is shown.

If a user is assigned to a command group and the group is subsequently deleted, the user will have operator privileges.

Creating a local user for a group

```
HP Switch(config)# aaa authentication local-user User1 group
Redgroup password plaintext
New password for User1: *******
```

Displaying command authorization information

To display information about users and command authorization for command groups, enter this command.

Syntax:

```
show authorization group [group-name]
```

Displays information about users and command authorization for command groups.

Specifying the group parameter without any group names displays information for all configured groups.

Showing command information for all groups

```
HP Switch(config)# show authorization group

 Local Management Groups - Authorization Information

  Group Name: Redgroup

  Username
  ----------------
  User1
  User2

  Sequence # | Permission Command Expression     Log
  ---------- + ---------- -------------------- -------
  100        | Permit     configure             Disable
  200        | Permit     vlan *                Disable

  Group Name: Bluegroup

  Username
  ----------------
  User3

  Sequence # | Permission Command Expression     Log
  ---------- + ---------- -------------------- -------
  100        | Permit     configure             Disable
  200        | Permit     telnet                Disable
  300        | Permit     menu                  Disable
```

Dynamic removal of authentication limits

In some situations, it is desirable to configure RADIUS attributes for downstream supplicant devices that allow dynamic removal of the 802.1X, MAC, and web-based authentication limits on the associated port of the authenticator switch. This eliminates the need to manually reconfigure ports associated with downstream 802.1X-capable devices, and MAC relay devices such as IP phones, on the authenticator switches. When the RADIUS

authentication ages out, the authentication limits are dynamically restored. This enhancement allows a common port policy to be configured on all access ports by creating new RADIUS vendor-specific attributes (VSAs) that will dynamically override the authentication limits. The changes are always applied to the port on the authenticator switch associated with the supplicant being authenticated.

> **NOTE**
>
> All the changes requested by the VSAs must be valid for the switch configuration. For example, if either MAC or web-based port access is configured while 802.1X port access is in client mode, a RADIUS client with a VSA to change the 802.1X port access to port-based mode is not allowed. 802.1X in port-based mode is not allowed with MAC or web-based port access types. However, if the authenticating client has VSAs to disable MAC and web-based authentication in conjunction with changing 802.1X to port-based mode, then client authentication is allowed.

Messages related to RADIUS operation

Message	Meaning
`Can't reach RADIUS server < x.x.x.x >.`	A designated RADIUS server is not responding to an authentication request. Try pinging the server to determine whether it is accessible to the switch. If the server is accessible, then verify that the switch is using the correct encryption key and that the server is correctly configured to receive an authentication request from the switch.
`No server(s) responding.`	The switch is configured for and attempting RADIUS authentication, however it is not receiving a response from a RADIUS server. Ensure that the switch is configured to access at least one RADIUS server. (Use `show radius`.) If you also see the message `Can't reach RADIUS server < x.x.x.x >` , try the suggestions listed for that message.
`Not legal combination of authentication methods.`	Indicates an attempt to configure `local` as both the primary and secondary authentication methods. If `local` is the primary method, then `none` must be the secondary method.

Security event log

The Joint Interoperability Test Command (JITC) is a United States military organization that tests technology that pertains to multiple branches of the armed services and government. The JITC requires that access to security logs be provided through security user authentication.

Security user log access

Security user logs are accessible when both the authentication and authorization are local. A default group called the default-security-group is available in manager mode and has the privileges to execute the commands copy security-log, show security-logging and clear security-logging. When a security user is attached to the group, they will only be able to execute the three commands. Any other user will not be able to execute the commands, no matter whether they are an operator or manager.

Creating a security user

Syntax

```
aaa authentication local-user user1 group default-security-group password plaintext
```

Security user commands

Syntax

```
copy security-log sftp | tfp | xmodem user IP-Address <FILENAME-STR>
```

Syntax

```
show security-logging
```

Syntax

```
clear security-logging
```

Authentication and Authorization through RADIUS

For RADIUS authentication and authorization, the security user will be able to access to security log by configuring the file located on RADIUS server.

Accessing the security log

```
/etc/raddb/users
steve Cleartext-Password := "testing"
Service-Type = Administrative-User,
HP-Command-Exception=0,
HP-Command-String="copy security log;show security-logging;clear security-logging"
```

Authentication and Authorization through TACACS+

For TACACS+ authentication and authorization, the user can access to security log by configuring the file located on TACACS+ server.

Security user access for TACACS authentication and authorization

```
/etc/tacacs/tac_plus.cfg
group = admin {
# default service = permit
service = exec {
priv-lvl = 15
}
cmd = copy security-log {

permit .*
}
cmd = copy show security-logging {
permit .*
}
cmd = clear security-logging {
permit .*
```

Restrictions

In the case of local authentication and authorization, the default-security-group group only applies to manager logins for CLI actions; menu interface and Web UI capabilities.

- There is no WebUI and Menu support for this feature.

- The same mechanism should be used for authentication and authorization when using this feature. Cross combination is not supported.

  For example: If authentication is local, then the authorization should also be set to local. Similarly, if authentication is RADIUS, authorization should also be set to RADIUS. TACACS authentication works in the same fashion.

- Security logging is supported via Syslog.

Event log wrap

To add a new log message when the event buffer wraps, a new RMON log is generated when the buffer wrapping is identified.

Configuring concurrent sessions

The following commands configure the max concurrent sessions allowable.

For non-stackable switches

Syntax

```
console max-sessions 1-6
```

The default value is 6. The [no] for the command restores the default value.

For non-stackable devices the allocations is as :

> 1 local console and 5 SSH/Telnet sessions.
> Minimum value of 1 is set since 1 local console is always available.

Configuring concurrent sessions per user

Configures the max concurrent sessions allowable per user. The [no] command restores the default value.

For non-stackable switches

Syntax

```
console max-sessions 1-6
```

The default value is 6.

Allocations

> 1 local console and 5 SSH/Telnet sessions.
> Minimum value of 1 is set since 1 local console is always available.

Configuring concurrent sessions per user

Configures the max concurrent sessions allowable per user in HPE switches. The [no] command restores the default value.

Syntax

```
console max-sessions 1-7
```

> **NOTE**
>
> The default value in non-stackable and HPE 5400R series switches is 6. The default value for stackable devices is 6.

Maximum session for a manager/operator

Configuring a value of 4 means that the Manager can have a maximum of 4 concurrent sessions and the Operator can also have a maximum of 4 concurrent sessions.

The allocations will be based on the current available free session (since the system as a whole can have only 6 to 7 concurrent sessions by default.

A session is allocated to a Manager or Operator only if the current allocation to the user is within the configured value. The session allocation thus happens based on the user sessions as well the max sessions available in the system.

Failed login attempts delay

Authentication happens by sending authentication request to the Authentication sub system. The Authentication sub system then authenticates and sends back the authentication result. If an authentication failure is identified, a delay is introduced so that the next authentication request is not serviced immediately. This is applicable for console, telnet and ssh sessions.

Features

- Enables Port ACL applications to be grouped.

- Allows end users to programmatically control grouping on a per port ACL basis.

- CLI support for new "shared" keyword.

- ACLs will be treated as unshared when upgrading to new release.

- Downgrading shared ACLs to previous release will not be supported, the ACLs will be disabled.

- Enables ACL applications to VLAN to be grouped.

- Grouped ACLs applied to multiple VLANs will use only a single TCAM resource.

- Allow users to control grouping on a per VLAN ACL basis.

Overview

ACL grouping is an extension of the ACL feature. Each ACL application will consume "n" TCAM resources therefore "x" applications of an ACL will use "x . *n" resources. ACL grouping allows for grouping by an ACL. With ACL grouping, the TCAM usage would shrink to "n". ACL grouping can be applied to both ports and VLANs.

ACL grouping provides the following capabilities:

- Enables Port ACL applications to be grouped.

- Allows end users to programmatically control grouping on a per port ACL basis.

- Provides CLI support for the `shared` keyword.

- ACLs will be treated as unshared when upgrading to new release.

- Downgrading shared ACLs to previous release will not be supported, the ACLs will be disabled.

- Enables ACL applications to VLAN to be grouped.

- Grouped ACLs applied to multiple VLANs will use only a single TCAM resource.

- Allow users to control grouping on a per VLAN ACL basis.

Commands

These commands allow uers to control explicit groupings on PACLs, VACLs and RACLs applications which allows for TCAM resource consolidation. Allow for better network troubleshooting via an individual port or VLAN when reviewing statistics specifically for that port or VLAN.

IPv4 access-group (PACL)

Allows for the configuration of an IPv4 ACL on a port to be shared.

Syntax

```
[no] ip access-group ACL-ID in|out shared
```

Description

Apply the specified IPv4 ACL to inbound or outbound packets on this interface. When ACLs are shared, hardware resource usage is optimized where possible.

Options

shared

> Apply the IPv4 ACL so as to share hardware resources.

Restrictions

- Per-application statistics will not be available when ACLs are applied as shared.

- Connection rate filter ACLs cannot be applied on this interface.

---

ip access-group my-acl out shared

```
switch(config)# int a1
switch(eth-a1)# ip access-group my-acl out shared
```

IPv6 access-group (PACL)

Allows for the configuration of an IPv6 ACL on a port to be shared.

Syntax

```
[no] ipv6 access-group ACL-ID in|out shared
```

Description

Apply the specified IPv6 ACL to inbound or outbound packets on this interface. When ACLs are shared, hardware resource usage is optimized where possible.

Options

shared

> Apply the IPv6 ACL so as to share hardware resources.

Restrictions

- Per-application statistics will not be available when ACLs are applied as shared.

- Connection rate filter ACLs cannot be applied on this interface.

---

ipv6 access-group my-acl out shared

```
switch(config)# int a1
switch(eth-a1)# ipv6 access-group my-acl out shared
```

MAC access-group (PACL)

Allows for the configuration of a MAC ACL on a port to be shared.

Syntax

```
mac-access-group ACL-ID in|out shared
```

Description

Apply the MAC ACL to the traffic on a port. MAC ACLs can be used to filter the traffic based on the source MAC address, destination MAC address, EtherType, CoS priority, or VLAN number. When ACLs are shared, hardware resource usage is optimized where possible.

Options

shared

Apply the MAC ACL so as to share hardware resources.

Restrictions

Per-application statistics will not be available when ACLs are applied as shared.

---

mac-access-group my-acl out shared

```
Switch(config)# int a1
switch(eth-a1)# mac-access-group my-acl out shared
```

IPv4 access-group (VACL)

Allows for the configuration of an IPv4 ACL on a vlan to be shared. VACLs are applied from vlan context.

Syntax

```
[no] ip access-group ACL-ID in|out|vlan-in|vlan-out|connection-rate-filter shared
```

Description

Apply the specified IPv4 ACL on this VLAN interface. When ACLs are shared, hardware resource usage is optimized where possible.

Options

shared

Apply the IPv4 ACL so as to share hardware resources.

Restrictions

Per-application statistics will not be available when ACLs are applied as shared.

---

ip access-group my-acl out shared

```
switch(config)# vlan 1
switch(vlan-1)# ip access-group my-acl vlan-out shared
switch(vlan-1)# ip access-group my-acl out shared
```

IPv6 access-group (VACL)

Allows for the configuration of an IPv6 ACL on a VLAN to be shared. VACLs are applied from VLAN context.

Syntax

```
[no] ipv6 access-group ACL-ID in|out|vlan-in|vlan-out|connection-rate-filter shared
```

Description

Apply the specified IPv6 ACL on this VLAN interface. When ACLs are shared, hardware resource usage is optimized where possible.

Options

shared

Apply the IPv6 ACL so as to share hardware resources.

Restrictions

Per-application statistics will not be available when ACLs are applied as shared.

---

ipv6 access-group my-acl out shared

```
switch(config)# vlan 1
switch(vlan-1)# ipv6 access-group my-acl vlan-out shared
switch(vlan-1)# ipv6 access-group my-acl out shared
```

MAC access-group (VACL)

Allows for the configuration of a MAC ACL on a VLAN to be shared.

Syntax

```
mac-access-group ACL-ID in|out shared
```

Description

Apply the MAC ACL to the traffic on a VLAN. MAC ACLs can be used to filter the traffic based on the source MAC address, destination MAC address, EtherType, CoS priority, or VLAN number. When ACLs are shared, hardware resource usage is optimized where possible.

Options

shared

Apply the MAC ACL so as to share hardware resources.

Restrictions

Per-application statistics will not be available when ACLs are applied as shared.

---

mac-access-group my-acl out shared

```
switch(config)# vlan 1
switch(vlan-1)# mac-access-group my-acl out shared
```

Modify existing commands

ACL grouping feature does not introduce new show commands. Although it does modify the output content to indicate if an ACL is shared on specified interface lists.

show configuration

Another method to see if the ACL is shared by using `show running-config`. Add shared keyword after direction when the ACL is shared.

Syntax

```
show running-config
```

Output

```
; Configuration Editor; Created on release #x.x.x.x
; Ver #08:6b.ff.f7.fc.7f.ff.3f.ef:c7


ip access-list extended "my-acl"
     10 permit tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
    exit

interface A1
   ip access-group "my-acl" out shared
   exit

interface A2
   ip access-group "my-acl" out shared
   exit

interface A3
   ip access-group "my-acl" out shared
   exit
```

show statistics

Add shared keyword after direction when the ACL is shared. When the ACL is shared the hit counts listed will be for the total hit counts of all the interface lists that ACL is applied to. To debug the specified interface list the user has to remove it from the shared ACL.

Syntax

```
show statistics aclv4 my-acl port a1 out
```

Example output

```
switch# show statistics aclv4 my-acl port a1 out
 Hit Counts for ACL my-acl shared
  Total
(       0 )      10 permit tcp 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
```

show access-list

The new feature will modify some of the existing output content in 'show access-list ports' and 'show access-list vlan'. Added shared keyword after type. Yes/No to indicate whether the interface list is shared.

show access-list ports

Syntax

```
show access-list ports PORT-LIST | all
```

Description

Show access-list ports by port list or all.

Options

all

   Show all ports.

Example

```
switch# show access-list ports all

Access Lists for Port A1
  IPv4 Outbound: my-acl Type: Extended Shared: Yes

Access Lists for Port A2
  IPv4 Outbound: my-acl Type: Extended Shared: Yes

Access Lists for Port A3
  IPv4 Outbound: my-acl Type: Extended Shared: Yes
```

show access-list vlan

Syntax

```
show access-list vlan VLAN-ID | all
```

Description

Show the access-list for a VLAN by VLAN-ID or all.

Example

```
switch# show access-list vlan 1

Access Lists for VLAN 1
  IPv4 Router Inbound           : (None)
  IPv4 Router Outbound          : (None)
  IPv4 VLAN Inbound             : (None)
  IPv4 VLAN Outbound             : my-acl   Type: Extended Shared: Yes
  IPv4 Connection Rate Filter   : (None)
  IPv6 Router Inbound           : (None)
  IPv6 Router Outbound          : (None)
  IPv6 VLAN Inbound             : (None)
  IPv6 VLAN Outbound            : (None)
```

Error messages

Event — Event Type	Message
The user tries to apply an ACL when there is no TCAM Banks available. This situation should not happen if there is already an ACL applied.Event Type — Warning	Unable to apply the ACL as shared. Hardware resources are unavailable.

Overview

NOTE MACsec is supported on the 2930M switch family. It is not supported on the 2930F switch family.

Media Access Control security (MACsec) is an IEEE 802 standard specifying how to secure all or part of a LAN at the link layer transparently. MACsec PHY devices can do this while meeting the scalability and high-speed requirements set on such networks. MACsec is intended for wired LANs only; wireless networks use a different protocol set. To ensure wired network security, the MACsec functionality is required on the newer generation of network infrastructure switches.

The MACsec protocol provides:

• Connectionless data integrity — (each MAC frame carries a separate integrity verification code, hence the term connectionless).

• Data origin authenticity—(an assurance that only authorized MACsec stations send MAC frames).

• Confidentiality — (each MAC frame is encrypted to prevent it from being eavesdropped).

• Replay protection — (MAC frames copied from the LAN by an attacker cannot be resent into the LAN without being detected).

MACsec secures switch to switch infrastructure using the MKA (MACsec Key Agreement) protocol and the Static CAK (Connectivity Association Key) Mode. MACsec operation includes:

• Switch-to-Switch Pairwise Pre-Shared CAK mode with Single-User (CAK) per port.

• A new MACsec-PHY for faster processing through hardware.

• Supports MACsec Key Agreement protocol (MKA) for automatic MACsec peer discovery, peer-participant liveliness, Key-Server election and for distribution of SAKs

• Supports AES-GCM-128 bit Key-length (CAKs/ICKs/KEKs/SAKs).

• Configuration includes "Integrity Check Only" and "Integrity Check with Confidentiality at offset 0" modes.

• Supports MACsec CLI configurations through CLI and SNMP and over Telnet/SSH. MACsec configuration through the web interface is not supported.

MACsec switch support

All 2930M ports support MACsec, with the exception of the JL078A 1-port QSFP+ 40GbE Module.

MACsec configuration commands

For supporting the MACsec configuration, configure the following:

- MACsec Policy creation and configuration

- Apply MACsec policy on ports

- Configure the MKA parameters on ports

Create, modify or delete a MACsec policy

Syntax

```
[no] macsec policy <policy-name>
```

Configures the MAC Security (MACsec) protocol.

macsec

MAC Security (MACsec).

policy

Apply a MACsec policy.

policy-name

MACsec policy name up to 32 characters long.

Validation rules

Table 17: *MACsec policy validation rules*

Validation	Error/Warning/Prompt
While creating or editing a policy:	
Maximum policy name length is 32	The policy name exceeds the maximum permissible limit of %s characters.(where the %s is the Max length of the Policy Name, which is 32.)
Checks on the Policy Name string:Character validity (Alpha-numeric and hyphen. Must start and end with alphanumeric only).	Invalid policy name. The policy name may contain alphanumeric characters and hyphen, and must begin and end with a alphanumeric character.
Check for maximum number of policies (48).	Cannot create a new policy because the total number of policies on this device has reached the maximum limit of %s.(where %s, is the maximum number of policies, currently 48)
While removing the policy:	
A policy cannot be deleted if applied on any port at the time of removal	Policy %s is currently in use on one or more ports and cannot be deleted.(where %s is the name of the policy).
The policy name does not exist.	MACsec policy %s does not exist.

Table Continued

Validation	Error/Warning/Prompt
While modifying a policy:	
When a policy is applied on an interface, the mode, CKN and CAK cannot be modified.	Cannot modify mode, CAK or CKN when the policy is in use on one or more ports.

Configuring mode of MACsec policy

Configure the mode of this MACsec policy. The mode determines how the CA Key Name (CKN) and CA Key (CAK) are obtained.

Syntax

```
[no] mode pre-shared-key ckn <CKN> cak <CAK>
```

Configure the MACsec policy to use pre-shared key mode. In the pre-shared key mode, the CA Key Name (CKN) and the CA Key (CAK) are set manually.

Configure the CA Key Name (CKN) of this MACsec policy. A CKN must be specified before the policy can be applied. Enter the CKN as a string of hexadecimal digits up to 32 characters long. If the CKN configured is less than 32 digits, it will be padded up to 32 hexadecimal digits with 0s. A CAK must be specified before the policy can be applied. Enter the CAK as a string of hexadecimal digits up to 64 characters long. If the CAK is less than 64 digits, it will be padded up to 64 hexadecimal digits with 0s.

mode

   Configure the mode of this MACsec policy.

pre-shared-key

   Configure the MACsec policy to use pre-shared key mode.

cak

   Configure the CA Key (cak) of this MACsec policy.

   Example: `Mode pre-shared-key ckn 37c9c2c45ddd cak`

ckn

   Configure the CA Key Name (CKN) of this MACsec policy.

   The CKN as a string of hexadecimal digits up to 32 characters long.

   The CAK as a string of hexadecimal digits up to 64 characters long.

   Example: `Mode pre-shared-key ckn 37c9c2c45ddd cak 2c45ddd012`

Encrypted-credentials mode

As CAK is a key and needs to be protected, when in encrypt-credentials mode the value gets encrypted and stored in the configuration.

Syntax

```
[no] mode pre-shared-key ckn <CKN> encrypted-cak <ENC-CAK>
```

Configure the CA Key (CAK) of this MACsec policy in encrypted form. A CAK must be specified before the policy can be applied. The value is an encrypted string previously read from a compatible Networking device.

mode

> Configure the mode of this MACsec policy.

CAK

> Configure the CA Key (CAK) of this MACsec policy.

CKN

> Configure the CA Key Name (CKN) of this MACsec policy.

encrypted-cak

> Configure the CA Key (CAK) of the MACsec policy, specificed as a base64 encoded AES-256 encrypted string.

Validation rules

Validation	Error/Warning/Prompt
Length check on CKN: Up to 32 digits long hex stringCKN should have even number of digits.	"Invalid CKN. The CKN should be an even number of hexadecimal digits no longer than %s digits."Eg: Invalid CKN. The CKN should be an even number of hexadecimal digits no longer than 32 digits.
Length check on CAK: Up to 64 digits long hex stringCAK should have even number of digits	"Invalid CAK. The CAK should be an even number of hexadecimal digits no longer than %s digits." Eg: Invalid CAK. The CAK should be an even number of hexadecimal digits no longer than 64 digits.
Length check for encrypted CAK: Either 44 or 88 characters long ASCII string.	"Invalid encrypted key."
Validity check on the encrypted CAK: Decrypt and apply the CAK rules:"Length check on CAK: Up to 64 digits long hex stringCAK should have even number of digits"	"Invalid encrypted key."

MACsec policy: configuring confidentiality (policy context)

Syntax

```
[no] confidentiality
```

Enable confidentiality in this MACsec policy. When confidentiality is enabled, data packets are encrypted and verified. When confidentiality is disabled, data packets are not encrypted, but they are still verified. By default, confidentiality is enabled.

confidentiality

> Enable confidentiality in this MACsec policy.

Validation rules

Validation	Error/Warning/Prompt
When in FIPS(Enhanced Security) mode, confidentiality change from ICV only (no confidentiality) to confidentiality is not allowed, if the policy is already in use on any port.	"Cannot enable confidentiality in Enhanced Security mode, when the policy is already in use on a port."

Configuring replay protection

Syntax

```
[no] replay-protection <replaywindowsize>
```

Configure the Replay Protection feature on this MACsec policy. When Replay Protection is enabled, the receiving port checks the IP number of all received packets. If a packet arrives out of sequence and the difference between the packet numbers exceeds the Replay Protection window size, the packet is dropped. By setting the replay window size to 0, it is mandated that all packets arrive in order. The default value of Replay Protection is enabled and the default value of the Replay Protection window size is 0.

replay-protection

Enable Replay Protection in this MACsec policy.

0-1024

Configure the Replay Protection window size value.

Validation rules

Validation	Error/Warning/Prompt
Replay-protection window range validation. 0-1024	Invalid value %s. The Replay Protection window size ranges from %s to %s.Eg:Invalid value 2000. The Replay Protection window size ranges from 0 to 1024.

Configuring include-sci-tag

Syntax

```
[no] include-sci-tag
```

Description

Include Secure Channel Identifier (SCI) tag information in the Security TAG (`SecTAG`) field. The SCI tag information s included by default.

The `no` form of this command causes the SCI tag information to be omitted.

ⓘ   If MACsec is enabled and SCI tag information is omitted (using the `no include-sci-tag` command) on a link between a 10G SmartRate copper port and a 1G copper port, no traffic will pass through the link. Therefore, to avoid this problem, use the same type of port on each side of the link or enable the SCI tag information with `include-sci-tag`. The SCI tag information is included by default.

Apply policy on a port-list

Syntax

```
[no] macsec apply policy <policy-name> ethernet PORT-LIST
```

Apply a MACsec policy to a list of ports.

apply

Apply a MACsec policy to a list of ports.

policy

Configure a MACsec policy.

policy-name

The MACsec policy to apply.

ethernet PORT-LIST

The port on which to apply the MACsec policy.

Validation rules

NOTE When a validation check fails on any port in the port-list, the CLI command aborts and returns error message. No configuration changes take place on any of the ports in the port-list.

Validation	Error/Warning/Prompt
Check whether the policy with the	Policy %s does not exist.
PolicyName exists, and whether it can be applied on any port.	Cannot apply the MACsec policy on ports, because the policy is not complete.
Another policy is already applied on this port.	Cannot apply the MACsec policy on port %s, because another MACsec policy is already configured on this port.
Trunk (lacp, trunk, dt-trunk, dt-lacp)	Cannot configure the MACsec policy on port %s when it is part of a trunk.(where %s is the port name.)

MKA configuration on a port-list

Syntax

```
[no] aaa port-access mka key-server-priority PRIORITY transmit-interval INTERVAL ethernet PORT-LIST
```

```
[no] aaa port-access <authenticator ...|supplicant ...web-based ...|mac-based ...|mka ...>
```

Configure the MACsec Key Agreement (MKA) protocol parameters.

Configure 802.1X (Port Based Network Access), MAC address based network access, or web authentication based network access or the MACsec Key Agreement (MKA) protocol on the device.

NOTE See the help for the commands `aaa port-access authenticator`, `aaa port-access supplicant`, `aaa port-access mac-based`, `aaa port-access web-based` for further details on authenticator, supplicant, MAC address based, and web authentication based network access configuration.

Syntax

```
aaa port-access mka key-server-priority transmit-interval <INTERVAL> [ethernet] PORT-LIST
```

Configure the MKA key server Priority. The key server priority is used by MKA protocol in selecting a key server. The participant with the lower server priority is selected as the key server. The default value is 16.

Syntax

```
aaa port-access mka key-server-priority 18
```

Configure the MKA transmit interval. MKA sends the periodic MKA protocol data unit (PDU) at this interval to the connected device to maintain MACsec connectivity on the link. The default value is 2 seconds.

Syntax

Configure the MACsec Key Agreement (MKA) protocol parameters.

```
aaa port-access mka

aaa port-access mka key-server-priority 18 transmit-interval

aaa port-access mka key-server-priority 18 transmit-interval 4

aaa port-access mka key-server-priority 18 transmit-interval 4 A1

aaa port-access mka key-server-priority

aaa port-access mka key-server-priority 5

aaa port-access mka key-server-priority 10 transmit-interval 6 a3
```

key-server-priority

Configure the MKA key server priority.

transmit-interval

Configure the MKA transmit interval.

0-31

Enter a Key Server priority value.

[ethernet] port-list

Enter a port number, a list of ports or 'all' for all ports.

transmit-interval

Configure the MKA transmit interval.

2-6

Enter a transmit interval value.

[ethernet] PORT-LIST

Enter a port number, a list of ports or 'all' for all ports.

Validation rules

When a validation check fails on any port in the Port-List, the CLI command aborts and returns error message. No configuration changes take place on any of the ports in the Port-List.

NOTE

Validation	Error/Warning/Prompt
Range check for MKA server priority value [0-31].	Invalid value %s. The MKA server priority ranges from %s to %s.Eg:Invalid value 50. The MKA server priority ranges from 0 to 31.
Range check for MKA transmit-interval [2-6]	Invalid value %s. The MKA transmit interval ranges from %s to %s.Eg:Invalid value 1. The MKA transmit interval ranges from 2 to 6 seconds.

Clearing MKA statistics on ports

Syntax

```
clear statistics mka ethernet port-list
```

Reset statistics counters.

```
clear statistics <PORT-LIST>|global|aclv4| aclv6|policy|mka
```

Reset the MKA protocol statistics.

aclv4

Reset IPv4 Access Control List statistics.

aclv6

Reset IPv6 Access Control List statistics.

dldp

Reset Device Link Detection Protocol (DLDP) statistics.

global

Reset the port counters in all sessions.

mac

Reset MAC Access Control List statistics.

macsec

Reset the MACsec protocol statistics.

mka

Reset the MKA protocol statistics.

policy

Reset policy statistics.

[ethernet] PORT-LIST

Reset the port counters in the current session.

Clear statistics MKA

```
clear statistics mka <PORT-LIST>
```

[ethernet] PORT-LIST The port for which to reset statistics.

```
clear statistics mka A1
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Reset the MKA protocol statistics.

Validation rules

Validation	Error/Warning/Prompt
Check for valid logical ports entered.	Module not present for port or invalid port: %s.(Parser thrown error for a Port).
Check if MACsec is enabled on the port before letting it proceed to clear.	Cannot clear MKA statistics, because MACsec is not enabled on the port.

Clearing MACsec statistics on ports

Syntax

```
clear statistics PORT-LIST|global|aclv4 ...|aclv6 ...|policy|mka ...|macsec ...
```

aclv4

Reset IPv4 Access Control List statistics.

aclv6

Reset IPv6 Access Control List statistics.

dldp

Reset Device Link Detection Protocol (DLDP) statistics.

global

Reset the port counters in all sessions.

mac

Reset MAC Access Control List statistics.

macsec

Reset the MACsec protocol statistics.

mka

Reset the MKA protocol statistics.

policy

Reset policy statistics.

[ethernet] PORT-LIST

Reset the port counters in the current session.

Reset statistics counters

```
clear statistics macsec [ethernet] PORT-LIST
```

Reset statistics counters.

Reset the MACsec protocol statistics

```
clear statistics macsec
```

Reset the MACsec protocol statistics.

Reset the MACsec protocol statistics

```
clear statistics macsec A1
```

Reset the MACsec protocol statistics

Validation rules

Validation	Error/Warning/Prompt
Check for valid logical ports entered.	Module not present for port or invalid port: %s. (Parser thrown error for a Port).
Check if MACsec is enabled on the port before letting it proceed to clear.	Cannot clear MACsec statistics, because MACsec is not enabled on the port.

Show commands

Show command for MACsec policies

Syntax

```
show macsec policy <policy-name>
```

Shows one or more MACsec policies.

policy-name

   A MACsec policy name up to 32 characters long.

show macsec policy

```
switch(config)# show macsec policy
Configuration - MACsec Policy
Policy Name : policy1
Cipher Suite      : AES-GCM-128
Include-SCI       : Yes
Confidentiality   : On          Confidentiality offset   : 0
Replay-Protection : On          Replay-Protection Window : 0
Mode : pre-shared-key (PSK)
CKN  : abcd
CAK  : abcd
Policy Name : macsecpolicy5
Cipher Suite      : AES-GCM-128
Include-SCI       : No
Confidentiality   : Off         Confidentiality offset   : 0
Replay-Protection : On          Replay-Protection Window : 0
Mode : pre-shared-key (PSK)
CKN  : abcd1111111121212121212abcd3434
CAK  : abab121212121212abcd34343434121212121212abcd34343434abcdefabcdef
```

show macsec policy Policy1

```
switch(config)# show macsec policy Policy1
Configuration - MACsec Policy
Policy Name : policy1
Cipher Suite      : AES-GCM-128
```

```
Include-SCI      : Yes
Confidentiality  : On          Confidentiality offset  : 0
Replay-Protection : On          Replay-Protection Window : 0
Mode : pre-shared-key (PSK)
CKN  : abcd
CAK  : abcd
```

Command validations

Validation	Error/Warning/Prompt
Check whether the policy with the name exists.	MACsec policy %s does not exist.
There are no MACsec policies configured on the system.	No MACsec policy is found.

Details

NOTE   In Manager mode.

Condition	Behavior
Include-credentials enabled/disabled	CAK value is displayed in plaintext format.
Encrypt-credentials enabled/disabled	CAK value is displayed in plaintext format.
In Enhanced Secure Mode (FIPS)	A dialogue is provided to proceed with display of sensitive information and only on a consent to proceed, policy details are displayed.

Show command for MACsec status

Syntax

```
show macsec status
```

Show the status of all MACsec-enabled ports.

status

Show the status of all MACsec-enabled ports.

show macsec status

```
switch(config)# show macsec status

Status and Configuration - MACsec Protocol
Interface  Policy                            Mode     Status Protection
---------  --------------------------------  -------  ------ ---------------
 A2        policy1                           PSK      Up     Confidentiality
 L22       policy1                           PSK      Down   Confidentiality
```

Command validation

Validation	Error/Warning/Prompt
MACsec is not enabled on ports.	MACsec is not enabled on any port.

Show command for MACsec status on a port

Syntax

```
show macsec status <port-num>
```

Show the status of all MACsec-enabled ports.

show macsec status

```
switch(config)# show macsec status A1
```

show command output

```
switch(config)# show macsec status A1
Status and Configuration - MACsec Protocol
Interface : A1
Policy       : Policy1
Transmitting : Yes
Receiving    : Yes
Protection   : Confidentiality
Transmit secure Channel
SCI                 : 000C29F6A4380004c
Secure Association
Association Number : 1 (old)
KI               : 4F18CE25228178FD15976E4C
LPN              : 2
SA-Start-time    : 01:02:19
SA-Stop-time     : 02:04:29
Association Number : 0 (current)
KI               : 4F18CE25228178FD15976E4C
LPN              : 3
SA-Start         : 04:05:11
SA-Stop-time     : 04:10:12
Receive secure Channel
SCI                 : 000C29F6A4380003b
Secure Association
Association Number : 0 (current)
KI               : 4F18CE29456aefFD15976E4C
LPN              : 121198
SA-Start         : 04:05:12
SA-Stop-time     : 04:10:13
```

Validation	Error/Warning/Prompt
Check whether MACsec is enabled on the port.	MACsec is not enabled on port %s.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Show command for MACsec statistics

Syntax

```
show macsec statistics <port-num>
```

statistics

Show MACsec statistics.

[ethernet] PORT-NUM

The port to show MACsec statistics for.

Show macsec statistics

```
switch(config)# show macsec statistics

Status and Counters - MACsec Protocol
 Interface : A1
 Receive Statistics
Totals (Since boot or last clear) :
Bytes Received       : 234435
Unicast Packets      : 0
Multicast Packets    : 0
Broadcast Packets    : 0
Errors (Since boot or last clear) :
Discarded Packets    : 0
Crypto Overruns      : 0
Packets With No Tag  : 0
Erroneous Packets    : 0
Packets With Bad Tag : 0
Packets With No SCI  : 0

Transmit Statistics
Totals (Since boot or last clear) :
Bytes Transmitted    : 28733989
Unicast Packets      : 0
Multicast Packets    : 0
Broadcast Packets    : 0
Errors (Since boot or last clear) :
Erroneous Packets    : 0
Packets Too Long     : 0
Interface : A2
Receive Statistics
Totals (Since boot or last clear) :
Bytes Received       : 234435
Unicast Packets      : 0
Multicast Packets    : 0
Broadcast Packets    : 0
Errors (Since boot or last clear) :
Discarded Packets    : 0
Crypto Overruns      : 0
Packets With No Tag  : 0
Erroneous Packets    : 0
Packets With Bad Tag : 0
Packets With No SCI  : 0

Transmit Statistics
Totals (Since boot or last clear) :
Bytes Transmitted    : 28733989
Unicast Packets      : 0
Multicast Packets    : 0
```

```
Broadcast Packets     : 0
Errors (Since boot or last clear) :
Erroneous Packets     : 0
Packets Too Long      : 0
```

Show macsec statistics A1

```
switch(config)# show macsec statistics A1
Status and Counters - MACsec Protocol
Interface : A1
Receive Statistics
Totals (Since boot or last clear) :
Bytes Received        : 234435
Unicast Packets       : 0
Multicast Packets     : 0
Broadcast Packets     : 0
Errors (Since boot or last clear) :
Discarded Packets     : 0
Crypto Overruns       : 0
Packets With No Tag   : 0
Erroneous Packets     : 0
Packets With Bad Tag  : 0
Packets With No SCI    : 0

Transmit Statistics
Totals (Since boot or last clear) :
Bytes Transmitted     : 28733989
Unicast Packets       : 0
Multicast Packets     : 0
Broadcast Packets     : 0
Errors (Since boot or last clear) :
Erroneous Packets     : 0
Packets Too Long      : 0
```

Command validations

NOTE

In Manager mode.

Validation	Error/Warning/Prompt
Check whether MACsec is enabled on the port.	MACsec is not enabled on port %s.

Show command for detailed MACsec statistics on a port

Syntax

```
show macsec statistics <port-num> detail
```

Show detailed statistics for a MACsec-enabled port.

statistics

Show MACsec statistics.

detail

Show detailed statistics for a MACsec-enabled port.

[ethernet] PORT-NUM

The port to show MACsec statistics for.

show macsec statistics A1 detail

```
switch(config)# show macsec statistics A1 detail

Status and Counters - MACsec Protocol
Interface : A1
Receive Statistics
Totals (Since boot or last clear) :
Bytes Received         : 234435
Unicast Packets        : 0
Multicast Packets      : 0
Broadcast Packets      : 0
Errors (Since boot or last clear) :
Discarded Packets      : 0
Crypto Overruns        : 0
Packets With No Tag    : 0
Erroneous Packets      : 0

Packets With Bad Tag : 0
Packets With No SCI  : 0

Transmit Statistics
Totals (Since boot or last clear) :
Bytes Transmitted      : 28733989
Unicast Packets        : 0
Multicast Packets      : 0
Broadcast Packets      : 0
Errors (Since boot or last clear) :
Erroneous Packets      : 0
Packets Too Long       : 0

Secure Channel Transmit Statistics
Encrypted Packets    : 0
Bytes Protected      : 0
Bytes Encrypted      : 0

Secure Association Statistics
Association Number   : 0 (old)
Protected Packets  : 0

Encrypted Packets  : 0
Association Number   : 1 (current)
Protected Packets  : 0
Encrypted Packets  : 0

Secure Channel Receive Statistics
Not using SA         : 0
Late                 : 0
Not Valid            : 0
Delayed              : 0
Valid                : 0
Bytes Validated      : 0
Bytes Decrypted      : 0

Secure Association Statistics
Association Number   : 1 (current)
Not using SA         : 0
```

```
Not Valid          : 0
Valid              : 0
```

Command validations

In Manager mode.

Validation	Error/Warning/Prompt
Check whether MACsec is enabled on the port.	MACsec is not enabled on port %s.

Show command for MKA status

Syntax

```
show port-access mka status <port-num>
```

Show the MKA protocol status information.

```
show port-access authenticator [...]|supplicant [...]|summary [...]| mka...
```

Show 802.1X (Port Based Network Access) supplicant or authenticator current status and configuration.

[ethernet] PORT-LIST

Show Web/MAC Authentication statistics and configuration.

authenticator

Show 802.1X (Port Based Network Access) authenticator current status, configuration or last session counters.

config

Show status of 802.1X, Web Auth, and MAC Auth configurations.

local-mac

Show Local MAC Authentication statistics and configuration.

mac-based

Show MAC Authentication statistics and configuration.

mka

Show the MKA protocol information.

summary

Show summary configuration information for all ports, including that overridden by RADIUS attributes.

supplicant

Show 802.1X (Port Based Network Access) supplicant current status and configuration.

web-based

Show Web Authentication statistics and configuration.

statistics

Show the MKA statistics.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

status

Show the MKA protocol status information.

Show port-access mka status

```
switch(config)# show port-access mka status
Status and Configuration - MKA Protocol
Interface : A2
Port MAC Address      : f0921c-4576fe
MKA Session Status    : Secured
CKN                   : abcd
MI                    : 1c64f054f894b5482defdf81
MN                    : 86
Capability            : IC, Conf, Offset 0
Transmit Interval     : 2
Key Server Priority   : 16
Key Server            : No

Live Peer List:

MI                       MN       PRI Capability          Rx-SCI
------------------------ -------- --- -------------------- ---------------
fb7f82788e4cd38dbc65dc55 119      16  IC, Conf, Offset 0   a45d36489bfe0002

Potential Peer List:
MI                       MN       PRI Capability          Rx-SCI
 ------------------------ -------- --- -------------------- ---------------

Interface : L2
Port MAC Address      : f0921c-4576fe
MKA Session Status    : Secured
CKN                   : abcdefabcd
MI                    : 1c64f054f894b5482defdf81
MN                    : 86
Capability            : IC, Conf, Offset 0
Transmit Interval     : 2
Key Server Priority   : 16
Key Server            : No

Live Peer List:
   MI                       MN        PRI Capability          Rx-SCI
------------------------ -------- --- -------------------- ---------------
fb7f82788e4cd38dbc65dc55 119       16  IC, Conf, Offset 0   a45d36489bfe0002

Potential Peer List:
   MI                       MN       PRI Capability          Rx-SCI
------------------------ -------- --- -------------------- ---------------
```

Show port-access MKA status A2

```
switch(config)# show port-access mka status A2
Status and Configuration - MKA Protocol
Interface : A2
Port MAC Address      : f0921c-4576fe
MKA Session Status    : Secured
CKN                   : abcd
```

```
MI                    : 1c64f054f894b5482defdf81
MN                    : 86
Capability            : IC, Conf, Offset 0
Transmit Interval     : 2
Key Server Priority   : 16
Key Server            : No

Live Peer List:
MI                          MN        PRI Capability           Rx-SCI
------------------------ -------- --- -------------------- ----------------
fb7f82788e4cd38dbc65dc55 119       16  IC, Conf, Offset 0   a45d36489bfe0002

Potential Peer List:
MI                          MN        PRI Capability           Rx-SCI
------------------------ -------- --- -------------------- ----------------
```

Command validations

NOTE   In Operator mode.

Validation	Error/Warning/Prompt
Check whether MACsec is enabled on the port.	MACsec is not enabled on port %s.

Show command for MKA statistics

Syntax

```
show port-access mka statistics <port-num>
```

Show the MKA statistics. When a PORT-NUM is used, the MKA statistics of the selected port are shown.

[ethernet] PORT-LIST

Show Web/MAC Authentication statistics and configuration.

authenticator

Show 802.1X (Port Based Network Access) authenticator current status, configuration or last session counters.

config

Show status of 802.1X, Web Auth, and MAC Auth configurations.

local-mac

Show Local MAC Authentication statistics and configuration.

mac-based

Show MAC Authentication statistics and configuration.

mka

Show the MKA protocol information.

summary

Show summary configuration information for all ports, including that overridden by RADIUS attributes.

supplicant

Show 802.1X (Port Based Network Access) supplicant current status and configuration.

web-based

Show Web Authentication statistics and configuration.

statistics

Show the MKA statistics.

status

Show the MKA protocol status information.

[ethernet] PORT-NUM

Specify the port number.

Show port-access MKA statistics

```
switch(config)# show port-access mka statistics
Status and Counters - MKA Protocol
CAs Established : 32
CAs Deleted     : 1
Interface : A1
  Tx MKPDUs                     : 16534893
  Rx MKPDUs                     : 16534893
  SAKs Distributed              : 0
  SAKs Received                 : 0
  MKPDUs With Invalid Version   : 0
  MKPDUs With Invalid CKN       : 0
  MKPDUs With Invalid ICV       : 0
  MKPDUs With Duplicate MI      : 0
  MKPDUs With Invalid MN        : 0
 Interface : A2
  Tx MKPDUs                     : 16534893
  Rx MKPDUs                     : 16534893
  SAKs Distributed              : 0
  SAKs Received                 : 0
  MKPDUs With Invalid Version   : 0
  MKPDUs With Invalid CKN       : 0
  MKPDUs With Invalid ICV       : 0
  MKPDUs With Duplicate MI      : 0
  MKPDUs With Invalid MN        : 0
```

Show port-access MKA statistics A1

```
switch(config)# show port-access mka statistics A1
Status and Counters - MKA Protocol
Interface : A1
  Tx MKPDUs                     : 16534893
  Rx MKPDUs                     : 16534893
  SAKs Distributed              : 0
  SAKs Received                 : 0
  MKPDUs With Invalid Version   : 0
  MKPDUs With Invalid CKN       : 0
  MKPDUs With Invalid ICV       : 0
```

```
MKPDUs With Duplicate MI     : 0
MKPDUs With Invalid MN       : 0
```

Command validations

NOTE

In Operator mode.

Validation	Error/Warning/Prompt
Check whether MACsec is enabled on the port (in case where the Port-Num is given)	MACsec is not enabled on port %s.
Check whether MACsec is enabled on any port.	MACsec is not enabled on any port.

Show tech command

Syntax

```
show tech macsec status|statistics
```

Show tech MACsec for either status or statistics.

Mutually exclusive commands with MACsec configuration on a port

Validation	Error/Warning/Prompt
aaa port-access authenticator [ethernet] PORT-LIST	Cannot enable 802.1X authenticator on port x when MACsec is enabled on that port.
aaa port-access supplicant [ethernet] PORT-LIST	Cannot enable 802.1X authenticator on port x when MACsec is enabled on that port
aaa port-access mac-based [ethernet] PORT-LIST	Cannot enable MAC Authentication on port x when MACsec is enabled on that port
[no] aaa port-access web-based [ethernet] PORT-LIST	Cannot enable Web Authentication on port x when MACsec is enabled on that port
Mesh command	Cannot configure mesh on port x because MACsec is enabled on that port.

MACsec Log messages

Event	Message
CAK Mismatch (Note that CAK will not be displayed)	MACsec Connectivity Association failed on port %s: Mismatch in the Integrity Check Value (ICV).
Throttled messages for CAK mismatch	Ceasing 'Detection of Macsec CAK Mismatch' message for 5m.

Table Continued

Event	Message
CKN Mismatch (or missing Policy on a port)	MACsec Connectivity Association failed on port %s: Mismatch in the CA Key Name (CKN).
Throttled messages for CKN Mismatch	Ceasing 'Detection of Macsec CKN Mismatch' message for 5m.
MKA session start	The MACsec Connectivity Association established on port %s.
MKA session end	The MACsec Connectivity Association ended on port %s.
Detection of replay attack	Possible replay attack on MACsec port %s.
Throttled replay attack messages	Ceasing 'Detection of Replay Attack' for 5m.
More than 1 MACsec client on a MACsec enabled port.	More than one MACsec clients detected on port %s
Throttled message for more than 1 client	Ceasing 'Detection of More than one Macsec clients' for 5m.
If MACsec is running in integrity mode, as it might be configured for integrity on either of sides	The MACsec is operating in Integrity Check (IC) mode on port %s.
If MACsec is running in encryption mode, as it might be configured to confidentiality from integrity on both sides	The MACsec is operating in Encryption mode on port %s.
when boot up FIPS test failed	FIPS test failed on port %s in slot %s. (where 1st %s is Port-Name and 2nd %s is for Slot-Name)
when FIPS bypass test failed	FIPS bypass self-test failed on port %s in slot %s.(where 1st %s is Port-Name and 2nd %s is for Slot-Name)
When hard expiry limit is reached	MACsec Secure Association Key (SAK) expired in hardware. Port %s blocked by MACsec.
Macsec errors	MACsec errors detected on port %s.

Remote Authentication Dial-In User Service (RADIUS) is a networking protocol that provides centralized Authentication, Authorization, and Accounting (AAA) management for users that connect and use a network service. RADIUS is the transport for AAA services. The services can include the user profiles including storing user credentials, user access policies, and user activity statistics which can reside on the same server. Gateway devices that control network access, such as remote access servers, VPN servers, and network switches, can use the RADIUS protocol to communicate with a RADIUS server for:

• Authentication — verifying user credentials regarding granted access to their networks.

• Authorization — verifying user access policy on how much and what kind of resources are allowed for an authenticated user.

• Accounting — keeping statistic information about the user activities for accounting purpose.

Configuring

Configuring the switch to support RADIUS-assigned ACLs

An ACL configured in a RADIUS server is identified by the authentication credentials of the client or group of clients the ACL is designed to support. When a client authenticates with credentials associated with a particular ACL, the switch applies that ACL to the switch port the client is using. To enable the switch to forward a client's credentials to the RADIUS server, you must first configure RADIUS operation and an authentication method on the switch.

Procedure

1. Configure RADIUS operation on the switch:

   ```
   radius-server host <ipv4-address> key <key-string>
   ```

   This command configures the IPv4 address and encryption key of a RADIUS server. The server should be accessible to the switch and configured to support authentication requests from clients using the switch to access the network.

2. Configure RADIUS network accounting on the switch (optional).

   ```
   aaa accounting network <start-stop|stop-only> radius
   ```

   You can also view ACL counter hits using either of the following commands:

   ```
   show access-list radius <port-list>
   ```

   ```
   show port-access <authenticator|mac-based|web-based> <port-list> clients detailed
   ```

3. Configure an authentication method. Options include 802.1X, web-based authentication, and MAC authentication. You can configure 802.1X, web-based authentication, and/or MAC authentication to operate simultaneously on the same ports.

- 802.1X Option:

```
aaa port-access authenticator <port-list>
aaa authentication port-access chap-radius
aaa port-access authenticator active
```

These commands configure 802.1X port-based access control on the switch, and activates this feature on the specified ports. For more on 802.1X configuration and operation, see **User authentication methods**.

- MAC Authentication Option:

```
aaa port-access mac-based <port-list>
```

This command configures MAC authentication on the switch and activates this feature on the specified ports. For more on MAC authentication, see **Web and MAC Authentication**.

- Web Authentication Option:

```
aaa port-access web-based <port-list>
```

This command configures web-based authentication on the switch and activates this feature on the specified ports. For more on web-based authentication, see **Web and MAC Authentication**.

Viewing

Viewing the currently active per-port CoS and rate-limiting configuration

While a RADIUS-assigned client session is active on a given port, any RADIUS-imposed values for the settings listed in the following table are applied as shown:

Table 18: *Application of RADIUS-assigned values*

Dynamic RADIUS assignment options	Static per-port setting options	Application of dynamic RADIUS assignment
802.1p Priority (CoS)	`qos priority <0-7>`	Applies per-client; that is, only to client whose authentication triggered the assignment. (Up to 32 clients supported per-port.)
Inbound (Ingress) Rate-Limiting	`rate-limit <all|bcast| icmp|mcast> in <kbps| percent>`	
Outbound (Egress) Rate-Limiting	`rate-limit <all|bcast| icmp|mcast> out <kbps| percent>`	Applies per-port; that is, to all clients on the port.

Syntax:

```
show port-access
```

```
web-based clients [port-list] detail
```

```
mac-based clients [port-list] detail
```

```
authenticator clients [port-list] detail
```

If the switch receives an 802.1p priority (CoS) and/or rate-limit setting(s) from a RADIUS server as the result of a client authentication on a port, the above commands display the assigned values while the client's session is active. When the session ends, the values for that client are no longer displayed.

The priority and inbound (ingress) rate-limit are applied only to the inbound traffic of the client whose authentication triggered the assignment. The outbound (egress) rate-limit applies to all outbound traffic on the port.

```
web-based clients [port-list] detail
```

Displays, for a Web authenticated client (web-based authentication), the status of RADIUS-assignment details for that client. See **Show commands for web-based authentication**.

```
mac-based clients [port-list] detail
```

Displays, for a MAC authenticated client (MAC-Auth), the status of RADIUS-assignment details for that client.

```
authenticator clients [port-list] detail
```

Displays, for an 802.1X- authenticated client, the status of RADIUS-assignment details for that client.

Example:

Suppose port 4 has been statically configured from the CLI with the following:

- 802.1p priority: 7

- Inbound rate-limit: 50 percent

- Outbound rate-limit: 50 percent

The above, statically configured, per-port priority and inbound rate-limit settings will not apply to any clients who authenticate and receive different inbound priority and rate-limit settings from the RADIUS server. If the RADIUS server also assigns an outbound rate-limit setting, which is applied per-port instead of per-client, then the outbound traffic from the port to all connected clients will be rate-limited according to the value set by the server for the most recently authenticated client. Thus, if client "X" authenticates with web-based authentication on port 4 with a RADIUS server that assigns a priority of 3, an inbound rate-limit of 10,000 kbps, and an outbound rate-limit of 50,000 kbps, then:

- The inbound traffic from client "X" will be subject to a priority of 3 and inbound rate-limit of 10,000 kbps. Traffic from other clients using the port will not be affected by these values.

- The combined rate-limit outbound for all clients using the port will be 50,000 kbps until either all client sessions end, or another client authenticates and receives a different outbound rate-limit.

NOTE

Mixing CLI-configured and RADIUS-assigned rate-limiting on the same port can produce unexpected results. See **Per-port bandwidth override**.

Where multiple clients are currently authenticated on a given port where outbound (egress) rate-limiting values have been assigned by a RADIUS server, the port operates with the outbound rate-limit assigned by RADIUS for the most recently authenticated client. Any earlier outbound rate-limit values assigned on the same port for other authenticated client sessions that are still active are superseded by the most recent RADIUS-assigned value. For example, if client "X" is authenticated with an outbound rate-limit of 750 kbps, and client "Y" later becomes authenticated with an outbound rate-limit of 500 kbps while the session for client "X" is still active, then the port operates with an outbound rate-limit of 500 kbps for both clients.

Assignment method on port 10	802.1p	Inbound rate-limit	Outbound rate-limit
Statically Configured Values	7	100,000 kbps	100,000 kbps[1]
RADIUS-assigned when client "X" authenticates	3	10,000 kbps	50,000 kbps[1]

[1] Combined rate-limit output for all clients active on the port.

Figure 49: *Results of client authentication on port 4*

```
HP Switch(eth-10)# show port-access web-based clients 4 detail

                                            Indicates there is an authenticated
                                            client session running on port 10.
 Port Access Web-Based Client Status Detailed

  Client Base Details :
   Port           : 4
   Session Status : authenticated    Session Time(sec): 125
   Username       : client-X         MAC Address      : 0017a4-e6d787
   IP             : n/a

  Access Policy Details :
   COS Map        : 33333333          In Limit Kbps     : 10000
   Untagged VLAN  : 10                Out Limit Kbps    : 50000
   Tagged VLANs   : 20
   RADIUS-ACL List : No Radius ACL List

        Shows the values assigned to the client's traffic for inbound CoS     Shows the RADIUS-assigned value on the
        (802.1p priority) and client's maximum inbound traffic volume.         port for maximum outbound traffic volume.
```

Viewing CLI-configured rate-limiting and port priority for ports

Syntax:

```
show rate-limit all [port-list]
```

```
show qos port-priority
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

These commands show the CLI-configured rate-limiting and port priority for the selected ports. They also include indications of RADIUS-assigned rate-limiting and client traffic priority settings for any clients that may be authenticated on the same ports.

Figure 50: *Displaying rate-limiting for multiple ports (CLI and RADIUS)*

```
HP Switch# show rate-limit all 1-5

 All-Traffic Rate Limit Maximum %

         | Inbound                    Radius       | Outbound                   Radius
  Port   | Limit      Mode      Override     | Limit      Mode      Override
  ----- + --------- --------- ----------- + --------- --------- -----------
   1     | Disabled  Disabled  No-override  | Disabled  Disabled  No-override
   2     | Disabled  Disabled  No-override  | Disabled  Disabled  No-override
   3     | 1000      kbps      Override     | 1000      kbps      50000
   4     | 50        %         Override     | 50        %         50000
   5     | 50        %         No-override  | 50        %         No-override
```

Ports 3 - 5 have CLI-configured inbound per-port rate-limits and clients with RADIUS-assigned inbound per-client rate-limits. (To see the per-client RADIUS settings, use the command illustrated in figure 7-1.)

Ports 3 - 5 also have CLI-configured outbound per-port rate-limits and clients with RADIUS-assigned outbound (per-port) rate-limits.

Figure 51: *Displaying priority for multiple ports (CLI and RADIUS)*

```
HP Switch# show qos port-priority

  Port priorities

  Port Apply rule  | DSCP    Priority     Radius Override
  ---- ----------- + ------ ----------- ----------------
   1     No-override |            No-override No-override
   2     No-override |            No-override No-override
   3     No-override |            No-override No-override
   4     Priority    |    7                   Override
   5     No-override |            No-override No-override
```

Port 4 has CLI-configured per-port priority and a client with a RADIUS-assigned priority. (To see the RADIUS-assigned per-client priority settings, use the command illustrated in figure 7-1.)

Using

ACE syntax configuration options in a RADIUS server, using the standard attribute in an IPv4 ACL (Example)

The Standard attribute (92), when used in an ACL without the HP-Nas-Rules-IPv6 VSA, filters IPv4 traffic inbound from the authenticated client. (Any IPv6 traffic inbound from the client is dropped.) This example illustrates configuring RADIUS-assigned IPv4 ACL support on FreeRADIUS using the standard attribute for two different client identification methods (username/password and MAC address).

Procedure

1. Enter the ACL standard attribute in the FreeRADIUS `dictionary.rfc4849` file.

```
ATTRIBUTE Nas-FILTER-Rule 92
```

2. Enter the switch IP address, NAS (Network Attached Server) type, and the key used in the FreeRADIUS clients.conf file. For example, if the switch IP address is 10.10.10.125 and the key ("secret") is "1234", you would enter the following in the server's clients.conf file:

Figure 52: *Switch identity information for a freeRADIUS application*

```
client 10.10.10.125
nastype =  other
secret = 1234
```

Note: The **key** configured in the switch and the **secret** configured in the RADIUS server supporting the switch must be identical. See the chapter titled "RADIUS Authentication and Accounting" in the latest *Access Security Guide* for your switch.

3. For a given client username/password pair or MAC address, create an ACL by entering one or more ACEs in the FreeRADIUS "users" file. Remember that every ACL created automatically includes an implicit `deny in ip from any to any` ACE.

For example, to create identical ACL support for the following:

- Client having a username of "mobilE011" and a password of "run10kFast"

- Client having a MAC address of 08-E9-9C-4F-00-19

The ACL in this example must achieve the following:

- Permit http (TCP port 80) traffic from the client to the device at 10.10.10.101

- Deny http (TCP port 80) traffic from the client to all other devices

- Permit all other traffic from the client to all other devices

---

NOTE

For information on syntax details for RADIUS-assigned ACLs, see **Using HPE VSA 63 to assign IPv6 and IPv4 ACLs** on page 281.

To configure the above ACL, enter the username/password and ACE information shown in **Figure 53: Configuring the FreeRADIUS server to support ACLs for the indicated clients** on page 280.

Figure 53: *Configuring the FreeRADIUS server to support ACLs for the indicated clients*

Client's Username (802.1X or Web Authentication)   Client's Password (802.1X or Web Authentication)

```
mobilE011 Auth-Type:= Local, User-Password == run10kFast
        Nas-FILTER-Rule = "permit in tcp from any to host 10.10.10.101" 80,
        Nas-FILTER-Rule += "deny in tcp from any to any" 80,
        Nas-FILTER-Rule += "permit in ip from any to any"
```

Client's Username (MAC Authentication)   Client's Password (MAC Authentication)

```
08E99C4F0019 Auth-Type:= Local, User-Password == 08E99C4F0019
        Nas-FILTER-Rule = "permit in tcp from any to host 10.10.10.101" 80,
        Nas-FILTER-Rule += "deny in tcp from any to any" 80,
        Nas-FILTER-Rule += "permit in ip from any to any"
```

Note that when the client MAC address is used for authentication, it is used in both the username and password spaces in the entry.

Using HPE VSA 63 to assign IPv6 and IPv4 ACLs

The ACL VSA `HP-Nas-Rules-IPv6=1` is used in conjunction with the standard attribute (`Nas-Filter-Rule`) for ACL assignments filtering both IPv6 and IPv4 traffic inbound from an authenticated client. For example, to use these attributes to configure a RADIUS-assigned ACL on a FreeRADIUS server to filter both IPv6 and IPv4 ACLs, perform these steps:

Procedure

1. Enter the following in the FreeRADIUS `dictionary.hp` file:

   - vendor-specific ID

   - ACL VSA for IPv6 ACLs (63)

   - HP-Nas-Rules-IPv6 VALUE setting to specify both IPv4 and IPv6 (1)

Figure 54: *Configuring the VSA for RADIUS-assigned IPv6 and IPv4 ACLs in a FreeRADIUS server*

```
VENDOR          HP      11                              HP Vendor-Specific ID
BEGIN-VENDOR    HP
ATTRIBUTE       HP-Nas-Rules-IPv6 63 INTEGER
END-VENDOR      HP
                                                        VSA for RADIUS-Assigned IPv6 ACL
                                                        option.

      Note: If you were also using the RADIUS server to administer 802.1p (CoS) priority and/or Rate-Limiting,
      you would also insert the ATTRIBUTE entries for these functions above the END-VENDOR entry.
```

2. Enter the switch IPv4 address, NAS (Network Attached Server) type, and the key used in the FreeRADIUS clients.conf file. For example, if the switch IP address is 10.10.10.125 and the key ("secret") is "1234", you would enter the following in the server's clients.conf file:

Figure 55: *Switch identity information for a freeRADIUS application*

```
client 10.10.18.12            Note: The key configured in the switch and the
nastype =   other             secret configured in the RADIUS server
secret = 1234                 supporting the switch must be identical. See the
                              chapter titled "RADIUS Authentication and
                              Accounting" in the latest Access Security Guide
                              for your switch.
```

3. For a given client username/password pair, create an ACL by entering one or more IPv6 and IPv4 ACEs in the FreeRADIUS "users" file. Remember that the ACL created to filter both IPv4 and IPv6 traffic automatically includes an implicit deny in ip from any to any ACE at the end of the ACL in order to drop any IPv4 and IPv6 traffic that is not explicitly permitted or denied by the ACL. For example, to create ACL support for a client having a username of "Admin01" and a password of "myAuth9". The ACL in this example must achieve the following:

   a. Permit http (TCP port 80) traffic from the client to the device at FE80::a40.

   b. Deny http (TCP port 80) traffic from the client to all other IPv6 addresses.

   c. Permit http (TCP port 80) traffic from the client to the device at 10.10.10.117.

   d. Deny http (TCP port 80) traffic from the client to all other IPv4 addresses.

   e. Deny Telnet (TCP port 23) traffic from the client to any IPv4 or IPv6 addresses.

   f. Permit all other IPv4 and IPv6 traffic from the client to all other devices.

4. To configure the above ACL, enter the username/password and ACE information, as shown in this example:

Figure 56: *Configuring a FreeRADIUS server to filter IPv4 and IPv6 traffic for a client with correct credentials*

```
Client's Username (802.1X or Web Authentication)          Client's Password (802.1X or Web Authentication)

   Admin01 Auth-Type:= Local, User-Password == myAuth9
        HP-Nas-Rules-IPv6 = 1,                    ──── IPv6 VSA for the standard attribute-92.
        Nas-filter-rule = "permit in tcp from any to FE80::a40 80",
        Nas-filter-rule += "deny in tcp from any to ::/0 80",
        Nas-filter-rule += "permit in tcp from any to 10.10.10.117 80",
        Nas-filter-rule += "deny in tcp from any to 0.0.0.0/0 80",
        Nas-filter-rule += "deny in tcp from any to any 23",
        Nas-filter-rule += "permit in ip from any to any"

                                          In an ACL that filters both IP and IPv6 traffic
                                          type, the any keyword used for a destination
                                          includes both IPv4 and IPv6 traffic. As a result,
                                          this ACE denies Telnet traffic from the client to
                                          any IPv4 or IPv6 address.
```

Using HPE VSA 61 to assign IPv4 ACLs

The recommended use of this option is to support legacy ACL configurations that rely on VSA 61. Hewlett Packard Enterprise recommends using the standard attribute (92) for new, RADIUS-based IPv4 ACLs, see **Nas-filter-Rule attribute options** on page 302, and **ACE syntax configuration options in a RADIUS server, using the standard attribute in an IPv4 ACL (Example)** on page 279.

This example uses the VSA attribute 61 for configuring RADIUS-assigned IPv4 ACL support on FreeRADIUS for two different client identification methods (username/password and MAC address).

Procedure

1. Enter the vendor-specific ID and the ACL VSA in the FreeRADIUS dictionary file:

Figure 57: *Configuring the VSA for RADIUS-assigned IPv4 ACLs in a FreeRADIUS server*

```
VENDOR          HP      11        ──── HP Vendor-Specific ID
BEGIN-VENDOR    HP
ATTRIBUTE       HP-Nas-filter-Rule 61 STRING  ──── HP Vendor-Specific Attribute for
END-VENDOR      HP                                  RADIUS-Assigned ACLs

                                    Note that if you were also using the RADIUS server to
                                    administer 802.1p (CoS) priority and/or Rate-Limiting,
                                    you would also insert the ATTRIBUTE entries for these
                                    functions above the END-VENDOR entry.
```

2. Enter the switch IPv4 address, NAS (Network Attached Server) type, and the key used in the FreeRADIUS `clients.conf` file. For example, if the switch IP address is 10.10.10.125 and the key ("secret") is "1234", you would enter the following in the server's `clients.conf` file:

Figure 58: *Switch identity information for a freeRADIUS application*

```
client 10.10.18.12
nastype =  other
secret = 1234

                    Note: The key configured in the switch and the
                    secret configured in the RADIUS server
                    supporting the switch must be identical. See the
                    chapter titled "RADIUS Authentication and
                    Accounting" in the latest Access Security Guide
                    for your switch.
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

3. For a given client username/password pair, create an ACL by entering one or more IPv4 ACEs in the FreeRADIUS "users" file. Remember that the ACL created to filter IPv4 traffic automatically includes an implicit `deny in ip from any` to any ACE (for IPv4). For example, to create ACL support for a client having a username of "User-10" and a password of "auth7X". The ACL in this example must achieve the following:

- Permit http (TCP port 80) traffic from the client to the device at 10.10.10.117.

- Deny http (TCP port 80) traffic from the client to all other IPv4 addresses.

- Deny Telnet (TCP port 23) traffic from the client to any IPv4 address.

- Permit all other IPv4 traffic from the client to all other devices.

4. To configure the above ACL, you would enter the username/password and ACE information shown in **Figure 59: Configuring a FreeRADIUS server to filter IPv4 traffic for a client with the correct credentials** on page 283 into the FreeRADIUS "users" file.

Figure 59: *Configuring a FreeRADIUS server to filter IPv4 traffic for a client with the correct credentials*

Displaying the current RADIUS-assigned ACL activity on the switch

These commands output data indicating the current ACL activity imposed per-port by RADIUS server responses to client authentication.

Syntax:

```
show access-list radius <port-list>
```

For the specified ports, this command lists:

- Whether the ACL for the indicated client is configured to filter IPv4 traffic only, or both IPv4 and IPv6 traffic. See **Nas-Filter-Rule-Options** for more on this topic.

- The explicit ACEs, switch port, and client MAC address for each ACL dynamically assigned by a RADIUS server as a response to client authentication.

If `cnt` (counter) is included in an ACE, then the output includes the current number of inbound packet matches the switch has detected in the current session for that ACE, see **ACE syntax in RADIUS servers**.

Note: If there are no ACLs currently assigned to any port in `<port-list>`, executing this command returns only the system prompt. If a client authenticates but the server does not return a RADIUS-assigned ACL to the client port, then the server does not have a valid ACL configured and assigned to that client's authentication credentials.

Example:

The following output shows that a RADIUS server has assigned an ACL to port B1 to filter inbound traffic from an authenticated client identified by a MAC address of 00-17-A4-E6-D7-87.

Figure 60: *A RADIUS-assigned ACL application to a currently active client session*

```
HP Switch(config)# show access-list radius b1

Radius-configured Port-based ACL for
Port B1, Client -- 0017A4E6D787

IPv6 ACLs enabled (HP-Nas-Rules-Ipv6): FALSE
deny in tcp from any to 10.30.248.184 23 cnt
  Packet Hit Counter : 1
deny in tcp from any to 10.30.248.184 80 cnt
  Packet Hit Counter : 10
permit in tcp from any to 10.30.248.184 7
permit in udp from any to 10.30.248.184 7
deny in tcp from any to 10.30.248.184 161 cnt
  Packet Hit Counter : 25
deny in udp from any to 10.30.248.184 161 cnt
  Packet Hit Counter : 7
permit in ip from any to any
```

Indicates MAC address identity of the authenticated client on the specified port. This data identifies the client to which the ACL applies.	
Indicates that IPv6 traffic filtering is not enabled for the ACL assigned to the authenticated client.	
Lists "deny" ACE for Inbound Telnet (23 = TCP port number) traffic, with counter configured to show the number of matches detected.	
Lists current counter for the preceding "Deny" ACE.	
Lists "permit" ACEs for inbound TCP and UDP traffic, with no counters configured.	
Note that the implicit "deny any/any" included automatically at the end of every ACL is not visible in ACL listings generate by the switch.	

Syntax:

```
show port-access <web-based|mac-based|authenticator> clients <port-list> detailed
```

For ports in `<port-list>` configured for authentication, this command shows the details of the RADIUS-assigned features listed below that are active as the result of a client authentication. (Ports in `<port-list>` that are not configured for authentication are not listed.)

- **Client Base Details:**

  Port

    Port number of port configured for authentication.

  Session Status

    Indicates whether there is an authenticated client session active on the port. Options include `authenticated` and `unauthenticated`.

  Username

    During an authenticated session, shows the user name of the authenticated client. If the client is not authenticated, this field is empty.

  IP

    Shows the authenticated client's IP address, if available. Requires DHCP snooping enabled on the switch. When "n/a" appears in the field, the switch has not been able to acquire the client's IP address. Note: Where the client IP address is available to the switch, it can take a minute or longer for the switch to learn the address. For more on this topic, see **Configuring RADIUS accounting** on page 226.

  Session Time (sec)

    For an unauthenticated session, indicates the elapsed time in seconds since the client was detected on the port. For an authenticated session, this indicates the elapsed time in seconds since the client was authenticated on the port.

  MAC Address

    During an authenticated session, shows the MAC address of the authenticated client.

- **Access Policy Details:**

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

COS Map

Indicates the 802.1p priority assigned by the RADIUS server for traffic inbound on the port from an authenticated client. The field shows an eight-digit value where all digits show the same, assigned 802.1p number. For example, if the assigned 802.1p value is 5, then this field shows `55555555`. If an 802.1p priority has not been assigned by the RADIUS server, this field shows `Not Defined`.

Untagged VLAN

VLAN ID (VID) of the untagged VLAN currently supporting the authenticated connection.

Tagged VLANs

VLAN IDs (VIDs) of any tagged VLANs currently supporting the authenticated connection.

RADIUS ACL List

Lists the explicit ACEs in the ACL assigned to the port for the authenticated client. Includes the ACE "Hit Count" (matches) for ACEs configured with the `cnt` option, see **ACE syntax in RADIUS servers**. If a RADIUS ACL for the authenticated client is not assigned to the port, `No Radius ACL List` appears in this field.

In Limit Kbps

Indicates the ingress rate-limit assigned by the RADIUS server to the port for traffic inbound from the authenticated client. If there is no ingress rate-limit assigned, then `Not Set` appears in this field.

Out Limit Kbps

Indicates the egress rate-limit assigned by the RADIUS server to the port for traffic outbound to the authenticated client. If there is no egress rate-limit assigned, then `Not Set` appears in this field.

Output showing current RADIUS-applied features

```
HP Switch(config)# show port-access web-based clients 10 detailed

 Port Access Web-Based Client Status Detailed

  Client Base Details :
  Port           : 9
  Session Status : authenticated   Session Time(sec) : 5
  Username       : acluser1         MAC Address        : 0017a4-e6d787
  IP             : n/a

  Access Policy Details :
  COS Map        : 77777777         In Limit Kbps      : 1000
  Untagged VLAN  : 10               Out Limit Kbps     : Not Set
  Tagged VLANs   : 20
  RADIUS-ACL List :
     deny in 23 from any to 10.0.8.1/24 23 CNT
        Hit Count: 1
     permit in 1 from any to 10.0.10.1/24 CNT
        Hit Count: 112
     deny in udp from any to any 67-68 CNT
        Hit Count: 7
     permit in ip from any to any CNT
        Hit Count: 125
```

Table 19: *ICMP type numbers and keywords*

IPv4 ICMP		IPv6 ICMP	
#	Keyword	#	Keyword
0	echo reply	1	destination unreachable
3	destination unreachable	2	packet too big
4	source quench	3	time exceeded
5	redirect	4	parameter problem
8	echo request	128	echo request
9	router advertisement	129	echo reply
10	router solicitation	130	multicast listener query
11	time-to-live exceeded	131	multicast listener reply
12	IP header bad	132	multicast listener done
13	timestamp request	133	router solicitation
14	timestamp reply	134	router advertisement
15	information request	135	neighbor solicitation
16	information reply	136	neighbor advertisement
17	address mask request	137	redirect message
18	address mask reply	138	router renumbering
		139	icmp node information query
		140	icmp node information response
		141	inverse neighbor discovery solicitation message
		142	inverse neighbor discovery advertisement message
		143	version 2 multicast listener report
		144	home agent address discovery request message
		145	home agent address discovery reply message
		146	mobile prefix solicitation

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

IPv4 ICMP		IPv6 ICMP	
		147	mobile prefix advertisement
		148	certification path solicitation message
		149	certification path advertisement message
		151	multicast router advertisement
		152	multicast router solicitation
		153	multicast router termination

Overview

The following information provides an overview about RADIUS services supported on a switch, including CoS (802.1p priority), ingress and egress rate-limiting, and ACL client services on a RADIUS server. For information on configuring client authentication capability on the switch, see **RADIUS Authentication, Authorization, and Accounting** on page 194.

Table 20: *RADIUS services supported on the switch*

Service	Application	Standard RADIUS attribute	Vendor-specific RADIUS attribute (VSA)
Cos (Priority)	per-user	59	40
Ingress Rate-Limiting	per-user	—	46
Egress Rate-Limiting	per-port[1]	—	48
ACLs			
IPv6 and/or IPv4 ACEs(NAS-Filter-Rule)	per-user	92	61
NAS-Rules-IPv6 (sets IP mode to IPv4-only or IPv4 and IPv6)	per-user	—	63

[1] If multiple clients are authenticated on a port where per-port rules are assigned by a RADIUS server, then the most recently assigned rule is applied to the traffic of all clients authenticated on the port.

> **NOTE**
> Hewlett Packard Enterprise recommends using the Standard RADIUS attribute, if available. Where both a standard attribute and a VSA are available, the VSA is maintained for backwards compatibility with configurations based on earlier software releases.

About RADIUS server support

RADIUS client and server requirements

- Clients can be dual-stack, IPv4-only or IPv6 only.

- Client authentication can be through 802.1X, MAC authentication, or web-based authentication. (clients using web-based authentication must be IPv4-capable.)

RADIUS server configuration for CoS (802.1p priority) and rate-limiting

The following information provides general guidelines for configuring RADIUS servers, so that the features listed in the following table can be dynamically applied on ports that support authenticated clients.

Table 21: *CoS and rate-limiting services*

Service	Control method and operating notes
802.1p Assigns a RADIUS-configured 802.1p priority to the inbound packets received from a specific client authenticated on a switch port. **NOTE** This attribute is assigned per authenticated user.	Standard Attribute used in the RADIUS server: 59 (This is the preferred attribute for new or updated configurations.) Vendor-Specific Attribute used in the RADIUS server. (This attribute is maintained for legacy configurations.) vendor-specific ID: 11 VSA: 40 Setting: User-Priority-Table=xxxxxxxx where: x=desired 802.1p priority Note: This is an eight-digit field. Enter the same x-value for all eight digits. Requires a port-access authentication method (802.1X, Web Auth, or MAC Auth) configured on the client's port on the switch. For more on 802.1p priority levels, see "Quality of Service (QoS)" in the advanced traffic management guide for your switch.
Ingress (inbound) rate-limiting per-user Assigns a RADIUS-configured bandwidth limit to the inbound packets received from a specific client authenticated on a port. Note: This attribute is assigned per-authenticated-user, not per-port. To assign a per-port inbound rate-limit, use the rate-limit all in CLI command instead of this option.	Vendor-Specific Attribute used in the RADIUS server. vendor-specific ID: 11 VSA: 46 Setting: HP-Bandwidth-Max-Egress=< bandwidth-in-Kbps > Note: RADIUS-assigned rate-limit bandwidths must be specified in Kbps. (Bandwidth percentage settings are not supported.) Using a VSA on a RADIUS server to specify a per-user rate-limit requires the actual Kbps to which you want to limit ingress (inbound) traffic volume. For example, to limit inbound traffic on a gigabit port to half of the port's bandwidth capacity requires a VSA setting of 500,000 Kbps. Requires a port-access authentication method (802.1X, Web Auth, or MAC Auth) configured on the client's port on the switch. The actual bandwidth available for ingress traffic from an authenticated client can be affected by the total bandwidth available on the client port. See **Per-port bandwidth override**.
Egress (outbound) rate-limiting per-port Assigns a RADIUS-configured bandwidth limit to the outbound traffic sent to a switch port.	Vendor-Specific Attribute used in the RADIUS server. vendor-specific ID: 11 VSA: 48 (string=HP) Setting: HP-RATE-LIMIT=< bandwidth-in-Kbps >

Service	Control method and operating notes
	Note: RADIUS-assigned rate-limit bandwidths must be specified in Kbps. (Bandwidth percentage settings are not supported.) Using a VSA on a RADIUS server to specify a per-port rate-limit requires the actual Kbps to which you want to limit outbound traffic volume. For example, to limit outbound traffic on a gigabit port to half of the port's bandwidth capacity requires a VSA setting of 500,000 Kbps.
	In instances where multiple, authenticated clients are using this feature on the same switch port, only one (per-port) rate limit will be applied. In this case, the actual rate used is the rate assigned by the RADIUS server to the most recently authenticated client. This rate remains in effect as long as any authenticated client remains connected on the port.
	Requires a port-access authentication method (802.1X, Web Auth, or MAC Auth) configured on the client's port on the switch.
	The actual bandwidth available for egress traffic from an authenticated client can be affected by the total bandwidth available on the client port. See **Per-port bandwidth override**.

To configure support for services on a specific RADIUS server application, see the documentation provided with the RADIUS application.

Applied rates for RADIUS-assigned rate limits

A particular switch does not necessarily support all rates. The actual rate applied is rounded down to the nearest value that is supported by switch. The increments are in graduated steps, as described in the following table.

Table 22: *RADIUS-assigned rate-limit increments*

RADIUS-assignedbits-per-second rate limit	Applied rate-limiting increment
1 - 10,999,999	100 Kbps
11,000,000 - 100,999,999	1 Mbps
101,000,000 - 999,999,999	10 Mbps
1,000,000,000 - 10 Gbps	100 Mbps

For example, some of the following RADIUS-assigned rates fall between their respective incremental values, resulting in applied rates lower than the RADIUS-assigned rates. However, others match their respective incremental values, resulting in no difference between the RADIUS-assigned rate limits and the applied rate limits.

Table 23: *Assigned and applied rate limits example*

RADIUS-assigned bandwidth (Kbps)	Applied increments	Applied rate limit (Kbps)	Difference/Kbps
5,250	100 Kbps	5,200	50
50,250	1 Mbps	50,000	250 Kbps
51,000	1 Mbps	51,000	0
525,000	10 Mbps	520,000	5,000 Kbps
530,000	10 Mbps	530,000	0
1,250,000	100 Mbps	1,200,000	50,000 Kbps
1,300,000	100 Mbps	1,300,000	0

Per-port bandwidth override

Hewlett Packard Enterprise recommends that rate-limiting be configured either solely through RADIUS assignments or solely through static CLI configuration on the switch unless the potential for the override described below is specifically desired.

Ingress (inbound) traffic

RADIUS-assigned ingress rate-limits are applied to individual clients instead of to the client's port. But if you use the CLI to configure a per-port ingress rate-limit on the same port where an authenticated client receives a RADIUS-assigned ingress rate-limit, the client's assigned ingress limit can be reduced by the CLI-configured port ingress limit. This occurs if the port reaches its CLI-configured rate-limit maximum before the client reaches its RADIUS-assigned rate-limit maximum, thus denying the client its intended maximum.

Egress (outbound) traffic

The most recent RADIUS-assigned egress rate-limit specifies the maximum egress rate-limit for a port, even if the CLI has also been used to configure an egress rate limit on the port.

Rate-limit assignment method		Rate-limit actions and restrictions
Inbound	CLI ingress rate-limit per-port `rate-limit all in`	Determines the maximum ingress bandwidth available on the port, regardless of any RADIUS-assigned per-client rate-limits dynamically assigned to the same port.
	RADIUS ingress rate-limit per-client VSA 46	Each client is allowed the inbound bandwidth individually assigned to it by the RADIUS server, up to the port's physical capacity, unless the available bandwidth on the port has been reduced by a CLI-assigned per-port bandwidth limit.

Table Continued

Rate-limit assignment method		Rate-limit actions and restrictions
Outbound	CLI egress rate-limit per-port `rate-limit all out`	Determines the maximum egress bandwidth available on the port, unless there is also a RADIUS-assigned per-port rate limit on the port.
	RADIUS egress rate-limit per client VSA 48	The most recent client to authenticate determines the maximum egress bandwidth on the port for all outbound traffic, regardless of any CLI-assigned per-port outbound rate-limit.

For example, suppose the CLI is used to configure a gigabit port to have an ingress rate limit of 500,000 Kbps (50% of available bandwidth), and is receiving 450,000 Kbps of traffic from existing clients. If a RADIUS server then authenticates a new client with an ingress rate-limit of 100,000 Kbps, the maximum ingress rate limit actually available for the new client is 50,000 Kbps as long as the bandwidth usage by the other clients already on the port remains at 450,000 Kbps.

For more on static rate-limiting, see "Rate-Limiting" in the "Port Traffic Controls" in the management and configuration guide for your switch.

Configuring and using dynamic (RADIUS-assigned) access control lists

A RADIUS-assigned ACL is configured on a RADIUS server and dynamically assigned by the server to filter IP traffic from a specific client after the client is authenticated by the server.

The information in this section describes how to apply RADIUS-assigned ACLs on the switch, and assumes a general understanding of ACL structure and operation. If you need information on ACL filtering criteria, design, and operation, see the following:

- **IPv4 Access Control Lists (ACLs)** on page 398

- "IPv6 Access Control Lists (ACLs)" in the latest IPv6 configuration guide for your switch.

Overview of RADIUS-assigned, dynamic ACLs

RADIUS-assigned ACLs enhance network and switch management access security and traffic control by permitting or denying authenticated client access to specific network resources and to the switch management interface. This includes preventing clients from using TCP or UDP applications, ICMP packet types, and IGMP (IPv4 only) if you do not want their access privileges to include these capabilities.

Traffic applications

The switch supports RADIUS-assigned ACLs for the following traffic applications:

- Inbound IPv4 traffic only

- Inbound IPv4 and IPv6 traffic

This feature is designed for use on the network edge to accept RADIUS-assigned ACLs for Layer-3 filtering of IP traffic entering the switch from authenticated clients. A given RADIUS-assigned ACL is identified by a unique username/password pair or client MAC address, and applies only to IP traffic entering the switch from clients that authenticate with the required, unique credentials. The switch allows multiple RADIUS-assigned ACLs on a given port, up to the maximum number of authenticated clients allowed on the port. Also, a RADIUS-assigned ACL for a given client's traffic can be assigned regardless of whether other ACLs assigned to the same port are statically configured on the switch.

A RADIUS-assigned ACL filters IP traffic entering the switch from the client whose authentication caused the ACL assignment. Filter criteria is based on:

- Destination address

- IPv4 or IPv6 traffic type (such as TCP and UDP traffic)

Implementing the feature requires:

- RADIUS authentication using the 802.1X, web-based authentication, or MAC authentication available on the switch to provide client authentication services

- Configuring one or more ACLs on a RADIUS server (instead of the switch), and assigning each ACL to the username/password pair or MAC address of the client(s) you want the ACLs to support

Using RADIUS to dynamically apply ACLs to clients on edge ports enables the switch to filter IP traffic coming from outside the network, thus removing unwanted IP traffic as soon as possible and helping to improve system performance. Also, applying RADIUS-assigned ACLs to the network edge is likely to be less complex than configuring static port and VLAN-based ACLs in the network core to filter unwanted IP traffic that could have been filtered at the edge.

> **NOTE**
>
> A RADIUS-assigned ACL filters inbound IP traffic on a given port from the client whose authentication triggered the ACL assignment to the port.
>
> A RADIUS-assigned ACL can be applied regardless of whether IP traffic on the port is already being filtered by other, static ACLs that are already assigned. The following table lists the supported per-port ACL assignment capacity (Subject to resource availability on the switch. For more information, see the appendix titled "Monitoring Resources" in the latest management and configuration guide.).
>
> ACLs enhance network security by blocking selected IP traffic, and can serve as one aspect of network security. However, because ACLs do not protect from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete edge security solution.
>
> Depending on the ACL configuration in the RADIUS server, the ACLs described in this section filter either IPv4 traffic only or both IPv4 and IPv6 traffic. These ACLs do not filter non-IP traffic such as AppleTalk and IPX.

The values in the following table are subject to resource availability on the switch. For more information, see the appendix titled "Monitoring Resources" in the latest management and configuration guide for your switch.

Table 24: *Simultaneous ACL activity supported per-port for your switch.*

ACL type	Function	IPv4	IPv6
VACL	Static ACL assignment to filter inbound IP traffic on a specific VLAN.	1 in 1 out	1
Port ACL	Static ACL assignment to filter inbound IP traffic on a specific port.	1	1
RADIUS-assigned ACL	Dynamic ACL assignment to filter inbound IP traffic from a specific client on a given port.	1-32[1]	1-32[1]

Table Continued

ACL type	Function	IPv4	IPv6
RACL (IPv4 only)	static ACL assignment to filter routed IPv4 traffic entering or leaving the switch on a specific VLAN	1 in 1 out	n/a
Connection-rate ACL	Static ACL assignment for virus-throttling on a specific port	1	n/a

[1] One per authenticated client, up to a maximum of 32 clients per-port for 802.1X, web-based authentication, and MAC-Authentication methods combined.

ACLs enhance network security by blocking selected IP traffic, and can serve as one aspect of network security. However, because ACLs do not protect from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete edge security solution.

Depending on the ACL configuration in the RADIUS server, the ACLs described in this section filter either IPv4 traffic only or both IPv4 and IPv6 traffic. These ACLs do not filter non-IP traffic such as AppleTalk and IPX.

RADIUS filter-id

IP traffic filter rules, also known as IP ACLs, provide a user access policy that defines what IP traffic from the user is permitted. IP ACLs can be specified in two ways:

- By using the filter-id attribute that gives the ID of a pre-defined ACL. A filter-id is an alphabetic-string identifier, or name, corresponding to an IP ACL that is pre-configured on the access-control device.

- By using the NAS-filter-rule attribute which explicitly defines a set of filter rules.

Filter-id attributes and NAS-Filter-Rule attributes may be intermixed in the RADIUS user entry. Filter-id attributes are expanded as they are read so they are added to the ACL in the correct order.

> **NOTE**
> This feature does not modify any existing commands. CLI `show` commands currently display the applied RADIUS defined ACL rules. ACL rules specified by a filter-id attribute are expanded and displayed as if they were NAS-Filter-Rule entries. The list of rules will be a snapshot of the CLI ACL at the time of authentication. Updates to the ACL are not applied until the client reauthenticates.

A filter-id name may refer to an IPv4 ACL, an IPv6 ACL, or both. ACLs for both families are checked and expanded if found. All other ACL types, including MAC and router ACLs, are ignored when processing filter-id attributes. Any number of filter-id attributes may be specified subject to length limitations of a RADIUS packet. The limit for all platforms is 100 ACEs per client ACL.

> **NOTE**
> RADIUS ACL rules do not support source IP or source L4 port qualifiers. If any source IP or source L4 port qualifiers are found in the CLI ACL, the client will fail authentication and an error will be logged.
>
> CLI ACLs include an optional `log` keyword that captures rule hits for debugging. No logging for ACL rules that are applied via filter-id is available. However, all rules from ACLs have an implicit `cnt` keyword which allows the administrator to see the hit count for each rule.

RADIUS user entry

```
NAS-Filter-Rule += "permit in 10 from any to any cnt",
Filter-ID += "104",
```

```
NAS-Filter-Rule += "permit in 30 from any to any cnt",
Filter-ID += "106",
NAS-Filter-Rule += "permit in 55 from any to any cnt",
Filter-ID += "146",
NAS-Filter-Rule += "permit in 70 from any to any cnt",
```

Forcing reauthentication

Syntax:

```
aaa port-access < authenticator | mac-based | web-based ><port-list> reauthenticate
```

A manager may force a reauthentication by using this command.

> **NOTE**
> RADIUS Filter-Rule entries are only allowed to contain IPv6 addresses if the `hp-nas-rules-ipv6` attribute is set. This does not apply to filter-id ACLs. If there is an IPv6 ACL of the name given, it will be applied even if `hp-nas-rules-ipv6` is not set.

`show access-list radius`

Syntax:

```
show running config
```

System configuration for `show running config`

```
ip access-list extended "104"
10 permit 20 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255 log
exit
```

```
ip access-list extended "146"
10 permit 64 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
exit
```

```
ipv6 access-list "106"
10 permit 40 ::/0 ::/0 log
exit
```

```
ipv6 access-list "146"
10 permit 66 ::/0 ::/0
exit
```

`show access-list` (NAS rule) and (filter-id)

> **NOTE**
> There is a legacy attribute named `hp-nas-filter-rule` that was in use before the `nas-filter-rule` was standardized in RFC 4849. Switches still support the `hp-nas-filter-rule` for backwards compatibility, but this rule should not be mixed with the newer `nas-filter-rule` or `filter-id` attributes. With mixed ACEs, rules will not be applied in the order listed, which may block traffic that should be permitted or may permit traffic that should be blocked. No error message is produced to inform the user that mixing current and legacy attributes will lead to unexpected results.

Syntax:

```
show access-list radius
```

```
show access-list radius (NAS rule)

Radius-configured Port-based ACL for
Port 1/1, Client - 24BE05-76DA40

IPv6 ACLs enabled (HP-Nas-Rules-Ipv6): FALSE
permit in 10 from any to any cnt
Packet Hit Counter 0
permit in 20 from any to 0.0.0.0 255.255.255.255 cnt
(IP ACL 104, rule 10)
Packet Hit Counter 0
permit in 30 from any to any cnt
Packet Hit Counter 0
permit in 40 from any to ::/0 cnt
(IPv6 ACL 106, rule 10)
Packet Hit Counter 0
permit in 55 from any to any cnt
Packet Hit Counter 0
permit in 64 from any to 0.0.0.0 255.255.255.255 cnt
(IP ACL 146, rule 10)
Packet Hit Counter 0
permit in 66 from any to ::/0 cnt
(IPv6 ACL 146, rule 10)
Packet Hit Counter 0
permit in 70 from any to any cnt
Packet Hit Counter 0
```

NOTE: The output shows IPv6 rules with a prefix of IPv6 and shows IPv4 rules with a prefix of IP.

Log messages

Event	Message
dca_filter_id_match_not_found – This event is logged when the ACL name given in a filter-id attribute does not match any existing 'ip' or 'ipv6' access-list.	Authentication failed for client <mac> on port <port>: unknown ACL name in attribute filter-id.
rmon_dca_acl_has_source_qualifier – This event is logged when the ACL given in a filter-id attribute contains an ACE that has a source IP address or source tcp/udp port qualifier.	Authentication failed for client <mac> on port <port>: the ACL specified by the filter-id attribute contains a source address or application port qualifier.

Event message

```
W 10/20/14 15:26:17 03214 dca: Authentication failed for client 0025618D7920 on
port 1: unknown ACL name in attribute filter-id.

W 10/20/14 15:26:17 03215 dca: Authentication failed for client 0025618D7920 on
port 1: the ACL specified by the filter-id attribute contains a source address or
application port qualifier.
```

Contrasting RADIUS-assigned and static ACLs

The following table highlights several key differences between the static ACLs configurable on switch VLANs and ports, and the dynamic ACLs that can be assigned by a RADIUS server to filter IP traffic from individual clients.

Table 25: *Contrasting dynamic (RADIUS-assigned) and static ACLs*

RADIUS-assigned ACLs	Static port and VLAN ACLs
Configured in client accounts on a RADIUS server.	Configured on switch ports and VLANs.
Designed for use on the edge of the network where filtering of IP traffic entering the switch from individual, authenticated clients is most important and where clients with differing access requirements are likely to use the same port.	Designed for use where the filtering needs focus on static configurations covering: • switched IP traffic entering from multiple authenticated or unauthenticated sources (VACLs or static port ACLs) • routed IPv4 traffic (RACLs) • IP traffic from multiple sources and having a destination on the switch itself
Implementation requires client authentication.	Client authentication not a factor.
Identified by the credentials (username/password pair or the MAC address) of the specific client the ACL is intended to service.	Identified by a number in the range of 1-199 or an alphanumeric name.
Supports dynamic assignment to filter only the IP traffic entering the switch from an authenticated client on the port to which the client is connected. (IPv6 traffic can be switched; IPv4 traffic can be routed or switched. For either IP traffic family, includes traffic having a DA on the switch itself.)	Supports static assignments to filter: • switched IPv6 traffic entering the switch • switched or routed IPv4 traffic entering the switch, or routed IPv4 traffic leaving the switch.
When the authenticated client session ends, the switch removes the RADIUS-assigned ACL from the client port.	Remains statically assigned to the port or VLAN.

Table Continued

RADIUS-assigned ACLs	Static port and VLAN ACLs

Allows one RADIUS-assigned ACL per authenticated client on a port. (Each such ACL filters traffic from a different, authenticated client.)Note: The switch provides ample resources for supporting RADIUS-assigned ACLs and other features. However, the actual number of ACLs supported depends on the switch current feature configuration and the related resource requirements. For more information, see the appendix titled "Monitoring Resources" in the management and configuration guide for your switch.	Simultaneously supports all of the following static assignments affecting a given port:
	
	• IPv4 traffic:
	
	  ◦ inbound RACL
	
	  ◦ outbound RACL
	
	  ◦ VACL
	
	  ◦ static port ACL
	
	• IPv6 traffic:
	
	  ◦ VACL
	
	  ◦ static port ACL

Supports IPv6 ACLs and IPv4 extended ACLs. "IPv6 Access Control Lists (ACLs)" in the *IPv6 Configuration Guide* for your switch.	Supports IPv6 ACLs and standard, extended, and connection-rate IPv4 ACLs
A given RADIUS-assigned ACL operates on a port to filter only the IP traffic entering the switch from the authenticated client corresponding to that ACL, and does not filter IP traffic inbound from other authenticated clients. (The traffic source is not a configurable setting.)	An RACL applied to inbound traffic on a VLAN filters routed IPv4 traffic entering the switch through a port on that VLAN, as well as any inbound traffic having a DA on the switch itself. An RACL can be applied to outbound IPv4 traffic on a VLAN to filters routed IPv4 traffic leaving the switch through a port on that VLAN (and includes routed IPv4 traffic generated by the switch itself). A VACL can be applied on a VLAN to filter either IPv4 or IPv6 traffic entering the switch through a port on that VLAN.A static port ACL can be applied on a port to filters either IPv4 or IPv6 traffic entering the switch through that port.
Requires client authentication by a RADIUS server configured to dynamically assign an ACL to a client on a switch port, based on client credentials.	No client authentication requirement.
ACEs allow a counter (cnt) option that causes a counter to increment when there is a packet match.	The show statistics command includes options for displaying the packet match count, see **Monitoring static ACL performance** on page 455. Also, ACEs allow a log option that generates a log message whenever there is a packet match.

CAUTION

Regarding the Use of IPv4 Source Routing:

IPv4 source routing is enabled by default on the switch and can be used to override IPv4 ACLs. For this reason, if you are using IPv4 ACLs to enhance network security, the recommended action is to use the `no ip source-route` command to disable source routing on the switch. (If source routing is disabled in the running-config file, the `show running` command includes "`no ip source-route`" in the running-config file listing.)

How a RADIUS server applies a RADIUS-assigned ACL to a client on a switch port

A RADIUS-assigned ACL configured on a RADIUS server is identified and invoked by the unique credentials (username/password pair or a client MAC address) of the specific client the ACL is intended to service. Where the username/password pair is the selection criteria, the corresponding ACL can also be used for a group of clients that all require the same ACL policy and use the same username/password pair. Where the client MAC address is the selection criteria, only the client having that MAC address can use the corresponding ACL. When a RADIUS server authenticates a client, it also assigns the ACL configured with that client's credentials to the client's port. The ACL then filters the client's inbound IP traffic and denies (drops) any such traffic that is not explicitly permitted by the ACL.

• If the filter rule used for a RADIUS-based ACL is one of the options that specifies only IPv4 traffic, then the ACL will implicitly deny any inbound IPv6 traffic from the authenticated client.

• If the filter rule used for a RADIUS-based ACL is the option for specifying both IPv4 and IPv6 traffic, then the ACL filter both IP traffic types according to the ACEs included in the RADIUS-assigned ACL.

When the client session ends, the switch removes the RADIUS-assigned ACL from the client port.

NOTE

Implicit Deny

Every RADIUS-assigned ACL ends with an implicit `deny in` ACE for both IPv4 and IPv6 traffic. This implicit ACE denies any IP traffic that is not specifically permitted. To override this default, configure an explicit `permit in ip from any to any` as the ACL's last explicit ACE.

Multiple clients in a RADIUS-assigned ACL environment

Where multiple clients are authenticated on the same port, if any of the clients has a RADIUS-assigned ACL, then all of the authenticated clients on the port must have a RADIUS-assigned ACL. In this case, the switch drops the IP traffic from any authenticated client that does not have a RADIUS-assigned ACL, and deauthenticates that client.

Multiple clients sharing the same RADIUS-assigned ACL

When multiple clients supported by the same RADIUS server use the same credentials, they will all be serviced by different instances of the same ACL. (The actual IP traffic inbound from any client on the switch carries a source MAC address unique to that client. The RADIUS-assigned ACL uses this MAC address to identify the traffic to be filtered.)

Effect of multiple ACL application types on an interface

The switch allows simultaneous use of all supported ACL application types on an interface. Thus, a static ACL assigned to an interface filters authenticated client traffic, regardless of whether a RADIUS-assigned ACL is also filtering the client's traffic. For more information, see **Multiple ACLs on an interface** on page 401.

General ACL features, planning, and configuration

These steps suggest a process for using RADIUS-assigned ACLs to establish access policies for client IP traffic.

Procedure

1. Determine the polices you want to enforce for authenticated client traffic inbound on the switch.

2. Plan ACLs to execute traffic policies:

   a. Apply ACLs on a per-client basis where individual clients need different traffic policies or where each client must have a different username/password pair or will authenticate using MAC authentication.

   b. Apply ACLs on a client group basis where all clients in a given group can use the same traffic policy and the same username/password pair.

3. Configure the ACLs on a RADIUS server accessible to the intended clients.

4. Configure the switch to use the desired RADIUS server and to support the desired client authentication scheme. Options include 802.1X, web-based authentication, or MAC authentication. (Note that the switch supports the option of simultaneously using 802.1X with either web-based or MAC authentication.)

5. Test client access on the network to ensure that your RADIUS-assigned ACL application is properly enforcing your policies.

For further information common to all IPv4 or IPv6 ACL applications, see the IPv4 configuration guide or IPv6 configuration guide for your switch.

The packet-filtering process

Packet-filtering in an applied ACL is sequential, from the first ACE in the ACL to the implicit `deny any any` following the last explicit ACE. This operation is the same regardless of whether the ACL is applied dynamically from a RADIUS server or statically in the switch configuration.

> **CAUTION**
> ACLs can enhance network security by blocking selected IP traffic, and can serve as one aspect of maintaining network security. However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete security solution.

> **NOTE**
> If a RADIUS-assigned ACL permits an authenticated client's inbound IP packet, but the client port is also configured with a static port ACL and/or belongs to a VLAN for which there is an inbound, VLAN-based ACL configured on the switch, then the packet will also be filtered by these other ACLs. If there is a match with a deny ACE in any of these ACLs, the switch drops the packet.

Operating rules for RADIUS-assigned ACLs

- Relating a client to a RADIUS-assigned ACLA RADIUS-assigned ACL for a particular client must be configured in the RADIUS server under the authentication credentials the server should expect for that client. If the client must authenticate using 802.1X and/or web-based authentication, the username/password pair forms the credential set. If authentication is through MAC Authentication, then the client MAC address forms the credential set. See **Configuring an ACL in a RADIUS server** on page 301.

- Multiple clients using the same username/password pairMultiple clients using the same username/password pair will use duplicate instances of the same ACL.

- Limits for ACEs in RADIUS-assigned ACLsThe switch supports up to 80 characters in a single ACE. Exceeding this limit causes the related client authentication to fail.

- Effect of other, statically configured ACLsSuppose that port B1 belongs to VLAN "Y" and has a RADIUS-assigned ACL to filter inbound traffic from an authenticated client. Port B1 is also configured with IPv4 and IPv6 static port ACLs, and VLAN "Y" is statically configured with IPv4 and IPv6 VACLs.

◦ IP traffic entering the switch on port B1 from the client and having a match with a `deny` ACE configured in any of the ACLs mentioned above will be dropped.

◦ If an inbound RACL was also configured on VLAN "Y", then a `deny` match in the RACL would apply to any inbound, routed IPv4 traffic from the client (and to any inbound, switched traffic having a destination on the switch itself).

◦ If an outbound RACL was also configured on VLAN "Y", then any outbound, routed IPv4 traffic leaving the switch through the port B1 would be filtered by the outbound RACL.

Configuring an ACL in a RADIUS server

The following information provides general guidelines for configuring a RADIUS server to specify RADIUS-assigned ACLs. It also provides an example configuration for a FreeRADIUS server application. To configure services on a specific RADIUS server application, see the documentation provided with that application.

A RADIUS-assigned ACL configuration in a RADIUS server includes the following elements:

• Nas-Filter-Rule attributes — standard and vendor-specific

• ACL configuration, entered in the server, and associated with specific username/password or MAC address criteria, and comprised of ACEs entered in the server

A RADIUS-assigned ACL includes:

• One or more explicit `permit` and/or `deny` ACEs

• An implicit `deny in ip from any to any` ACE automatically applied after the last operator-created ACE

Nas-filter-Rule attribute options

Table 26: *Nas-filter-Rule attribute options*

Service	Control method and operating notes

ACLs Applied to Client Traffic Inbound to the Switch
 	
Assigns a RADIUS-configured ACL to filter inbound packets received from a specific client authenticated on a switch port.
	Standard Attribute: 92
	
	This is the preferred attribute for use in RADIUS-assigned ACLs to configure ACEs to filter IPv4 and IPv6 traffic.
	
	Entry for IPv4-Only ACE To Filter Client Traffic:
	
	Nas-filter-Rule="< permit or deny ACE >"(Standard Attribute 92) For example:
	
	`Nas-filter-Rule=permit in tcp from any to any`
	
	Entries for IPv4/IPv6 ACE To Filter Client Traffic: `HP-Nas-Rules-IPv6 <1 2>` (VSA, where 1=IPv4 and IPv6 traffic, and 2=IPv4-only traffic.)
	
	Nas-filter-Rule="<permit or deny ACE>"(Standard Attribute 92). For example:
	
	`HP-Nas-Rules-IPv6=1`
	`Nas-filter-Rule="permit in tcp from any to any"`
	
	Note: If `HP-Nas-Rules-IPv6` is set to 2 or is not present in the ACL, IPv6 traffic from the client will be dropped.

Set IP Mode Used with the Nas-filter-Rule attribute described above to provide IPv6 traffic-filtering capability in an ACE.	HP-Nas-Rules-IPv6: 63 (Vendor-Specific Attribute) When using the standard attribute (92) described above in a RADIUS-assigned ACL to support both IPv4 and IPv6 traffic inbound from an authenticated client, one instance of this VSA must be included in the ACL. Note that this attribute supports either of the following IP modes for Nas-filter-Rule ACEs: • Both IPv6 and IPv4 traffic • Only IPv4 traffic vendor-specific ID: 11VSA: 63 (string=HP-Nas-Rules-IPv6) • IPv6 and IPv4 ACLs: integer = 1(Using this option causes the ACL to filter both IPv4 and IPv6 traffic.) • IPv4-only ACLs: integer=2 (Using this option causes the ACL to drop any IPv6 traffic received from the authenticated client.) Setting: `HP-Nas-Rules-IPv6=< 1 2 > Nas-filter-Rule "< permit or deny ACE >"` Note: When the configured integer option is "1", the `any` keyword used as a destination applies to both IPv4 and IPv6 destinations for the selected traffic type (such as Telnet). Thus, if you want the IPv4 and IPv6 versions

Table Continued

Service	Control method and operating notes
	of the selected traffic type to both go to their respective "any" destinations, then a single ACE is needed for the selected traffic type. For example: ``` HP-Nas-Rules-IPv6=1 Nas-filter-Rule="permit in tcp from any to any 23" ``` However, if you do not want both the IPv4 and IPv6 traffic of the selected type to go to their respective "any" destinations, then two ACEs with explicit destination addresses are needed. In this case, do one of the following: • Use `0.0.0.0/0` in one ACE to specify the "any" destination for IPv4 traffic, and use a specific IPv6 address for the destination in the other ACE. • Use `::/0` in one ACE to specify the "any" destination for IPv6 traffic, and use a specific IPv4 address for the destination in the other ACE. For example, if you want to allow the IPv4 Telnet traffic from a client to go to any destination, but you want the IPv6 Telnet traffic from the same client to go only to a specific address or group of addresses, you will need to distinguish the separate destinations. This is done by using explicit addresses for the "any" destinations. For example: ``` HP-Nas-Rules-IPv6=1 Nas-filter-Rule="deny in tcp from any to 0.0.0.0/0 23" Nas-filter-Rule="deny in tcp from any to fe80::b1 23" ``` The above example sends IPv4 Telnet traffic to its "any" destination, but allows IPv6 Telnet traffic only to fe80::b1 23. To reverse this example, you would configure ACEs such as the following: ``` HP-Nas-Rules-IPv6=1 Nas-filter-Rule="deny in tcp from any to 10.10.10.1 23" Nas-filter-Rule="deny in tcp from any to ::/0 23" ``` In cases where you do not want the selected traffic type for either IPv4 or IPv6 to go to the "any" destination, you must use two ACEs to specify the destination addresses. For example: ``` HP-Nas-Rules-IPv6=1 Nas-filter-Rule="deny in tcp from any to 10.10.10.1 23" Nas-filter-Rule="deny in tcp from any to fe80::23 23" ``` To use the IPv6 VSA while allowing only IPv4 traffic to be filtered, you would use a configuration such as the following: ``` HP-Nas-Rules-IPv6=2 Nas-filter-Rule="permit in tcp from any to any" ```
IPv4-only ACLs applied to client traffic inbound to the switch. (Assigns a RADIUS-configured IPv4 ACL to filter inbound IPv4 packets received from a specific	HP-Nas-filter-Rule (Vendor-Specific Attribute): 61 This attribute is maintained for legacy purposes. However, for new or updated configurations (and any configurations supporting IPv6 traffic filtering) Hewlett Packard Enterprise recommends using the Standard Attribute (92) described earlier in this table instead of the HP-Nas-filter-

Service	Control method and operating notes
client authenticated on a switch port.)	Rule attribute described here.vendor-specific ID: 11VSA: 61 (string=HP-Nas-filter-RuleSetting: HP-Nas-filter-Rule="< permit or deny ACE >"Note: An ACL applying this VSA to inbound traffic from an authenticated client drops any IPv6 traffic from the client.

ACE syntax in RADIUS servers

The following information describes ACE syntax configuration options in a RADIUS server.

ACE syntax (standard attribute-92)	`Nas-filter-Rule = " <permit|deny> in <ip|ip-protocol-value> from any to <any|host|<ip-addr>|ipv4-addr/mask|IPv6-address/prefix> [tcp/udp-port|tcp/udp-port range>] [cnt] "`
ACE syntax (legacy VSA-61)	`HP-Nas-filter-Rule= " <permit|deny> in <ip|ip-protocol-value> from any to <any|host|<ip-addr>|ipv4-addr/mask|IPv6-address/prefix> [<tcp/udp-port|tcp/udp-port range> ] [cnt"]`
	`Nas-filter-Rule=` Standard attribute for filtering inbound IPv4 traffic from an authenticated client. When used without the VSA option (below) for filtering inbound IPv6 traffic from the client, drops the IPv6 traffic. See also **Nas-filter-Rule attribute options** on page 302.
	`[HP-Nas-Rules-IPv6=|<1|2>]` VSA used in an ACL intended to filter IPv6 traffic. Settings include: • 1: ACE filters both IPv4 and IPv6 traffic. • 2: ACE filters IPv4 traffic and drops IPv6 traffic. • VSA not used: ACE filters IPv4 traffic and drops IPv6 traffic. This VSA must be present in an ACL where the Nas-filter-Rule= attribute is intended to filter inbound IPv6 traffic from an authenticated client. See also **Nas-filter-Rule attribute options** on page 302.

	`HP-Nas-filter-Rule=Legacy`
	
	VSA for filtering inbound IPv4 traffic only from an authenticated client. Drops inbound IPv6 traffic from the client. See also **Nas-filter-Rule attribute options** on page 302.
	
	Must be used to enclose and identify a complete permit or deny ACE syntax statement. For example:
	
	`Nas-filter-Rule="deny in tcp from any to 0.0.0.0/0 23"`

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

	`<permit|deny>` Specifies whether to forward or drop the identified IP traffic type from the authenticated client. (For information on explicitly permitting or denying all inbound IP traffic from an authenticated client, or for implicitly denying all such IP traffic not already permitted or denied, see **Configuration notes**.)
	`in` Required keyword specifying that the ACL applies only to the traffic inbound from the authenticated client.
	`<ip|ip-protocol-value>` Options for specifying the type of traffic to filter.
	`ip` Applies the ACE to all IP traffic from the authenticated client.
	`ip-protocol-value` This option applies the ACE to the type of IP traffic specified by either a protocol number or by tcp, udp, icmp, or (for IPv4-only) igmp. The range of protocol numbers is 0-255. (Protocol numbers are defined in RFC 2780. For a complete listing, see "Protocol Registries" on the Website of the Internet Assigned Numbers Authority at (**http://www.iana.com**). Some examples of protocol numbers include: 1=ICMP 17=UDP 2=IGMP (IPv4 only) 41=IPv6 6=TCP
	`from any` Required keywords specifying the (authenticated) client source. (Note that a RADIUS-assigned ACL assigned to a port filters only the inbound traffic having a source MAC address that matches the MAC address of the client whose authentication invoked the ACL assignment.)
	`to` Required destination keyword.

Table Continued

```
any
```

- Specifies any IPv4 destination address if one of the following is true:

    ◦ the ACE uses the standard attribute (`Nas-filter-Rule`) and the IPv6 VSA (`HP-Nas-Rules-IPv6`) is not included the ACL. For example:

      ```
      Nas-filter-Rule="permit in tcp from any to any 23"
      ```

      ```
      Nas-filter-Rule+="permit in ip from any to
      10.10.10.1/24"
      ```

      ```
      Nas-filter-Rule+="deny in ip from any to any"
      ```

    ◦ the ACE uses the standard attribute (`Nas-filter-Rule`) and the IPv6 VSA (`HP-Nas-Rules-IPv6`) is included in the ACL with an integer setting of `2`. For example, all of the following destinations are for IPv4 traffic:

      ```
      HP-Nas-Rules-IPv6=2
      ```

      ```
      Nas-filter-Rule="permit in tcp from any to any 23"
      ```

      ```
      Nas-filter-Rule+="permit in ip from any to
      10.10.10.1/24"
      ```

      ```
      Nas-filter-Rule+="deny in ip from any to any"
      ```

    ◦ the HP-Nas-Filter-Rule VSA is used instead of either of the above options. For example, all of the following destinations are for IPv4 traffic:

      ```
      HP-Nas-filter-Rule="permit in tcp from any to any 23"
      ```

      ```
      HP-Nas-filter-Rule+="permit in ip from any to
      10.10.10.1/24"
      ```

      ```
      HP-Nas-filter-Rule+="deny in ip from any to any"
      ```

- Specifies any IPv4 or IPv6 destination address if the ACL uses the HP-Nas-Rules-IPv6 VSA with an integer setting of `1`. See **Nas-filter-Rule attribute options** on page 302. For example, the `any` destinations in the following ACL apply to both IPv4 and IPv6 traffic:

  ```
  HP-Nas-Rules-IPv6=1Nas-filter-Rule="permit in tcp from
  any to any 23"
  ```

  ```
  Nas-filter-Rule+="permit in ip from any to
  10.10.10.1/24"
  ```

  ```
  Nas-filter-Rule+="permit in ip from any to
  fe80::d1:1/120"
  ```

  ```
  Nas-filter-Rule+="deny in ip from any to any"
  ```

Table Continued

	`host <ipv4-addr>` Specifies a single destination IPv4 address.
	`<ipv4-addr/<mask>` Specifies a series of contiguous destination addresses or all destination addresses in a subnet. The `<mask>` is CIDR notation for the number of leftmost bits in a packet's destination IPv4 address that must match the corresponding bits in the destination IPv4 address listed in the ACE. For example, a destination of 10.100.17.1/24 in the ACE means that a match occurs when an inbound packet (of the designated IPv4 type) from the authenticated client has a destination IPv4 address where the first three octets are 10.100.17. (The fourth octet is a wildcard, and can be any value up to 255.)
	`host <ipv6-addr>` Specifies a single destination IPv6 address. Note: Filtering IPv6 traffic requires the Standard Attribute(Nas-Filter-Rule)with the HP-Nas-Rules-IPv6 VSA set to 1. See **Nas-filter-Rule attribute options** on page 302.
	`<ipv6-addr/<prefix>` Specifies a series of contiguous destination addresses or all destination addresses in a subnet. The `< prefix >` specifies the number of leftmost bits in a packet's destination IPv6 address that must match the corresponding bits in the destination IPv6 address listed in the ACE. For example, a destination of FE80::1b:127/112 in the ACE means that a match occurs when an inbound packet (of the designated IPv6 type) from the authenticated client has a destination IPv6 address where the first 112 are FE80::1b. (The last 16 bits in the address configured in the ACE form a "wildcard", and can be any value from 0 to FFFF.) Also, see `Note`, above.

	`[<tcp/udp-port>|<tcp/udp-port-range>]`
	
	Optional TCP or UDP port specifier. Used when the ACE is intended to filter client TCP or UDP traffic with one or more specific TCP or UDP destination port numbers. You can specify port numbers as individual values and/or ranges. For example, the following ACE shows two ways to deny any UDP traffic from an authenticated client that has a DA of any address and a UDP destination port of 135, 137-139, or 445:
	
	`deny in udp from any to any 135, 137-139, 445`
	
	`deny in 17 from any to any 135, 137-139, 445`

Table Continued

	`[<icmp-type>|icmpv6-type]` Optional ICMP type specifier. This can be either a keyword or an ICMP type number. For a listing of numbers and types, see **Option for IGMP in extended ACLs** on page 433.
	`[cnt]` Optional counter specifier for a RADIUS-assigned ACE. When used, the counter increments each time there is a "match" with the ACE. This option does not require that you configure the switch for RADIUS accounting.

Configuration notes
Explicitly permit IPv4 and IPv6 traffic from an authenticated client

This option for ending a RADIUS-assigned ACL permits all of the client's inbound IPv4 and IPv6 traffic not previously permitted or denied.

```
Nas-filter-Rule += permit in ip from any to any
HP-Nas-Rules-IPv6=1
```

See **Nas-filter-Rule attribute options** on page 302 for information on the above attributes.

Explicitly permit only the IPv4 traffic from an authenticated client

Any of the following three options for ending a RADIUS-assigned ACL explicitly permit all of the client's inbound IPv4 traffic not previously permitted or denied. These options also deny any of the client's IPv6 traffic not previously permitted or denied.

```
Nas-filter-Rule += permit in ip from any to any
```

(Using this attribute to permit IPv4 traffic from the client while denying any IPv6 traffic from the client assumes that `HP-Nas-Rules-IPv6=1`does not exist elsewhere in the ACL. See **Nas-filter-Rule attribute options** on page 302 for more on `HP-Nas-Rules-IPv6`.)

- `HP-Nas-Filter-Rule += permit in ip from any to any`

- `Nas-filter-Rule += permit in ip from any to any`
  `HP-Nas-Rules-IPv6=2`

Implicitly denying any IP traffic

For any packet being filtered by a RADIUS-assigned ACL, there will always be a match. That is, any packet that does not have a match with an explicit permit or deny ACE in the list will match with the implicit `deny any any` ACE automatically included at the end of the ACL. That is, a RADIUS-assigned ACL includes an implicit `deny in ip from any to any` ACE at the end of the ACL to deny any IPv4 and IPv6 traffic not previously permitted or denied.

Monitoring shared resources

Currently active RADIUS-based authentication sessions using RADIUS-assigned ACLs share internal switch resources with several other features. The switch provides sufficient resources for all features. However, if the internal resources do become fully subscribed, new RADIUS-based sessions using RADIUS-assigned ACLs cannot be authenticated until the necessary resources are released from other applications.

- For information on determining the current resource availability and usage, see "Monitoring Resources" in the management and configuration guide for your switch.

- For a summary of ACL resource limits, see the topics covering scalability in the latest management and configuration guide for your switch.

Event log messages

See the event log message reference guide for information about event log messages.

Causes of client deauthentication immediately after authenticating

- ACE formatted incorrectly in the RADIUS server examples:

  - `from`, `any`, or `to` keyword missing.

  - An IPv4 or IPv6 protocol number in the ACE exceeds 255.

  - An optional UDP or TCP port number is invalid, or a UDP/TCP port number is specified when the protocol is neither UDP or TCP.

  - An ACE in the ACL for a given authenticated client exceeds 80 characters.

- A RADIUS-assigned ACL limit has been exceeded:

  - The TCP/UDP port-range quantity of 14 per slot or port group has been exceeded.

  - The rule limit has been exceeded.

- An IPv6 ACE has been received on a port and either the `HP-Nas-Rules-IPv6` attribute is missing or `HP-Nas-Rules-IPv6=2` is configured. See **Nas-Filter-Rule-Options** for more on this attribute.

Open Authentication allows a device, such as an IP Phone, to have network access before the device is authenticated. Open Authentication is triggered when a mac-based client is connected to an Aruba switch before being authenticated by the RADIUS Server. To provide network connectivity for devices, they must be assigned a VLAN. Two new VLANs are created for Open Authentication functionality, one for voice traffic and one for data traffic. Open Authentication VLANs can be configured on the switch individually or within a user-role. Devices that can be connected to the switch without authentication are divided into two categories:

• Devices that send voice traffic.

• Devices that send data traffic.

NOTE

Either one of open authentication VLAN (voice and/or data) or open authentication user-role can be configured for a port. However, both a VLAN and user-role cannot coexist for an interface. Initial traffic on the port is restricted only by ACLs configured for the port or for VLANs or ACLs in the user-role.

Impact of Open Authentication on existing features

Unauthenticated devices

Configuring open authentication VLAN will change the behavior of unauthenticated devices. Normally, authentication-enabled ports will not provide unauthenticated client any network access until the device is authenticated by the RADIUS Server. With open authentication VLAN configured, the client will be put in open authentication VLAN until the RADIUS Server authenticates the device.

Unauthenticated clients will be placed into the VLAN specified in the open authentication command string. After authentication by the RADIUS server, the client will be placed into the VLAN specified by the RADIUS authentication command string or as specified in the RADIUS authentication accept string.

LLDP-Bypass

When LLDP-bypass is enabled on the switch, Aruba APs are not authenticated therefore open authentication VLAN is not applicable.

ACLs applied on an Interface

If an ACL rule is applied on an interface which is part of an open authentication VLAN, traffic coming through that interface will be affected. Traffic will be affected based on the rule in the ACL. For more information, see the *Access Security Guide* for your switch.

ACLs applied on a VLAN

If an ACL rule is applied on an open authentication VLAN, traffic entering that VLAN will be affected. Traffic will be affected based on the rule in the ACL. For more information, see the *Access Security Guide* for your switch.

Rate-limiting on an interface

If the traffic is rate-limited on an interface as part of an open authentication VLAN, the traffic will be impacted. The traffic will be affected based on the rule in the rate-limiting configuration command. For more information, see the *Management and Configuration Guide* for your switch.

Authenticated or rejected clients

Clients which are either authenticated or rejected by the RADIUS server are given different VLANs. These clients are moved from open authentication to new VLANs based on authentication by the RADIUS Server.

MAC pinning

Clients whose MAC addresses are pinned and have undergone authentication will always be treated as authenticated. Open authentication VLAN is not applicable in this scenario.

Effect of RADIUS tracking on open authentication

If RADIUS tracking is enabled and no RADIUS server is available for authentication, the port will be changed from an open authentication VLAN to a critical VLAN. The time taken to move from open authentication VLAN to Critical VLAN depends on the time it takes for RADIUS tracker to inform the subsystem.

Impact of disabling open authentication feature

When a device is in an open authentication VLAN and the open authentication feature is disabled at the switch, the device will be moved to the PVID. All tagged traffic to that device will be dropped while untagged traffic will be assigned to the PVID.

Restrictions

This feature will not support more than one tagged or untagged VLAN membership either through direct VLAN configuration or through user-roles.
This feature is not applicable for authentication methods other than mac-based.
This feature is not available to be configured from WebUI, Menu, or REST.

aaa port-access open-auth voice-vlan

Syntax

```
aaa port-access <port> open-auth  voice-vlan <VLAN-ID>

no aaa port-access <port> open-auth voice-vlan
```

Description

A voice VLAN is configured on the switch using the port number and open authentication. "A voice VLAN is configured on a port for open authentication.

The `no` form of this command disables the open authentication on that port for that VLAN

Command context

```
manager
```

Parameters

<port>

Specifies the port number of the device.

<VLAN-ID>

Specifies the ID of the VLAN being mapped to the device.

aaa port-access open-auth data-vlan

Syntax

```
aaa port-access <port> open-auth data-vlan <VLAN-ID>

no aaa port-access <port> open-auth data-vlan
```

Description

A data VLAN is configured on a port for open authentication.

The `no` form of this command disables open authentication on that port for that VLAN

Command context

```
manager
```

Parameters

<port>

Specifies the port number of the device.

<VLAN-ID>

Specifies the ID of the VLAN being mapped to the device.

aaa port-access open-auth user-role

Syntax

```
aaa port-access <port> open-auth user-role <ROLE-NAME>

no aaa port-access <port> open-auth user-role <ROLE-NAME>
```

Description

Assign an open-auth user-role to a port by role-name instead of assigning individual VLANs. Role-names can be created by the user.

The `no` form of this command disassociates the open-auth user-role from the port.

Command context

```
manager
```

Parameters

<port>

Specifies the port number of the device.

<ROLE-NAME>

The role-name is created when the user role is named and configured to contain a VLAN.

show port-access clients

Syntax

```
show port-access clients [detail]
```

Description

Show clients with OpenAuth access for both voice and data VLANs.

Command context

config

Example

Show the port-access clients on the switch.

```
switch(config)# show port-access clients

  Port Access Client Status
  Port  Client Name  MAC Address    IP Address     User Role         Type     VLAN
  ----- -----------  ------------ ---------------  -----------------  -------- ----
  23                 005056-bd1374   n/a                               MAC      20
```

Example

Show the port-access clients on the switch in detail.

```
switch(config)# show port-access clients detailed

 Port Access Client Status Detail

  Client Base Details :
   Port          : 23                 Authentication Type : mac-based
   Client Status : open-auth          Session Time        : 11 seconds
   Client Name   :                    Session Timeout     : 7 seconds
   MAC Address   : 005056-bd1374
   IP            : n/a

  Access Policy Details :
   COS Map       : Not Defined        In Limit Kbps       : Not Set
   Untagged VLAN : 20                 Out Limit Kbps      : Not Set
   Tagged VLANs  : No Tagged VLANs
   Port Mode     : 1000FDx
   RADIUS ACL List : No Radius ACL List

Port Access Client Status Detail
Client Base Details :
   Port          : 2                  Authentication Type : 802.1x
   Client Status : OpenAuth           Session Time        : 0 seconds
   Client name   :                    Session Timeout     : 0 seconds
   MAC Address   : 005056-bd37e1
   IP            : n/a

  Access Policy Details :
   COS Map       : Not Defined        In Limit Kbps       : Not Set
   Untagged VLAN : 10                 Out Limit Kbps      : Not Set
   Tagged VLANs  : 0
   Port Mode     : 1000FDx
   RADIUS ACL List : No Radius ACL List
```

Example

Show the port-access clients when a user-role is configured.

```
switch(config)# show port-access clients

 Port Access Client Status

  Port    Client Name  MAC Address      IP Address     User Role       Type      VLAN
  -----   -----------  ------------    ------------   --------------  --------   ----
```

| 23 | 005056-bd1374 | n/a | OpenAuth | MAC |

Critical authentication

Critical authentication provides alternative VLAN authentication for a client when the remote authentication server is not reachable. When remote authentication is not available, the client is placed in Critical VLAN instead of being blocked from access. A Critical VLAN can be configured per-port for both voice and data traffic and can be applied to Mac-based or 802.1x authentication. A critical user-role is configured which accepts the client when authentication fails due to an unreachable authentication server.

Critical voice (tagged) VLAN

When the remote authentication server is not reachable, clients sending tagged traffic will be placed in a Critical (tagged) VLAN.

When a client is sending a MED device advertisement (such as an IP phone) using CDP, the switch sends the VLAN information in the "TIA TR-41 Committee - Network Policy" of the LLDP packet with auto-VLAN-negotiation capability. The MED device uses the VLAN to tag traffic. A Critical VLAN can be tagged, untagged or a combination of both. To enable this VLAN advertisement in LLDP, the Critical VLAN must be a voice VLAN. A Critical (tagged) VLAN is called a Critical voice VLAN.

There are two ways to configure a Critical voice VLAN:

- Directly assign the VLAN.

- Assign a user-role containing the tagged VLAN as a critical-role.

Critical data (untagged) VLAN

For clients sending untagged traffic, if the RADIUS server is unreachable, the client is placed in a Critical VLAN.

There are two ways to configure a Critical Data VLAN:

1. Directly assigning the VLAN using the command `aaa port-access <port> critical-auth data-vlan <VLAN-ID>`.

2. Assign a user-role containing untagged VLAN as critical-role using the command `aaa port-access <port> critical-auth user-role <ROLE-NAME>`.

Restrictions

- Aruba switches will only support one Critical VLAN.

- Either a Critical VLAN or a Critical user-role can be configured for a port. However, both VLAN and user-role cannot coexist for a port.

- This feature is configurable per-port and only applies to RADIUS-based authentication mechanisms.

- Web-based authentication is applicable to web-aware clients. However, if the port connected to the client (either phone or PC behind phone) has web-based authentication enabled, the switch will initiate web-based authentication for all devices.

Examples of Behaviors

Unreachable RADIUS server

A device, such as an IP phone or PC, goes to a RADIUS server and is unable to authentication. The authentication of the device is then applied to a Critical VLAN or a critical user-role.

```
Stack(config)# show port-ac clients

 Port Access Client Status

 Port    Client Name    MAC Address      IPAddress    User Role Type VLAN
 -----  ------------  ------------- ---------- --------- ---------
 1/1         b4b0178db6a2 b4b017-8db6a2          n/a     critical   MAC
```

Tagged critical role

When a critical-role has tagged VID and configured as voice, the port-connected to the MED device (IP phone) will be a tagged member of the voice VLAN. The switch will only support one tagged VLAN as critical. For clients with auto-VLAN-negotiation capabilities (MED devices), the switch sends the VLAN information in the "TIA TR-41 Committee – Network Policy" of the LLDP packet. If the MED device advertising is using CDP, the switch sends the VLAN information in the "VOIP VLAN Reply" field of CDP. The MED devices will use that VLAN to tag their traffic. To enable this VLAN advertisement in LLDP, we need to make the Critical VLAN as 'voice' VLAN.

For clients which send tagged traffic, switch can put them in Critical Tagged-VLAN:

1. Create tagged VLAN.

2. Make the tagged VLAN voice.

3. Create a user-role.

4. Make the tagged VLAN a member of the user-role.

5. Make the user-role a critical user-role with the command `aaa authorization user-role name <CRITICAL-VOICE> vlan-id-tagged <ID>`

```
Stack(config)# show vlan 10
       VLAN ID : 10
  Name : VLAN10
  Status : Port-based
  Voice : Yes
  Jumbo : No
  Private VLAN : none
  Associated Primary VID : none
  Associated Secondary VIDs : none

  Port Information Mode      Unknown VLAN Status
  ---------------- -------- ------------ ----------
  1/1              MACAUTH  Learn        Up

  Overridden Port VLAN configuration


  ------ ------------
  1/1    MACAUTH
```

```
show lldp info remote-device
```

If we execute show LLDP info remote, we can see that the phone has learned which tag to apply for traffic. if we run show lldp info remote, the results are as follows:

```
Stack(config)# show lldp info remote-device 1/1

LLDP Remote Device Information Detail

  Local Port   : 1/1
  ChassisType  : network-address
  ChassisId    : 0.0.0.0
  PortType     : mac-address
  PortId       : b4 b0 17 8d b6 a2
  SysName      : AVX8DB6A2
  PortDescr    :
  Pvid         :

  System Capabilities Supported  : bridge, telephone
  System Capabilities Enabled    : bridge

  Remote Management Address
     Type    : ipv4

  MED Information Detail
    EndpointClass         :Class3
    Media Policy Vlan id   :10
    Media Policy Priority  :6
    Media Policy Dscp      :46
    Media Policy Tagged    :True
    Poe Device Type        :PD
    Power Requested        :2.6 W
    Power Source           :From PSE
```

Run Packet Captures to show the switch advertising which VLAN phone to use or that the phone is advertising which VLAN to use.

Deploying Critical VLAN

When a RADIUS server is configured as an authentication server, the clients are authenticated based on responses from the RADIUS server. The critical authentication feature does not operate unless the RADIUS server does not respond to authentication requests and only on those ports where critical authentication is configured. RADIUS Authentication can be deployed as is or as a module within policy managers such as ClearPass Policy Manager (CPPM).

Procedure

1. **Creating a critical VLAN**

2. **Creating a user-role**

3. **Making a critical-role**

Creating a VLAN for voice traffic.

Creates a critical VLAN for voice traffic.

Procedure

1. Create a voice VLAN.

   > **NOTE**
   >
   > For more information on how to create a VLAN see **VLANs**.

2. Name the new VLAN using the following convention: `Critical_Voice_VLAN-NAME`.

3. Exit

Creating a user-role

Creates a user-role.

Prerequisites

Create a critical VLAN.

Procedure

1. Create a user-role.

   For information about creating user-roles, refer to **Local user-roles**.

2. Make the user-role a voice VLAN by using the command `vlan-id-tagged`.

   For information about the user-roles, refer to **vlan-id-tagged**.

Associating a critical user-role to the critical VLAN

Prerequisites
Create a user-role

Procedure

1. Name the user-role as a critical voice user with the command `aaa authorization user-role name critical-voice`.

2. Enable the user-role by using the command `aaa authorization user-role enable`

3. Associate the port with the user-role critical voice VLAN by using the command `aaa port-access port critical user-role critical-voice`.

aaa port-access critical-auth

Syntax

```
aaa port-access <PORT-LIST> critical-auth {voice-vlan <VLAN-ID> | data-vlan <VLAN-ID>
 | user-role <ROLE-NAME>}

no aaa port-access <PORT-LIST> critical-auth {voice-vlan <VLAN-ID> | data-vlan <VLAN-ID>
 | user-role <ROLE-NAME>}
```

Description

Configures and enables critical authentication for clients due to nonreachable authentication server.

The `no` form of this command disables the critical authentication.

Command context

`manager`

Parameters

<PORT-LIST>

   Specifies the port or list of ports to configure with Critical Authentication.

<VLAN-ID>

   Specifies the IP of the voice or data VLAN being configured with Critical Authentication.

<ROLE-NAME>

   Specifies the role name assigned to the user-role for Critical Authentication.

Restrictions

Critical authentication is only available for MAC-based and 802.1x authentication.

---

show port-access clients

Use the show commands to display Critical Authentication and Open Authentication information and status.

```
switch# show port-access clients

  Port Access Client Status

Port  Client Name    MAC Address     IP Address User Role   Type  VLAN
----- ------------- -------------- ---------- ---------- ----- ----
A1    b4b0178db6a2  b4b017-8db6a2    n/a      critical_role    MAC
A2    b4b0178db6a3  b4b017-8db6a3    n/a      open-auth_role   MAC
```

show port-access authenticator clients

```
switch# show port-access authenticator clients

  Port Access Authenticator Client Status

Port  Client Name    MAC Address     IP Address Session Status
----- ------------- -------------- --------- -------------
A1    b4b0178db6a2  b4b017-8db6a2    n/a      critical
A2    b4b0178db6a3  b4b017-8db6a3    n/a      open-auth
```

show port-access mac-based clients

```
switch# show port-access mac-based clients

  Port Access MAC-Based Client Status

      Port  Client Name   MAC Address    IP Address Session Status
      ----- ------------- ------------- ---------- --------------
      A1    b4b0178db6a2  b4b017-8db6a2  n/a     critical-auth
      A2    b4b0178db6a3  b4b017-8db6a3  n/a     open-auth
```

```
switch# show port-access mac-based clients A1 detailed

      Port Access MAC-Based Client Status Detailed

    Client Base Details:
  Port          : A1
  Session Status : critical auth  Session Time (Sec) : 6
  Username       : client1  MAC Address    : b4b0178db6a2
  IP             : n/a
        ...

switch# show port-access mac-based clients A2 detailed

      Port Access MAC-Based Client Status Detailed

    Client Base Details:
  Port          : A2
  Session Status : open-auth Session Time (Sec) : 6
  Username       : client1   MAC Address   : b4b0178db6a3
  IP             : n/a
           ...
```

show running config

```
switch# show runnig-config

aaa port-access A1 critical-auth voice-vlan 10
aaa port-access A2 critical-auth user-role guest_role
aaa port-access A4 open-auth voice-vlan 10
aaa port-access A5 open-auth user-role guest_role
```

Cached reauthentication

Currently Aruba network switches support primary and fallback authentication mechanism for mac-authentication & dot1x authentication methods. This fallback method will be applied for these authentication methods only if RADIUS is not reachable or down. Following option can be enabled as fallback mechanism.

- Authorized

- Cached Re-Auth

- None

When `authorized` is set, authenticated clients will be allowed as authenticated during reauthentication.

When `cached Re-Auth` is set, authenticated client will be authenticated for the configured cached reauth period or RADIUS server reachability.

`none` is the default option, authenticated client will be de-authenticated.

A CLI is introduced after the fallback mechanism cached reauthentication for 8021x and MAC authentication to authorize. Once cached–reauth has expired, the client will be triggered to authorized if the CLI option has been enabled.

aaa authentication mac-based cached-reauth authorized

Syntax

```
aaa authentication mac-based [chap-radius | peap-mschapv2 | none] cached-reauth authorized
```

Description

When `authorized` is set for MAC-based, authenticated clients, clients are allowed authentication during re-authentication.

Command context

config

show authentication

```
switch # show authentication

Status and Counters - Authentication Information
Authorized enabled as backup for secondary login are preceded by *

  Login Attempts : 3
  Lockout Delay : 0
  Respect Privilege : Disabled
  Bypass Username For Operator and Manager Access : Disabled

                 | Login       Login        Login
  Access Task    | Primary     Server Group Secondary
  -------------- + ----------- ------------ ----------
  Console        | Local                    None
  Telnet         | Local                    None
  Port-Access    | Local                    *Cached
  Webui          | Local                    None
  SSH            | Local                    None
  Web-Auth       | ChapRadius  radius       None
  MAC-Auth       | ChapRadius  radius       *Cached
  SNMP           | Local                    None
  Local-MAC-Auth | Local                    None

                 | Enable      Enable       Enable
  Access Task    | Primary     Server Group Secondary
  -------------- + ----------- ------------ ----------
  Console        | Local                    None
  Telnet         | Local                    None
  Webui          | Local                    None
  SSH            | Local                    None
```

aaa authentication port-access cached-reauth authorized

Syntax

aaa authentication port-access [chap-radius | eap-radius | none] cached-reauth authorized

Description

When `authorized` is set for port-access, clients are allowed authentication during re-authentication.

Command context

config

Configuring a client for `retain-unauth-clients`

A series of steps must be undertaken to configure a client for enforce-cache reauthentication.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Procedure

1. `(config)# aaa port-access mac-based <PORT-LIST>`

   Associates the specified port with the port-access on a MAC-based client.

2. `(config)# no aaa port-access mac-based addr-format [no-delimiter | single-dash | multi-dash | multi-colon | no-delimiter-uppercase | single-dash-uppercase | multi-dash-uppercase | multi-colon-uppercase]`

   Sets the MAC address format to use. The same format is used for all ports in the system.

3. `(config)# no aaa port-access mac-based <PORT-LIST> [addr-limit <Limit> | addr-moves | quiet-period <1-65535> | retain-unauth-clients | server-timeout <1-300> | mac-pin | max-requests <1-10> | logoff-period <1-9999999> | reauth-period <0-999999999> | unauth-period <0-255> | auth-vid <VLAN-ID> | unauth-vid <VLAN-ID> | reauthenticate]`

   Specifies parameters and limits on the configured client authentication.

4. `no aaa port-access mac-based password <PASSWORD>`

   The *password'* form of the command sets the global password for all MAC authentication clients. This password is used instead of the client's MAC address in the RADIUS request.

5. `aaa port-access mac-based <port-list> retain-unauth-clients`

   Retain `unanth-vid` is not enabled .

Resilient 802.1x cached-reauth

802.1x authenticated clients are placed in cached-reauthentication phase when a RADIUS server is not reachable. The switch sends an EAPOL (Extensible Authentication Protocol Over LAN) start message to reauthenticate the client before RADIUS connection timeout occurs or the server-times out. When configured, the client may be authorized to use a cached reauthentication as a backup method for access to the RADIUS server. Currently Aruba switches support primary and fallback authentication for both MAC authentication & DOT1x authentication. If the RADIUS server is down or unreachable, the fallback method is applied using one of the three methods available:

Authorized

When configured, authenticated clients are authorized.

Cached reauthentication

When configured, the client is authorized for the configured cached reauthentication period or RADIUS server reachability.

None

If none is configured, for the client, the client will be de-authenticated. None is the default.

Configuring a client for `retain-unauth-clients`

A series of steps must be undertaken to configure a client for enforce-cache reauthentication.

Procedure

1. `(config)# aaa authentication port-access eap-radius cached-reauth`

   Enable cache-reauth as secondary authentication method

2. `(config)# aaa port-access authenticator <PORT-LIST>`

Associate the specific port with port-access authenticator for 802.1x authentication

3. Configure server timeout < (no. of retransmit+1)*timeout default is [(3+1)*5] 20sec

```
(config)# show radius

 Dead RADIUS server are preceded by *

  Deadtime (minutes)            : 0
  Timeout (seconds)             : 5
  Retransmit Attempts           : 3
  Global Encryption Key         :
  Dynamic Authorization UDP Port : 3799
  Source IP Selection           : Outgoing Interface
  Tracking                      : Disabled
  Tracking Period (seconds)     : 300
  CPPM Identity                 :


                Auth  Acct  DM/ Time   |
  Server IP Addr  Port  Port  CoA Window | Encryption Key  OOBM
  --------------- ----- ----- --- ------ + -------------- ----
  <Server IP>      1812  1813  No   300  | <encryption-key>   No


(config)# aaa port-access authenticator <PORT-LIST> server-timeout
```

4. `(config)# aaa port-access authenticator *<PORT-LIST>* enforce-cache-reauth`

    Enable `enforce-cache-reauth` on the 802.1x authentication associated port.

5. `(config)# aaa port-access authenticator *<PORT-LIST>* cached-reauth-period`

    Set the `cache-reauth-period` for 802.1x associated port.

    a. Time in seconds, *<1-2147483647>* , during which cached reauthentication is allowed on the port. The minimum reauthentication period should be greater than 30 seconds.

6. `(config)# aaa port-access authenticator *<PORT-LIST>* reauth-period`

    Set the reauth-period for the 802.1x associated port.

    a. Enter a number, *<0-999999999>* .

7. `(config)# aaa port-access authenticator *<PORT-LIST>* [auth-vid *<VLAN-ID>* | cached-reauth-period | clear-statistics | client-limit *<1-32>* | control | enforce-cache-reauth | initialize | logoff-period | max-requests *<1-10>* | quiet-period *<1-65535>* | reauth-period *<0-999999999>*| reauthenticate | server- timeout *<1-300>* | supplicant-timeout | tx-period | unauth-period *<0-255>* | unauth-vid *<VLAN-ID>*]`

    Specifies parameters and limits on the configured client authentication.

8. `(config)# aaa port-access authenticator active`

    Initializes the authenticator.

RBAC Overview

The Role Based Access Control (RBAC) is a runtime database that consists of roles and rules that are mapped to users. RBAC lets you secure the management of your network infrastructure by defining the roles for each network administrator for their specific function. The resource access permissions ensure that the network administrator of one department cannot modify the configuration of another department. The feature access permission allows you to create roles based on the function of the user.

Every user is mapped to a role in the RBAC database and every role has one or more rules. RBAC supports 64 roles and you can configure a maximum of 1000 rules per role.

This figure shows the mapping of users, roles, and rules. In this example, `User 3` and `User 8` share the same role, `Role 3`. In turn, `Role 3` points to the various rules it was configured to support.

Figure 61: *RBAC role and rule mapping*
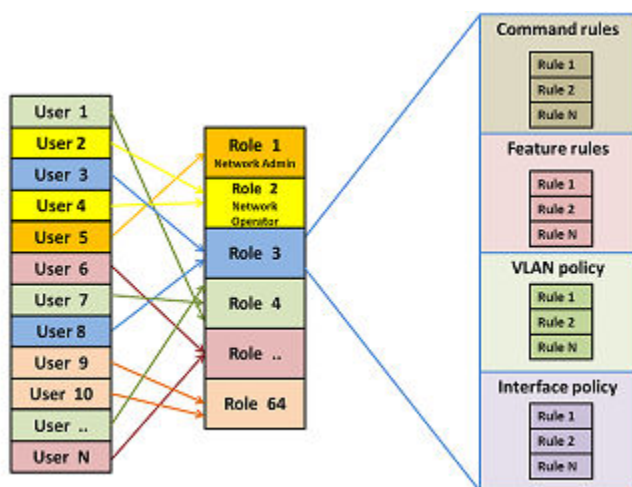
Limitations

- A user can only be configured to one role.

- You can give access to the `"command:write memory" deny` rule by saving your changes when logging out of your session.

- You cannot add the `default-security-group` rules to any other group.

- The command strings are not validated. You must provide a valid command string.

- If you configure multiple interface policy rules, only the last entry is taken into effect. All other interface policy rules are ignored.

- If you configure multiple VLAN policy rules, only the last entry is taken into effect. All other VLAN policy rules are ignored.

- RBAC supports a maximum of 1000 rules per role, which equals to 64000 rules per system (1000 rules x 64 roles).

Roles

You can configure a maximum of 64 roles in a system and for each role, you can assign one or more rules. Roles are categorized as follows:

- 3 default roles: `operator`, `manager`, and `default-security-group`

- 16 predefined roles: Level-0 to Level-15

- 45 user roles

> **NOTE**
>
> When a user is not mapped to any role, the user gets mapped to the predefined `Network-Operator` role (Level-1).

Predefined roles

RBAC offers 16 predefined roles in the system (Level-0 to Level-15) as follows:

- The `Network-Diagnostic` role (Level-0)This role can perform the following commands:

  - `ping`

  - `tracert`

  - `ssh`

  - `telnet`

  The superuser can configure the access rights for this role.

- The `Network-Operator` role (Level-1)This role has the same access rights as the `Operator` role and can perform the following commands:

  - `ping`

  - `traceroute`

  - `traceroute6`

  - `ssh`

  - `telnet`

  - All `show` commands, except for `show history`

  - All `display` commands, except for `display history`

  The superuser can configure the access rights for this role.

- User modifiable roles (Level-2 to Level-8 and Level-10 to Level-14)By default, these roles have no access to any commands. The superuser can configure the access rights for these roles.

- The `Designated-Administrator` role (Level-9)This role can perform all commands except for user management commands (such as : `deny rwx aaa`, `deny rwx tacas`, `deny radius`, `deny configure password`, `deny configure authentication`, `deny show authorization`). You cannot configure the access rights for this role.

- The `Administrator` role (Level-15)This role has the same access rights as the `Manager` role and it can perform all commands, features, and policies in the system. You cannot configure the access rights for this role.

Rules

RBAC supports a maximum of 1000 rules per role. With RBAC, you can configure the access of a user to a limited set of VLAN, interfaces, features, and commands rules.

When a user logs into the system, the role and rules are mapped to their session data structure.

Figure 62: *RBAC rule mapping based on role per session*
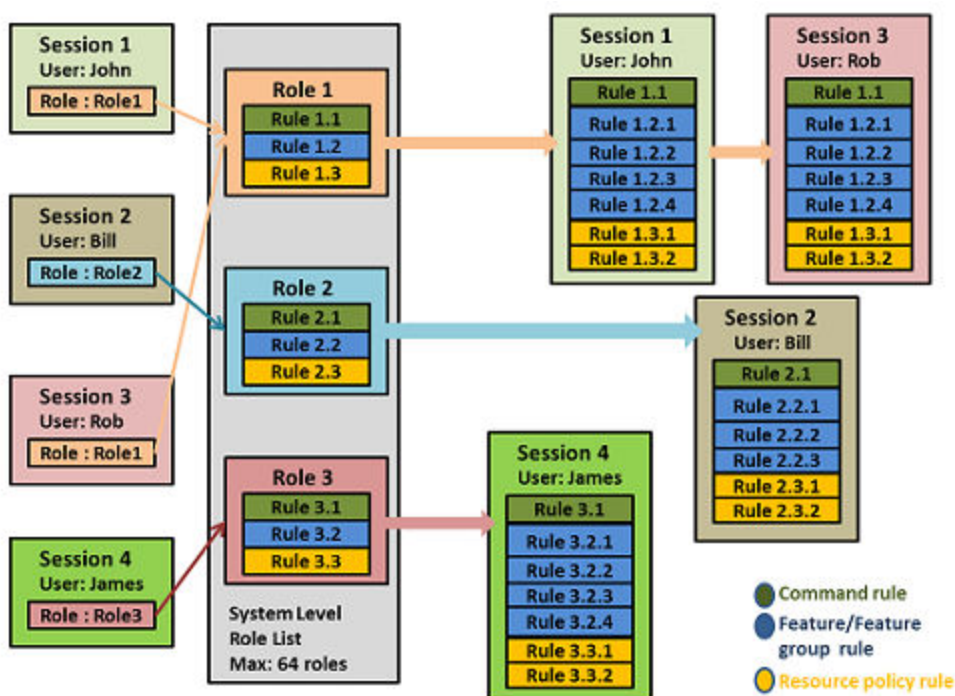
There are four types of rules:

- **Command rules**

- **Feature rules**

- **VLAN policy rules**

- **Interface policy rules**

Command rules

The command rule indicates the absolute command path, including the command context that is taken into consideration while validating the commands. The command rule is specific to each user session.

The `command` parameter must contain the command context separated with a `;` delimiter. For example, the command string that indicates the configuration of an IP address on any VLAN is as follows:

`"configure;vlan;ip address"`

> **NOTE:** The command strings are not validated. You must provide a valid command string.

Feature rules

The feature rule indicates that the feature is related to a command set. There are 40 predefined features. Each feature can have read, write, and execute privileges. You can configure multiple features for a single role. When you add a feature to a role, the command rule entries are included automatically for all the commands associated with that feature.

A feature can have the following permissions:

- `r`: The read permission displays the configuration and maintenance information. For example, the `display` and `show` commands.

- `w`: The write permission configures the feature in the system. For example, the ACL and the OSPF configuration commands.

- `x`: The execute permission executes specific functions. For example, the `ping` and the `copy` commands.

VLAN policy rules

To configure a VLAN policy rule, set the `policy` parameter to `vlan`. Only one VLAN policy rule is allowed per role. The opposite VLAN rule is applied to the rest of the VLAN IDs. For example, a policy rule "`policy:vlan:2-4`" `permit` gives access permission to user for VLANs 2 to 4 only and denies access to rest of the VLANs available in the system.

If you configure multiple VLAN policy rules, only the last entry is taken into effect. All other VLAN policy rules are ignored.

> **NOTE:** By default, VLAN policy rules allow all commands.

Interface policy rules

To configure an interface policy rule, set the `policy` parameter to the `interface` value. Only one interface policy rule is allowed per role. The opposite interface rule is applied to the rest of the interface IDs. For example, a policy rule "`policy:interface:A2-A4`" `deny` denies access permission to user for interfaces A2 to A4 only and permits access to rest of the interfaces available in the system.

If you configure multiple interface policy rules, only the last entry is taken into effect. All other interface policy rules are ignored.

> **NOTE:** By default, interface policy rules allow all commands.

Creating roles and assigning rules

To create roles and assign rules to the configured roles, follow these steps:

Procedure

1. **Enabling authorization**

2. **Creating a role**

3. Assign one or more rules:

   a. **Configuring command rules**

   b. **Configuring VLAN policy**

   c. **Configuring interface policy**

   d. **Configuring feature policy**

Enabling authorization

Before you can create a role, you must enable the authorization commands for local users as follows:

1. Run the `aaa authorization commands` command.

2. Specify the `local` parameter.

Enabling authorization commands

```
# aaa authorization commands local
```

Creating a role

Create a role as follows:

1. Run the `aaa authorization local-user` command.

2. Specify the `local-user` parameter.

3. Specify the `group` parameter.

4. Press `enter`.

5. Enter a password.

In this example, a local-user `user1` is assigned to the `network-admin` role.

Assigning a local-user

```
# aaa authentication local-user "user1" group "network-admin" password plaintext
```

Configuring command rules

Assign one or more command rules to a user as follows:

1. Run the `aaa authorization group` command.

2. Specify the `group` parameter.

3. Specify the `match-command` parameter. You can specify one or more rules.

4. Specify the access: `permit` or `deny`.

In the following example, the `network-admin` role is given access to the `router ospf` and `ip address` commands.

Permit rule

`# aaa authorization group "network-admin" 1 match-command "command:router ospf;ip address" permit log`

In the following example, the `network-admin` role is denied access to the `configure router ospf enable` command.

Deny rule

`# aaa authorization group "network-admin" 1 match-command "command:configure router ospf enable" deny log`

Configuring VLAN policy

Procedure

1. Run the `aaa authorization group` command.

2. Specify the `group` parameter.

3. Specify the `match-command` parameter for the desired VLAN policy.

4. Specify the access: `permit` or `deny`.

> **NOTE**
>
> If a command must be preceded by the execution of another command, you must first permit both commands for the command authorization group. You can then configure the rule.

In this example, the `network-admin` role is denied access to the `"policy:vlan:10-12,20,30-40"` VLAN policy. The `sequence` parameter is used to give order to the sequence of commands to be executed. **See: example**

Since only one VLAN policy rule can be assigned per role, if access is permitted for VLAN IDs 10 to 12, access to the rest of the VLAN IDs is denied for the same role. Similarly, if access is denied for VLAN IDs 10 to 12, then access to the rest of the VLAN IDs is permitted for the same role.

Configuring VLAN policy rules

`# aaa authorization group "network-admin" 1 match-command "command:^configure$" permit`

`# aaa authorization group "network-admin" 2 match-command "command:configure vlan" permit log`

`# aaa authorization group "network-admin" 3 match-command "policy:vlan: 10-12,20,30-40" deny log`

Configuring interface policy

1. Run the `aaa authorization group` command.

2. Specify the `group` parameter.

3. Specify the `match-command` parameter for the desired interface policy.

4. Specify the access: `permit` or `deny`.

NOTE If a command must be preceded by the execution of another command, you must first permit both commands for the command authorization group. You can then configure the rule.

In this example, the `network-admin` role is denied access to the "`policy:interface:A10-A12,A20,L20-L24`" interface policy. The `sequence` parameter is used to give order to the sequence of commands to be executed.

Configuring interface policy rules

```
# aaa authorization group "network-admin" 1 match-command "command:^configure$" permit

# aaa authorization group "network-admin" 2 match-command "command:configure interface" permit log

# aaa authorization group "network-admin" 3 match-command "policy:interface:A10-A12,A20,L20-L24" deny log
```

Since only one interface policy rule can be assigned per role, if access is permitted for A10 to A12, access to the rest of the interfaces is denied for the same role. Similarly if access is denied for A10 to A12, then access to rest of the interfaces is permitted for the same role.

Configuring feature policy

Procedure

1. Run the `aaa authorization group` command.

2. Specify the `group` parameter.

3. Specify the `match-command` parameter. You can specify one or more features.

4. Specify the access: `permit` or `deny`.

NOTE If a command must be preceded by the execution of another command, you must first permit both commands for the command authorization group. You can then configure the rule.

In this example, the `network-admin` role is granted access to the "`feature:rwx:ospf`" feature policy. The `sequence` parameter is used to give order to the sequence of commands to be executed.**See: example**

Configuring feature rules

```
# aaa authorization group "network-admin" 1 match-command "command:^configure$" permit
```

```
# aaa authorization group "network-admin" 2 match-command "command:configure
feature" permit log
```

```
# aaa authorization group "network-admin" 1 match-command "feature:rwx:ospf" permit
log
```

Displaying rules for predefined roles

1. Run the `show authorization` command.

2. Specify the `group` parameter.

In this example, the authorization is displayed for the predefined roles.

Displaying rules

 # show authorization group

```
Local Management Groups - Authorization Information

  Group Name: Level-0
  Group Privilege Level: 18

  Users: Tom, Bill, Will

  Seq. Num.  | Permission Rule Expression                                Log
  ---------- + ---------- ------------------------------------------- -------
  999        | Permit     ping *                                       Disable
  1000       | Permit     ping6 *                                      Disable
  1001       | Permit     traceroute *                                 Disable
  1002       | Permit     traceroute6 *                                Disable
  1003       | Permit     ssh *                                        Disable
  1004       | Permit     telnet *                                     Disable
  1005       | Permit     telnet-server *                              Disable
  1006       | Deny       .*                                           Disable

...

  Group Name: Level-15
  Group Privilege Level: 33

  Users
  ----------------

  Seq. Num.  | Permission Rule Expression                                Log
  ---------- + ---------- ------------------------------------------- -------
  999        | Permit     configure .*                                 Disable
  1000       | Permit     .*                                           Disable
```

Displaying predefined features

Procedure

1. Run the `show authorization` command.

2. Specify the `feature` parameter.

3. Specify the `detailed` option.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

In the following example, the details of the `access-list` feature are displayed.

Figure 63: *Displaying predefined features*

```
# show authorization feature access-list detailed
 access-list            IP access list related commands
    (W) command:debug acl
    (W) command:configure access-list
    (W) command:configure ip access-list
    (W) command:configure ipv6 access-list
    (W) command:configure s?vlan .* ip access-group
    (W) command:configure s?vlan .* ipv6 access-group
    (W) command:configure interface .* ip access-group
    (W) command:configure interface .* ipv6 access-group
    (W) command:configure interface tunnel .* ipv6 access-group
    (R) command:show access-list
    (R) command:show statistics aclv4
    (R) command:show statistics aclv6
    (R) command:display acl
```

To view all the predefined features in your system, enter:

```
# show authorization feature all detailed
```

Troubleshooting

Cannot modify group name

Symptom

The default group '<group-name>' cannot be modified.

Cause

User tries to modify a predefined group name.

Action

Do not attempt to change the name of a predefined group.

Cannot delete a group

Symptom

The default group '<group-name>' cannot be deleted.

Cause

User tries to delete a predefined group.

Action

Do not attempt to delete a predefined group.

Unable to run a command

Symptom

User is not authorized to execute this command.

Cause

The user is not getting access to the command.

Action

Superuser must execute the command `show logging -r` to check the sequencing of rules and arrange the rules in the proper sequence.

Unable to add a rule

Symptom

User is unable to add a rule.

Cause

Adding a new rule fails if the existing rules exceed the limit (1000).

Action

If you have exceeded the limit, you can only add a new rule if you remove an existing rule.

aaa authorization group

Syntax

```
[no] aaa authorization group <GROUPNAME> <SEQ-NUM>
    match-command {command | feature | policy} {deny | permit} [log]
```

Description

Assigns rules to existing roles. Rules can be permitted or denied for a specified user.

Parameters

GROUPNAME

The name of the role.

SEQ-NUM

When more than one rule matches the command entered, the rule with the lowest sequence number gets precedence over the other rules.

command

Indicates that the rule requires context level information to validate the command string following this parameter.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

```
feature
```

Indicates that it is a feature related to a command set. A feature can have the following permissions:

- `r`: The read feature displays the configuration and maintenance information. For example, the `display` and `show` commands.

- `w`: The write feature configures the feature in the system. For example, the ACL and the OSPF configuration commands.

- `x`: The execute feature executes specific functions. For example, the `ping` and the `copy` commands.

There are 40 predefined features. Multiple features can be configured for a single role. When a feature is added to a role, the command rule entries are included automatically for all the commands for that feature.

```
policy
```

Indicates that it is a resource policy rule. There are two resource policies: VLAN and interface.

```
deny
```

The specified match-command is denied for the specified group.

```
permit
```

The specified match-command is permitted for the specified group.

```
log
```

Generates a log message in the show logging output for the rule that is permitted or denied.

Predefined features

Feature	Description
aaa	AAA service-related commands.
arp	ARP protocol-related commands.
cdp	Cisco Discovery Protocol-related commands.
ping	Network reachability test commands.
snmp	SNMP related commands.
radius	Radius configuration and show commands.
syslog	Syslog related commands.
tacacs	TACACS configuration and show commands.
access-list	IP access list related commands.
vlan	Virtual LAN related commands.
spanning-tree	Spanning Tree protocol-related commands.

Table Continued

Feature	Description
dhcp	DHCP related commands.
gvrp	GVRP related commands.
igmp	IGMP related commands.
router	Routing related Commands.
port-security	Port security related commands.
dldp	DLDP related commands.
lldp	LLDP related commands.
crypto	Crypto related commands.
mac-access-list	MAC related commands.
telnet	Telnet related commands.
smart-link group	smart-link group related commands.
sntp	SNTP related commands.
mirror	Mirror diagnostic related commands.
rmon	RMON feature related commands.
interface	Interface related commands.
ip	IP related commands.
ipv6	IPv6 related commands.
qos	QoS related commands.
mesh	mesh related commands.
policy	classifier policy commands.
redundancy	Redundancy management related commands.
sflow	sFlow related commands.
rate-limit	Rate limit related commands.
trunk	Trunk related commands.
terminal	Terminal related commands.

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Feature	Description
tftp	TFTP related commands.
ssh	SSH related commands.
copy	copy related commands.
macsec	MAC security-related commands.

Password complexity overview

Password Complexity enforces the use and configuration of a complex password, and offers more stringent password policies. This feature complies with the UCR-2008 standard for system passwords. Password Complexity performs checks while configuring the password and provides user alerts based on the configuration of the password expiration. By default, Password Complexity is disabled.

The Password Complexity feature offers the following:

* Enable or disable password configuration and complexity features.

* Configure minimum password length.

* Configure password history specifications. Password modification requires re-authentication of user identity where the old password is required to change the password.

* Configure global as well as per user specific password aging interval.

* Notification for password expiration (alert before expiry, at expiry, and grace period).

* Configure additional number of subsequent logon attempts after password expiry. By default, three attempts within a configurable grace period (default 30 days).

* Minimum wait period before password change (default 24 hours).

* When the user establishes a session for the first time, they are prompted to change the password and the session is denied if the user does not comply.

* Enabling or disabling the display of the last successful or unsuccessful log-on information

Password expiration periods

The Password Complexity feature includes the following expiration periods:

* Aging Period: The aging period is the password expire period. This is the validity period of the password.

* Grace periods:

  ◦ The grace period before expiry: During this period, the user is informed of how many days are left for the password to expire.

  ◦ The grace period after expiry: During this period, the user is informed that the password has expired and how much more time is left after which the user will not be allowed to login if the password is not re-configured.

Requirements

The requirements to enable the Password Complexity feature are as follows:

- The manager's password must be configured. It can be done using the `password manager` command.

- The minimum length of the password must be set to a value greater than or equal to the `sum-of-compositions` value. Since the sum of the default value of compositions is 8, the minimum length of the password must be at least 8. This can be set using the `password minimum-length` command.

- The WebUI and REST interface must be disabled. Only the CLI can be used to configure the password.

The requirements to configure the password are as follows:

- Password consists of a minimum of eight characters using at least two characters from each of the four character sets: uppercase letters, lowercase letters, numbers, and special characters.

- Password cannot be the same value or the reverse form of the associated user ID.

- Password cannot have three consecutive identical characters.

- Password cannot be empty.

- The new password must differ from the old password by at least four characters.

Limitations

- This feature is supported only for local authentication.

- The password configuration is not applicable for other clients such as 802.1X clients.

- This feature is not supported for the following interfaces:

  ◦ WebUI

  ◦ Menu

  ◦ Cloud

  ◦ REST

- The `setup mgmt-interfaces` command is mutually exclusive when the Password Complexity feature is enabled.

- The `display` command does not support this feature.

- Password complexity checks are not performed when the password is in the SHA format. But the password expiry feature will still be applied.

Configuring Password Complexity

Enable the Password Complexity feature

Procedure

1. Set the minimum password length to at least 8.

2. The `manager` credentials must be configured. If they are not, use the `password manager` command to create the password.

3. Enable the Password Complexity feature (**Enable Password Complexity**)

Configure the Password Complexity parameters

• Configure the password aging, history, and log-on details. (**Configure the Password Complexity parameters**)

• Configure the password minimum length (**Configure password minimum length**)

• Configure the password composition (**Configure password composition**)

• Configure complexity checks on a new user password (**Configure password complexity checks**)

Viewing the password configuration

Procedure

To view the password configuration, enter the `show password-configuration` command.

The `Password Control` parameter displays the status and parameters of the feature, either enabled or disabled.

---

show password-configuration

In this example, the password configuration is disabled and the minimum password length is 8.

```
switch# show password-configuration
 Global password control configuration

  Password control                    : Disabled
  Password history                     : Disabled
  Number of history records            : 8
  Password aging                       : Disabled
  Aging time                           : 90 days
  Early notice on password expiration  : 7 days
  Minimum password update interval     : 24 hours
  Expired user login                   : 3 login attempts in 30 days
  Password minimum length              : 8
  User login details checking          : Enabled
  Password composition
          Lower case                   : 2 characters
          Upper case                   : 2 characters
          Special character            : 2 characters
          Number                       : 2 characters
  Repeat password checking             : Disabled
  Username checking                    : Disabled
  Repeat characters checking           : Disabled
```

Enable Password Complexity

To enable the Password Complexity feature, enter the `password configuration-control` command.

Example

In this example, the WebUI is enabled and the user enters `y` to disable the WebUI and enable the Password Complexity feature.

```
switch# password configuration-control
The password configuration feature cannot be enabled when the WebUI is enabled.
Would you like to disable WebUI and REST protocol? [y/n]: y
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Configure the Password Complexity parameters

You can configure the Password Complexity parameters at any time but they will only take effect if the Password Complexity feature is enabled.

Example

In this example, we enable the `aging` and `history` parameters.

```
switch# password configuration aging
switch# password configuration history
```

The `Password history` and `Password aging` are now enabled.

We then set the aging parameter, `aging-period`, to 60 days:

```
switch# password configuration aging-period 60
```

The `show password-configuration` displays the configuration changes:

```
switch# show password-configuration
 Global password control configuration

  Password control                       : Enabled
  Password history                        : Enabled
  Number of history records               : 8
  Password aging                          : Enabled
  Aging time                              : 60 days
  Early notice on password expiration  : 7 days
  Minimum password update interval     : 24 hours
  Expired user login                      : 3 login attempts in 30 days
  Password minimum length                 : 8
  User login details checking          : Enabled
  Password composition
          Lower case                      : 2 characters
          Upper case                      : 2 characters
          Special character               : 2 characters
          Number                          : 2 characters
  Repeat password checking                : Disabled
  Username checking                       : Disabled
  Repeat characters checking              : Disabled
```

Configure password minimum length

To configure the password minimum length, enter the `password minimum-length` command. The minimum password length must be equal or greater than the sum of the password composition. You can set the password minimum length for all users or per user.

Example

To configure the password minimum length to 10 for all users, enter:

```
switch# password minimum-length 10
```

To configure the password minimum length to 10 for the operator user `operatorABC` only, enter:

```
switch# password operator user-name operatorABC minimum-length 10
```

To configure the password minimum length to 10 for the local user `localuserXYZ` only, enter:

```
switch# aaa authentication local-user localuserXYZ min-pwd-length 10
```

Configure password composition

To configure the password composition policy for all users, which includes the minimum number of characters from the set of lowercase letters, uppercase letters, special characters, and numbers, enter the `password composition` command. The minimum password length must be equal or greater than the sum of the password composition.

Example

In this example, we set the password special characters parameter to 3.

```
switch# password composition specialcharacter 3
```

The `show password-configuration` command displays the configuration changes:

```
switch# show password-configuration
 Global password control configuration

  Password control                   : Enabled
  Password history                   : Enabled
  Number of history records          : 8
  Password aging                     : Enabled
  Aging time                         : 60 days
  Early notice on password expiration : 7 days
  Minimum password update interval   : 24 hours
  Expired user login                 : 3 login attempts in 30 days
  Password minimum length            : 10
  User login details checking        : Enabled
  Password composition
          Lower case                 : 2 characters
          Upper case                 : 2 characters
          Special character          : 3 characters
          Number                     : 2 characters
  Repeat password checking           : Disabled
  Username checking                  : Disabled
  Repeat characters checking         : Disabled
```

Configure password complexity checks

To configure password complexity checks on a new user password, enter the `password complexity` command.

Example

In this example, we set the password complexity check to all, which includes `repeat-password-check`, `repeat-char-check`, and `user-name-check`.

```
switch# password complexity all
```

password configuration commands

Use the following password configuration commands to configure the Password Complexity feature:

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Command name	Description
`password configuration`	Enables the aging, logon, and history checks. Configures aging and history parameters.You can configure the password parameters even if the password configuration feature is not enabled. But they will only take effect once the password configuration feature is enabled.The `no` option disables the Password Complexity configuration.
`password configuration aging`	Enables the password configuration aging check.The `no` option disables aging.
`password configuration history`	Enables the password history check.After authentication, the history is updated. The history is also updated whenever there is reconfiguration of the password. Once the maximum number of configured entries are reached, the oldest entry is overwritten.The `no` option disables the history check.
`password configuration log-on-details`	Disables the display of user login details. The `no` option enables the same.
`password configuration aging-period`	Configures the global password aging time for a system. The `no` option sets the global aging time to the default value of 90 days.
`password configuration alert-before-expiry`	Specifies the number of days for which the user is warned of the pending password expiration. The default value is 7 days.
`password configuration expired-user-login`	Configures additional login attempts allowed or a delay period during which the user is allowed to login after the password expiry. The default value is 30 days. The maximum number of login attempts is 10, the default is 3.
`password configuration update-interval-time`	Configures the minimum period of waiting, in hours, before an existing password can be changed.

Table Continued

Command name	Description
`password configuration history-record`	Configures the maximum number of history password records for each user.
`password minimum-length`	Configures the minimum password length and completes the `password configuration` command. When changing the password for the manager, operator, and local management users, the new password must be at least the length of this parameter. The possible values are 0 to 64.The `no` option sets the minimum password length to the default value of 0. When the Password Complexity feature is enabled, the minimum password length is 15 for the `manager` user and 8 for all other users, including the `operator`. The range is 15 to 64 for the `manager` and 8 to 64 for all other users.

password configuration-control

The `password configuration-control` command globally enables or disables the Password Complexity feature.

Syntax

`[no] password configuration-control`

Description

Enables the Password Complexity feature to follow the UCR-2008 standard. The `no` option disables the password complexity feature.

Options

`no`

Disables the Password Complexity feature.

Limitations

This command has the following prerequisites:

- The manager's password must be configured. You can configure the manager's password using the `password manager` command.

- You must set the minimum length of the password to a value greater than or equal to the sum-of-compositions. Since the sum of the default value of compositions is 8, the minimum length of the password must be at least 8. You can set this value with the `password minimum-length` command.

- The WebUI and REST interface must be disabled.

password configuration

The `password configuration` command enables the aging, logon, and history checks and configures the aging and history parameters.

The `no` option disables the password configuration feature.

Syntax

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

```
[no] password configuration [aging |
          history |
          log-on-details |
          aging-period <aging-time> |
          alert-before-expiry <alert-time> |
          expired-user-login [days <delay> ] [attempts <time> ] |
          update-interval-time <time> |
          history-record <max-record-num>]
```

Parameters

`aging`

Enables the password configuration aging check.

`history`

Enables the password history check.

`log-on-details`

Disables execution of the `show authentication last-login` command to display the logon details.

`aging-period`

Configures the password aging time for a system.

`alert-before-expiry`

Sets the number of days before password aging during which the user is warned of the pending password expiration.

`expired-user-login`

Configures additional login attempts within a specified period during which a user is allowed to access the switch without changing an expired password.

`update-interval-time`

The period of waiting, in hours, before an existing password can be changed.

`history-record`

Configures the maximum number of history password records for each user.

Options

`no`

Disables the password configuration feature.

`attempts`

The number of subsequent login attempts allowed after the password expiry. Possible values are 0 to 10, the default value is 3.

`days`

The period during which subsequent login attempts are allowed after the password expiry. Possible values are 1 to 90, the default value is 30 days.

`update-interval-time <time>`

The minimum period of waiting, in hours, before an existing password can be updated. Possible values are 0 to 168, the default value is 24 hours.

`max-record-num`

Maximum number of history password records. Possible values are 2 to 15 for each user, the default value is 8.

`aging-time`

Password aging time, in days. Possible values are 1 to 365, the default value is 90 days.

```
alert-time
```

Sets the number of days before password aging during which the user is warned of the pending password expiration. The `no` option sets the alert time to the default value of 7.

password minimum-length

Syntax

```
[no] password minimum-length <length>
```

Description

Configures the minimum password length. When changing the password for the manager, operator, or local management users, the new password must be at least the length of this parameter.

Parameters

```
length
```

When the Password Complexity feature is enabled, the minimum password length is 15 for the `manager` user and 8 for all other users, including the `operator`. The range is 15 to 64 for the `manager` and 8 to 64 for all other users.

Options

```
no
```

Sets the minimum password length to the default value of 0. If the Password Complexity feature is disabled, the `length` parameter is set to 0. The manager minimum password length is 15.

password

Syntax

```
[no] password [manager|operator]
                    [user-name ASCII-STR] [{{plaintext|sha1} ASCII-STR} |
                    {min-pwd-length{length}} |
                    {aging-period {value}} |
                    {clear-history-record}}]
```

Description

Configures the local password and username for an access level. If no password is specified on the command line, the user will be prompted for the new password and for confirmation. The `port-access` password is only configurable when `include-credentials` is enabled.

Parameters

```
aging-period
```

Configures the password aging time for a user. This will override the global set value. The `no` option applies the global aging time to the user password expiry.

Password aging time, in days. Possible values are 1 to 365, the default value is 90 days.

```
clear-history-record
```

Clears history records of passwords for a user. The `no` option results in no change.

```
min-pwd-length
```

Configures the minimum password length for a user. The `no` option applies the default minimum length to the user. If the Password Complexity feature is enabled, the default minimum password length is 15 for a manager user and 8 for all other users.

`plaintext`

Prompts for a plaintext password. The password can have a maximum of 64 characters. It must not contain spaces and is case-sensitive. Plaintext is the default type.

Options

`no`

Removes the specified password.

aaa authentication local-user

Syntax

```
[no] aaa authentication local-user <USERNAME>
                {{group <GROUPNAME> password {plaintext|sha1 <PASSWORD> }}|
                {aging-period <aging-time> }
                | {min-pwd-length <length> }
                | {clear-password-history}}
```

Description

Configures the aging period, minimum password length, and clear password history for a local user.

Parameters

`local-user`

The local user being added to the authorization group. The username can be up to 16 characters. The username must not contain spaces and should be case-sensitive.

`group`

Name of the authorization group to which the local user belongs. This must be a pre-existing group.

`aging-period`

Configures the password aging time in days for a user. Possible values are 1 to 365. The default value is 90 days. The `no` option applies the global aging time for the user.

`min-pwd-length`

Configures the password minimum length for a user. Possible values are 1 to 64. The default value is 8 characters. The `no` option applies the default minimum length for the user.

`clear-password-history`

Clears the password history for a user. The `no` option does not have any affect.

password complexity

Syntax

```
[no] password complexity [repeat-password-check |
        repeat-char-check |
        user-name-check|
        all]
```

Description

The `password complexity` command configures complexity checks on a new user password. You can enable the Password Complexity feature only if the password control is enabled. Use the `password configuration-control` command to enable it.

Parameters

`repeat-password-check`

Configures the repeat password character check.

`repeat-char-check`

Ensures that password does not contain three of the same characters used consecutively.

`user-name-check`

Ensures that the password does not contain the associated username or its reverse form.

Options

`no`

Disables the configuration.

Restrictions

The password control must be enabled. Use the `password configuration-control` command to enable it.

password composition

Syntax

```
password composition [lowercase | uppercase | specialCharacter | number] <value>
```

Description

Configures the password composition policy for all users. This includes the minimum number of characters from the set of lowercase letters, uppercase letters, special characters, and numbers.

Parameters

`lowercase`

Minimum number of lowercase characters. The default value is 2. Possible values are 2 to 15.

`uppercase`

Minimum number of uppercase characters. The default value is 2. Possible values are 2 to 15.

`specialCharacter`

Minimum number of special characters. The default value is 2. Possible values are 2 to 15.

`number`

Minimum number of number character type. The default value is 2. Possible values are 2 to 15.

show password-configuration

Syntax

```
show password-configuration [manager | operator | [group <group_name>]] <user_name>
```

Description

The `show password-configuration` command displays the global password information for all users as well as for a particular user.

Parameters

`manager`

Displays the password configuration for the `manager` user.

`operator`

Displays the password configuration for the `operator` user.

`group`

Displays the password configuration for the `group` user.

Example input

`switch# show password-configuration manager ABCD`

Example output

```
password settings for the user:
 Aging time:                        10 days
 Minimum password length:           15

Global password control configurations:

 Password control:                  Enabled

 Password history:                  Enabled
 Number of history records:         8

 Password aging:                    Enabled
 Aging time:                        80 days
 Early notice on password expiration: 7 days
 Minimum password update interval:   24 hours
 Expired user login:                3 login attempts in 30 days
 Password composition:
         Lower case:                2 characters
         Upper case:                2 characters
         Special character:         2 characters
         Number :                   2 characters
 Repeat password checking:          Enabled
 Username checking:                 Enabled
 Repeat characters checking:        Enabled
```

Troubleshooting

Unable to enable Password Complexity

Symptom

Getting an error when trying to enable the Password Complexity feature.

Cause

The username must be unique on the switch when the Password Complexity feature is enabled.

Action

Select a unique username.

Unable to download the configuration file

Symptom

Getting an error message when trying to download the configuration file.

Cause

When the password complexity feature is enabled, the configuration file that you are downloading must have a unique username for each privilege.

Action

Edit the configuration file to make sure that the usernames are unique for each privilege.

Validation rules

Validation	Error/Warning/Prompt
Fail updating the password if the old password entered is invalid.	`The old password entered is invalid. Operation aborted.`
Fail the password command when a given password string is not satisfying the password control requirement.	**During authentication**:`The password must include a minimum of two of these types: uppercase, lowercase, 0-9, and special characters.` **During password change from CLI**:`Password validation error: Password cannot be changed. It must have special characters, A-Z, a-z & 0-9.`
Password Minimum length check.	**During authentication**:`Password minimum length check failed; operation aborted.` **During password change from CLI**:`Password validation error: Password minimum length check failed; operation aborted.`
NULL password.	**During authentication**:`Blank password is not acceptable; operation aborted.` **During password change from CLI**:`Password validation error: Blank password is not acceptable; operation aborted.`

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Validation	Error/Warning/Prompt
Password contains username or the backwards form of the ID.	**During authentication**:`Password contains the username; operation aborted.` **During password change from CLI**:`Password validation error: Password contains the username.Password validation error: Password contains reverse of associated username.`
Fail the password command if given password contains three of the same characters used consecutively.	**During authentication**:`Password contains repetitive characters; operation aborted.` **During password change from CLI**:`Password validation error: Password contains repetitive characters; operation aborted.`
Fail the password command if the given password does not differ from the previous password by at least four characters.	**During authentication**:`Password cannot be changed. It must differ from the previous by four characters.` **During password change from CLI**:`Password validation error: Password cannot be changed. It must differ from the previous by four characters.`
Fail the password command if the given password is the same as that of a password configured within the password history period.	**During authentication**:`Password cannot be changed; the password entered was used previously.` **During password change from CLI**:Password validation error: `Password cannot be changed; the password entered was used previously.`
Fail the password command if it is executed before the password update interval time.	Password command will fail with an error message: `The minimum wait period for updating password is not expired. Operation aborted.`
Password history clear command: `clear password-history operator/manager/local group <name>`	If the name does not exist, the following error message is displayed:`User:user1 does not exist.`
Password minimum length should match the sum of the compositions (Lowercase + upper case + special characters + numbers).	`The minimum password length configured is %s less than the sum of password composition.Operation aborted.`

Table Continued

Validation	Error/Warning/Prompt
If user last login details display is disabled and the user executes the `sh authentication last-login` command	`The last login details cannot be displayed. Command execution is currently disabled by executing password configuration command 'password configuration log-on-details'.`
To enable the password configuration feature, the following should be configured: 1. Minimum password length should be set to 8 or greater. 2. Manager credentials should be configured. 3. Web UI should be disabled.	The precedence in the error message is as follows: 1. Password minimum length is < 8.`The minimum password length configured is 8 less than the sum of password composition. Operation aborted.` 2. Manager is not configured:`Manager credentials should be configured to enable the password configuration feature` 3. Web UI will ask for the following confirmation:`"The password configuration feature cannot be enabled when the WebUI is enabled.Would you like to disable WebUI and REST protocol? [y/n]:"`
When the Password Complexity feature is enabled and manager user is deleted from the system.	`Manager account cannot be deleted when the password configuration feature is enabled.`
WebUI and password configurations are mutually exclusive.	1. `WebUI cannot be enabled when the password configuration feature is enabled.` 2. `The password configuration feature cannot be enabled when the WebUI is enabled.`
When incorrect old password is entered during password change.	`The old password is invalid.`

While enabling the password configuration function:	`switch# password configuration-control`
	
	`The password configuration feature cannot be enabled when the WebUI is enabled.`
	
	`Would you like to disable WebUI and REST protocol? [y/n]:y`

Display messages

Validation	Error/Warning/Prompt

First time log-on, the user is prompted for entering new password along with old password.	`Please change the password to logon to the system.`
	`Old password: ********`
	`New password: ********`
	`Re-enter the new password: ********`

First time log-on, if the user fails to change the password the following error message is displayed.	`The respective error message related to password validation will be displayed and the existing session termination message is displayed.`
When the value of the alert before expiry is configured greater than the aging value.	`Alert before expiry is greater than password aging value.`
When the user is not configured on the switch.	`Password configuration feature is enabled. Configure the password for the user 'admin' to get access.`

When a password aging time is about to expire a warning message is displayed.	`Password ages out in %d day(s).`
	`Change it now [Y/Any key - No]?`
	`          If yes, the user will be prompted:`
	`          Old password: ********`
	`          New password: *************`
	`          Re-enter the new password:  *************`
	
	If the new password configuration fails, then the error message will be prompted and the user will be provided the access:
	
	`Password cannot be changed. It must have special characters, A-Z, a-z & 0-9. Your previous successful login (as manager) was on 1990-01-01 07:57:24 from the console`
	`switch#`

Table Continued

Validation	Error/Warning/Prompt
| When a password aging time expires, and the password is still not updated with the user still within the configured login attempt and period the following warning message is displayed. | ```
Password expired; %d login left in %d
day(s).
Change it now [Y/Any key - No]?
            If yes, the user will be
prompted:
            Old password: ********
            New password:
*************
            Re-enter the new
password:   *************
```<br><br>If the new password configuration fails, then the error message will be prompted and the user will be provided the access:<br><br>```
Password cannot be changed. It must
have special characters, A-Z, a-z & 0-9.
```<br><br>```
Your previous successful login (as
manager) was on 1990-01-01 07:57:24
 from the console
switch#
``` |
| When a password aging time expires and all the additional login attempts are exhausted, a warning message is displayed with a prompt for entering the new password. | ```
Password grace period is complete.
Please change password.
Old password: ********
New password: ********
Re-enter the new password: ********
```<br><br>If the above entry fails, the respective error message related to password validation will be displayed and the existing session termination message is displayed. |

Overview

The switches covered in this guide use Secure Shell version 2 (SSHv2) to provide remote access to management functions on the switches via encrypted paths between the switch and management station clients capable of SSH operation.

SSH provides Telnet-like functions but, unlike Telnet, SSH provides encrypted, authenticated transactions. The authentication types include:

- Client public-key authentication

- Switch SSH and user password authentication

NOTE

SSH in switches is based on the OpenSSH software toolkit. For more information on OpenSSH, visit **http://www.openssh.com**.

Client public-key authentication (login/operator level) with user password authentication (enable/manager level)

This option uses one or more public keys (from clients) that must be stored on the switch. Only a client with a private key that matches a stored public key can gain access to the switch, and the same private key can be stored on one or more clients.

Figure 64: *Client public-key authentication model*

Switch SSH and user password authentication

This option is a subset of full certificate authentication of the client public-key authentication shown in **Client public-key authentication (login/operator level) with user password authentication (enable/manager level)**. It occurs if the switch has SSH enabled but does not have login access (`login public-key`) configured to authenticate the client's key. As in **Client public-key authentication (login/operator level) with user password authentication (enable/manager level)**, the switch authenticates itself to SSH clients. Users on SSH browser then authenticate themselves to the switch (login and/or enable levels) by providing passwords stored locally on

the switch or on a TACACS+ or RADIUS server. However, the client does not use a certificate to authenticate itself to the switch.

Figure 65: *Switch/user authentication*

Prerequisite for using SSH

Before using the switch as an SSH server, install a publicly or commercially available SSH client application on the computers to be used for management access to the switch. For client public-key authentication the client program must have the capability to generate or import keys, see **Client public-key authentication (login/ operator level) with user password authentication (enable/manager level)** on page 353 for more details.

Public key formats

Any client application used for client public-key authentication with the switch must have the capability to export public keys. The switch can accept keys in the PEM-encoded ASCII format or in the non-encoded ASCII format.

Figure 66: *Public key in PEM-encoded ASCII format common for SSHv2 clients*

Steps for configuring and using SSH for switch and client authentication

For two-way authentication between the switch and an SSH client, you must use the login (operator) level.

Procedure

A. Client preparation

1. Install an SSH client application on a management station to be used for access to the switch. (See the documentation provided with your SSH client application.)

2. Optional–If you want the switch to authenticate a client public key on the client:

a. Either generate a public/private key pair on the client computer (if your client application allows) or import a client key pair generated using another SSH application.

b. Copy the client public key into an ASCII file on a TFTP server accessible to the switch and download the client public-key file to the switch. The client public-key file can hold up to 10 client keys. This topic is covered under **Creating a client public-key text file** on page 374.

B Switch preparation

3. Assign a login (operator) and enable (manager) password on the switch, see **Configuring the switch for SSH operation** for details.

4. Generate a public/private key pair on the switch, see **Configuring the switch for SSH operation** for details.

   You need to do this only once. The key remains in the switch even if you reset the switch to its factory-default configuration. You can remove or replace this key pair, if necessary.

5. Copy the switch public key to the SSH clients you want to access the switch, see for more details.

6. Enable SSH on the switch, see **Configuring the switch for SSH operation** for more details.

7. Configure the primary and secondary authentication methods for the switch to use. In all cases, the switch will use its host public key to authenticate itself when initiating an SSH session with a client.

   • SSH Login (operator) options:

   NOTE
   If you want the switch to perform client public-key authentication, you must configure the switch with Option B.

   ◦ Option A:

   Primary: Local, TACACS+, or RADIUS password

   Secondary: Local password or none.If the primary option is local, the secondary option must be none.

   ◦ Option B:

   Primary: Client public-key authentication (**SSH client public-key authentication notes** on page 373)

   Secondary: none

   • SSH Enable (manager) options:

   Primary: Local, TACACS+, or RADIUS

   Secondary: Local password or none. If the primary option is local, the secondary option must be none.

8. Use your SSH client to access the switch using the switch IP address or DNS name (if allowed by your SSH client application). See the documentation provided with the client application.

SSH options

Table 27: *SSH options*

Switch access level	Primary SSH authentication	Authenticate switch public key to SSH clients?	Authenticate client public key to the switch?	Primary switch password authentication	Secondary switch password authentication
operator(login)level	ssh login rsa	Yes	Yes[1]	No[1]	local or none
	ssh login Local	Yes	No	Yes	none
	ssh login TACACS	Yes	No	Yes	local or none
	ssh login RADIUS	Yes	No	Yes	local or none
manager(enable)level	ssh enable local	Yes	No	Yes	none
	ssh enable tacacs	Yes	No	Yes	local or none
	ssh enable radius	Yes	No	Yes	local or none

[1] For `ssh login public key`, the switch uses client public-key authentication instead of the switch password options for primary authentication.

General operating rules and notes

- Public keys generated on an SSH client must be exportable to the switch. The switch can store 100 client key pairs.

- The switch public/private key pair and the (optional) client public-key file are stored permanently in the switch memory and are unaffected by the `erase startup-config` command.

- Once you generate a key pair on the switch, you must avoid regenerating the key pair without a compelling reason. Otherwise, you must reintroduce the switch public key on all management stations (clients) you previously set up for SSH access to the switch. This can temporarily allow security breaches.

- With SSH running, the switch allows one console session and up to five other sessions (SSH and/or Telnet).

Configuring the switch for SSH operation

Procedure

1. Assign a local login (operator) and enable (manager) password.

   At a minimum, Hewlett Packard Enterprise recommends that you always assign at least a manager password to the switch. Otherwise, under some circumstances, anyone with Telnet, web, or serial port access could modify the switch configuration.

   You can configure both the operator and manager password with one command.

   ```
   password <manager|operator|all>
   ```

```
HP Switch(config)# password all
New password for operator: ********
Please retype new password for operator: ********
New password for manager: *******
Please retype new password for manager: *******
```

2. Generate the switch public and private key pair.

    See **Generating or erasing the switch public/private host key pair**.

3. Provide the switch public key to clients.

    See **Providing the switch public key to clients**.

4. Enable SSH on the switch and anticipate SSH client contact behavior.

    See **Enabling SSH on the switch and anticipating SSH client contact behavior**.

5. Configure the switch for SSH authentication.

    See **Configuring the switch for SSH authentication**.

6. Use an SSH client to access the switch.

    Test the SSH configuration on the switch to ensure that you have the level of SSH operation needed for the switch. If you have problems, see "RADIUS-related problems" in the management and configuration guide for your switch.

Generating or erasing the switch public/private host key pair

A public and private host key pair must be generated on the switch. The switch uses this key pair along with a dynamically generated session key pair to negotiate an encryption method and session with an SSH client trying to connect to the switch.

The host key pair is stored in the switch flash memory, and only the public key in this pair is readable. The public key should be added to a "known hosts" file (for example, `$HOME/.ssh/known_hosts` on UNIX systems) on the SSH clients which should have access to the switch. Some SSH client applications automatically add the switch public key to a "known hosts" file. Other SSH applications require you to manually create a known hosts file and place the switch public key in the file. See the documentation for your SSH client application for more details.

(The session key pair mentioned above is not visible on the switch. It is a temporary, internally generated pair used for a particular switch/client session, and then discarded.)

> **NOTE**
>
> When generating a host key pair on the switch, the switch places the key pair in flash memory and not in the running-config file. Also, the switch maintains the key pair across reboots, including power cycles. Consider this key pair to be "permanent" and avoid re-generating the key pair without a compelling reason. Otherwise, you must re-introduce the switch public key on all management stations you have set up for SSH access to the switch using the earlier pair.
>
> Removing (zeroing) the switch public/private key pair renders the switch unable to engage in SSH operation and automatically disables IP SSH on the switch. To verify whether SSH is enabled, execute `show ip ssh`. However, any active SSH sessions will continue to run, unless explicitly terminated with the CLI `kill` command.

Because the host key pair is stored in flash instead of the running-config file, it is not necessary to `use write memory` to save the key pair. Erasing the key pair automatically disables SSH.

Procedure

•   To generate the key, enter the appropriate `crypto key generate` command.

For the complete syntax, see **crypto key generate**.

- To show the key, enter the `show crypto host-public-key` command.

  For the complete syntax, see **show crypto host-public-key**.

  If you compare switch key to the key stored in your client's known-hosts file, note that the formatting and comments need not match.

- To erase the key, enter the appropriate `zeroize` command.

  For the complete syntax, see **zeroize** .

crypto key generate

Syntax

```
crypto key generate <autorun-key[rsa]|cert[rsa]
<keysize>|[ssh][dsa|rsa]bits <keysize>>
```

Description

Installs authentication files for ssh or https server, or for autorun.

Options

autorun-key

Install RSA key for autorun. See "Configuring Autorun on the Switch" in the management and configuration guide for more information.

cert

Install RSA key for https certificate.

Use your SSL enabled browser to access the switch using the switch IP address or DNS name (if allowed by your browser). See the documentation provided with the browser application for more information.

ssh [dsa|rsa]

Install host key for ssh server. Specify the key type as DSA or RSA.

bits <keysize>

Specify the key size (in bits).

The `crypto key generate ssh` command allows you to specify the type and length of the generated host key.

If `<keysize>` is not specified, the default value is used.

show crypto host-public-key

Syntax

```
show crypto host-public-key [babble|fingerprint]
```

Description

Displays switch public key. Displays the version 1 and version 2 views of the key. If no options are specified, the key is displayed in a nonencoded ASCII numeric string.

See **SSH client public-key authentication** on page 60 for information about public keys saved in a configuration file.

Options

babble

Displays hashes of the switch public using phonetic format, displayed as a short series of alphabetic character groups.

fingerprint

Displays fingerprints of the switch public key using hexadecimal format.

Example

```
Switch(config)# show crypto host-public-key babble          Phonetic "Hash" of Switch's Public Key
1024 xubep-neryf-typyz-necef-hekih-navir-bynyd-puzon-dumoc-gikyz-mexax
     host_ssh2.pub                                            Hexadecimal "Fingerprints"
                                                              of the Same Switch Public
Switch(config)# show crypto host-public-key fingerprint
1024 23:50:cb:73:f8:9d:09:bc:41:08:86:43:9c:42:e8:66 host_ssh2.pub
```

zeroize

Syntax

zeroize <ssh|cert|autorun[rsa]>

Description

Erases the switch public/private key pair and disables SSH operation.

Displaying the public key

The switch provides three options for displaying its public key. This is helpful if you need to visually verify that the public key the switch is using for authenticating itself to a client matches the copy of this key in the client's "known hosts" file:

• Non-encoded ASCII numeric string:

Requires a client ability to display the keys in the "known hosts" file in the ASCII format. This method is tedious and error-prone due to the length of the keys. See **Providing the switch public key to clients**.

The switch always uses an ASCII version of its public key, without babble or fingerprint conversion, for file storage and default display format.

• Phonetic hash:

Outputs the key as a relatively short series of alphabetic character groups. Requires a client ability to convert the key to this format.

• Hexadecimal hash:

Outputs the key as a relatively short series of hexadecimal numbers. Requires a parallel client ability.

Procedure

Enter the show crypto host-public-key command.

For example, on the switch, generate the phonetic and hexadecimal versions of the switch public key as follows:

Figure 67: *Visual phonetic and hexadecimal conversions of the switch public key*

```
Switch(config)# show crypto host-public-key babble                Phonetic "Hash" of Switch's Public Key
 1024 xubep-neryf-typyz-necef-hekih-navir-bynyd-puzon-dumoc-gikyz-mexax
     host_ssh2.pub                                                 Hexadecimal "Fingerprints"
                                                                   of the Same Switch Public
 Switch(config)# show crypto host-public-key fingerprint
 1024 23:50:cb:73:f8:9d:09:bc:41:08:86:43:9c:42:e8:66 host_ssh2.pub
```

The two commands shown in example convert the displayed format of the switch (host) public key for easier visual comparison of the switch public key to a copy of the key in a client's "known host" file. The switch has only one RSA host key.

The `babble` and `fingerprint` options produce two hashes for the key that corresponds to the challenge hash you will see if connecting with a v1 client, and the other corresponding to the hash you will see if connecting with a v2 client. These hashes do not correspond to different keys, but differ only because of the way v1 and v2 clients compute the hash of the same RSA key.

Providing the switch public key to clients

When an SSH client contacts the switch for the first time, the client will challenge the connection unless you have already copied the key into the client's "known host" file. Copying the switch key in this way reduces the chance that an unauthorized device can pose as the switch to learn your access passwords. The most secure way to acquire the switch public key for distribution to clients is to use a direct, serial connection between the switch and a management device (laptop, PC, or UNIX workstation), as described below.

The public key generated by the switch consists of three parts, separated by one blank space each:

Figure 68: *A public key generated by the switch*

```
Comment          Public Key


ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAAABAQCKzbhu89Lp5vmpuDwHmynUDnaBPPDtL9J8zcr4BLXi
k8G7VoMGghxJisLewYjFFUVVSOha8XCCGh4aVCur/HllP1enMGMGl/fu42oYDobmfgf/kJCK6U527rEU
gc+1VjxebjRGA5j0Bs8R2xo2k5ty6ZciMnDoEBeRcuoxd58dwzy7kU5hjQHSkubgwBBHYv7vmLtY0dPz
RHIRSoEezm5jFDhtF2EuubuB9adA43hggmKqciZZ8EfiOmKE9QwQCWLw2vUnXLhfFhxYwo7CoTir4yJA
lITSYFTrDd0G/pD67VfToz9DeHO163Yl9ukQQNXGES0LRK+8I1SZO33J7eb5
```

Prerequisites

•   A direct serial connection from a management station to the switch.

•   A terminal application such as HyperTerminal

Procedure

1. Use a terminal application such as HyperTerminal to display the switch public key with the `show crypto host public-key` command, see **Displaying the public key**.

2. Bring up the SSH client's "known host" file in a text editor such as Notepad as straight ASCII text, and copy the switch public key into the file.

3. Ensure that there are no changes or breaks in the text string. A public key must be an unbroken ASCII string. Line breaks are not allowed (changes in the line breaks will corrupt the Key.) For example, if you are using

Windows® Notepad, ensure that `Word Wrap` (in the `Edit` menu) is disabled, and that the key text appears on a single line.

Figure 69: *Example of a correctly formatted public key*

4. Add any data required by your SSH client application. For example, before saving the key to an SSH client's "known hosts" file you may have to insert the switch IP address:

Figure 70: *Example of a switch public key edited to include the switch's IP address*

For more on this topic, see the documentation provided with your SSH client application.

Enabling SSH on the switch and anticipating SSH client contact behavior

When configured for SSH, the switch uses its host public key to authenticate itself to SSH clients. For SSH clients to authenticate themselves to the switch, configure SSH on the switch for client public-key authentication at the login (operator) level. To enhance security also configure local, TACACS+, or RADIUS authentication at the enable (manager) level.

See **Configuring the switch for SSH authentication** and **SSH client contact behavior**.

Prerequisites

The public/private key pair for switch must have been generated. See **Generating or erasing the switch public/private host key pair**.

Procedure

Execute the `ip ssh` command.

For the complete syntax, see **ip ssh** .

Protect your private key file from access by anyone other than yourself. If someone can access your private key file, they can penetrate SSH security on the switch by appearing to be you.

SSH does not protect the switch from unauthorized access via the WebAgent, Telnet, SNMP, or the serial port. While WebAgent and Telnet access can be restricted by the use of passwords local to the switch, if you are unsure of the security this provides, you may want to disable web-based and/or Telnet access (`no web-management` and `no Telnet`). If you need to increase SNMP security, use SNMP version 3 only. To increase the security of your web interface see the section on SSL. For an additional security measure, see the authorized IP managers feature in the management and configuration guide for your switch. To protect against unauthorized access to the serial port (and the `Clear` button, which removes local password protection), keep physical access to the switch restricted to authorized personnel.

Figure 71: *Enabling IP SSH and displaying the SSH configuration*

```
HP Switch(config) ip ssh              Enable SSH

HP Switch(config)# show ip ssh

  SSH Enabled     : Yes               Secure Copy Enabled : No
  TCP Port Number : 22                Timeout (sec)       : 120
  IP Version      : IPv4orIPv6
  Host Key Type   : RSA               Host Key Size       : 1024

  Ciphers : aes128-cbc,3des-cbc,aes192-cbc,aes256-cbc,
            rijndael-cbc@lysator.liu.se,aes128-ctr,aes192-ctr,aes256-ctr
  MACs    : hmac-md5,hmac-sha1,hmac-sha1-96,hmac-md5-96

  Ses Type      | Source IP                                      Port
  --- -------- + --------------------------------------------- -----
  1   console  |
  2   telnet   |
  3   ssh      | 12.255.255.255
  4   inactive |
  5   inactive |
  6   inactive |
```

With SSH running, the switch allows one console session and up to five other sessions (SSH and/or Telnet). WebAgent sessions are also allowed, but do not appear in the **show ip ssh** listing.

ip ssh

Syntax

```
ip ssh [cipher <cipher-type>] [mac <mac-type>] [port <1-65535|default>] [timeout <5-120>]

[no] ip ssh [cipher <cipher-type>] [mac <mac-type>]
```

Description

Enables SSH on the switch.

The `no` form of the command disables SSH on the switch.

Options

```
[cipher <cipher-type>]
```

Specify a cipher type to use for connection.

Valid types are:

- aes128-cbc
- 3des-cbc

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

- aes192-cbc

- aes256-cbc

- aes128-ctr

- aes192-ctr

- aes256-ctr

Default: All cipher types are available.

Use the `no` form of the command to disable a cipher type.

`[filetransfer]`

Enable/disable secure file transfer capability.

SCP and SFTP secure file transfer will not function unless SSH is also enabled.

`[mac <mac-type>]`

Allows configuration of the set of MACs that can be selected.

Valid types are:

- hmac-md5

- hmac-sha1

- hmac-sha1-96

- hmac-md5-96

Default: All MAC types are available.

Use the `no` form of the command to disable a MAC type.

`[port <1-65535|default>]`

The TCP port number for SSH connections.

Default: 22.

`[timeout <5-120>]`

Sets the maximum length of time (in seconds) allowed for initial protocol negotiation and authentication.

Default: 120 seconds

> **NOTE**
>
> Hewlett Packard Enterprise recommends using the default TCP port number (22). However, you can use the `ip ssh port` command to specify any TCP port for SSH connections except those reserved for other purposes. Examples of reserved port numbers reserved IP ports are 23 (Telnet) and 80 (http). Some other reserved TCP ports on the switch are 49, 80, 1506, and 1513.

Disabling SSH on the switch

Procedure

1. Do one of the following:

- Execute the `no ip ssh` command (see **ip ssh**)

- Zeroize the switch existing key pair (see **Generating or erasing the switch public/private host key pair**)

Configuring the switch for SSH authentication

Either method you choose for this task results in authentication of the switch public key by an SSH client. However only Option B results in the switch also authenticating the client's public key.

See also **SSH client public-key authentication notes** on page 373.

Procedure

1. Use one of the following options to configure the switch for SSH authentication

- **Option A: Configuring SSH access for password-only SSH authentication**

  When configured with this option, the switch uses its public key to authenticate itself to a client, but uses only passwords for client authentication.

- **Option B: Configuring the switch for client Public-Key SSH authentication**

  When configured with this option, the switch uses its public key to authenticate itself to a client, but the client must also provide a client public key for the switch to authenticate. This option requires the additional step of copying a client public-key file from a TFTP or SFTP server into the switch.

  Using this option results in the switch also authenticating the client's public key.

Option A: Configuring SSH access for password-only SSH authentication

When configured with this option, the switch uses its public key to authenticate itself to a client, but uses only passwords for client authentication.

> **NOTE**
> Hewlett Packard Enterprise recommends that you always assign a manager-level (enable) password to the switch. Without this level of protection, any user with Telnet, web, or serial port access to the switch can change the switch configuration. If you configure only an operator password, entering the operator password through telnet, web, ssh or serial port access enables full manager privileges. See **1** on page 356.

Procedure

1. Configure a password method for the primary and secondary login (operator) access.

   If you do not specify an optional secondary method, it defaults to `none`. If the primary method is `local`, the secondary method must be `none`.

   ```
   aaa authentication ssh login <local|tacacs|radius|[<local|none>]
   ```

2. Configures a password method for the primary and secondary enable (manager) access.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

If you do not specify an optional secondary method, it defaults to `none`. If the primary method is `local`, the secondary method must be `none`.

```
aaa authentication ssh enable <local|tacacs|radius>[<local|none>]
```

Option B: Configuring the switch for client Public-Key SSH authentication

Using this option results in the switch also authenticating the client's public key.

If configured with this option, the switch uses its public key to authenticate itself to a client, but the client must also provide a client public key for the switch to authenticate. This option requires the additional step of copying a client public-key file from a TFTP or SFTP server into the switch.

With the prerequisite steps complete and SSH properly configured on the switch, if an SSH client contacts the switch, login authentication automatically occurs first, using the switch and client public keys. After the client gains login access, the switch controls client access to the manager level by requiring the passwords configured earlier by the `aaa authentication ssh enable` command.

NOTE   Hewlett Packard Enterprise recommends that you always assign a manager-level (enable) password to the switch. Without this level of protection, any user with Telnet, web, or serial port access to the switch can change the switch configuration. If you configure only an operator password, entering the operator password through telnet, web, ssh or serial port access enables full manager privileges. See **1** on page 356.

Prerequisites

Before you can use this option, you must do the following:

1. Create a key pair on an SSH client.

2. Copy the client's public key into a public-key file (which can contain up to 10 client public keys.)

3. Copy the public-key file into a TFTP or SFTP server accessible to the switch and download the file to the switch.

Procedure

1. Copy the public-key file into the switch.

```
copy tftp pub-key-file < ip-address > < filename >
```

2. Configure the switch to authenticate a client public key at the login level with an optional secondary password method.

```
aaa authentication ssh login public-key
```

Default: none

3. Configure a password method for the primary and secondary enable (manager) access. If you do not specify an optional secondary method, it defaults to `none`.

```
aaa authentication ssh enable <local|tacacs|radius> <local|none>
```

If the primary access method is `local`, you can only specify `none` for a secondary access method.

NOTE

The configuration of SSH clients' public keys is stored in flash memory on the switch. You also can save SSH client public-key configurations to a configuration file by entering the following commands:

a. `include-credentials`

b. `write memory`

For more information about saving security credentials to a configuration file, see **Saving security credentials in a config file** on page 52.

Configuring for SSH access requiring a client public-key match and manager passwords

Assume you have a client public-key file named Client-Keys.pub (on a TFTP server at 10.33.18.117) ready for downloading to the switch:

- For SSH access to the switch allow only clients having a private key that matches a public key found in Client-Keys.pub.

- For manager-level (enable) access for successful SSH clients, use TACACS+ for primary password authentication and `local` for secondary password authentication, with a manager username of "leader" and a password of "m0ns00n".

To set up this operation, configure the switch in a manner similar to the following illustration:


```
HP Switch(config)# password manager user-name leader
New password for Manager: ********
Please retype new password for Manager: ********
HP Switch(config)# aaa authentication ssh login public-key none
HP Switch(config)# aaa authentication ssh enable tacacs local
HP Switch(config)# coy tftp pub-key-file 10.33.18.117
HP Switch(config)# write memory
```

Configures Manager user-name and password.

Configures the switch to allow SSH access only for a client whose public key matches one of the keys in the public key file.

Copies a public key file named "Client-Keys.pub" into the

Configures the primary and secondary password methods for Manager (enable) access. (Becomes available after SSH access is granted

The following illustration shows how to check the results of the above commands.

```
HP Switch(config)# show authentication          ◄─────        Lists the current SSH
                                                              authentication
 Status and Counters - Authentication Information            configuration.

  Login Attempts : 2
  Respect Privilege : Disabled

            | Login       Login        Login
 Access Task | Primary     Server Group Secondary
 ----------- + ---------- ------------ ----------
 Console     | Local                    None
 Telnet      | Local                    None
 Port-Access | Local                    Authorized
 Webui       | Local                    None
 SSH         | PublicKey                None
 Web-Auth    | ChapRadius  radius       Authorized
 MAC-Auth    | ChapRadius  radius       None


            | Enable      Enable       Enable
 Access Task | Primary     Server Group Secondary      Shows the contents of
 ----------- + ---------- ------------ ---------       the public key file
 Console     | Local                    None           downloaded with the
 Telnet      | Local                    None           copy tftp command in
 Webui _ _ _ | Local_ _ _ _ _ _ _ _ _  None_           figure -10. In this
( SSH  _ _ _ | Tacacs _ _ _ _ _ _ _ _  Local )         example, the file
 `─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─       contains two client
                                                       public-keys.
           [ Client Key Index Number ]

HP Switch(config)# show crypto client-public-key
0, Maden name [1024-bit rsa, Local_cryto @ Local cryto, Thu Nov 07 2009
21:25:42]" ssh-rsa
AAAAB3NzaClyc2EAAAADAQABAAAAgQcz9oNfqxMHUFEC6frSu1Sa4Dh1EFznFhQqmgP2
9HXYp6NR/1QOumACtrFU+QD11Etm/X.xZH/
KIxTdEc5exFZXKIxTdEc5exFXS10tcRaFYxI9UjK80dBMavBGKBLvVebCVwlqdAqbkaEX3d/
WaPS2xArLCFHsTZhnCvqTZDOGAB1frlcw==1,"[768-bit rsa, Local_crypto@Localcrypto,
Mon Dec 16 2009 23:01:51]"ssh-rsaAAAAb3NzaClyc2EAAAADQABAAAAYQD0tmzA32JBgeu
FJN0iXI3bfooPKZ09JKCcpQcXEVK7N+eKf9M0XvnmfFuEpw/fpqhlvsE66n8FDu7W/B2tKH/tqQ
```

SSH client contact behavior

At the first contact between the switch and an SSH client, if the switch public key has not been copied into the client, then the client's first connection to the switch will question the connection and, for security reasons, provide the option of accepting or refusing. If it is safe to assume that an unauthorized device is not using the switch IP address in an attempt to gain access to the client's data or network, the connection can be accepted. (As a more secure alternative, the client can be directly connected to the switch serial port to download the switch public key into the client.)

When an SSH client connects to the switch for the first time, it is possible for a "man-in-the-middle" attack; that is, for an unauthorized device to pose undetected as the switch, and learn the usernames and passwords controlling access to the switch. This possibility can be removed by directly connecting the management station to the switch serial port, using a show command to display the switch public key, and copying the key from the display into a file. This requires a knowledge of where the client stores public keys, plus the knowledge of what key editing and file format might be required by the client application. However, if the first contact attempt between a client and the switch does not pose a security problem, this is unnecessary.

Disable username prompt for management interface authentication in the Quick Base system

Authentication bypass for the username when logging on to the switch.

Syntax:

```
[no] aaa authentication disable-username
```

Use the `disable-username` command in manager mode to by-pass prompting username at the authentication feature of the switch. Execution of the above command on the switch enables the switch to bypass prompting username if applicable and only password is used for authentication only when both operator and manager usernames are default. The `disable-username` command will display an acceptance option and security warning message similar to the following:

```
Bypass username during authentication while logging
into the device to get Manager or Operator access? (y/n)
```

Entering "**y/Y**" will complete setting the command.

The default configuration of the switch is to authenticate based on username and password. This command is disabled by default in the switch configuration. An event log message (warning) will be logged that relates to this command. The event log message with be similar to the following:

```
W 05/22/13 21:02:06 00419 auth: Bypassing the username
for Operator and Manager access level is enabled
```

> **NOTE**
>
> The protocols SFTP and SCP which are based on SSH will exhibit similar behavior as SSH.

Switch behavior with Telnet

The behavior of the switch software in different management interfaces when the `disable-username` command is enabled.

Telnet Example: 1:

```
Operator password: Set          Manager password: Not set
Operator username: Not set      Manager username: Not set
```

• Entering the operator password will log onto the switch in operator mode.

• Entering the `enable` command in operator mode will provide access to the switch in manager.

Telnet Example: 2:

```
Operator password: Not set      Manager password: Set
Operator username: Not set      Manager username: Not set
```

• Entering the manager password will log onto the switch in manager mode.

• Entering the `enable` command in the operator mode will prompt for the password to be entered for manager mode.

Telnet Example: 3:

```
Operator password: Set          Manager password: Set
Operator username: Not set      Manager username: Not set
```

- Entering the operator password will log onto the switch in operator mode.

- Entering the `enable` command in the operator mode will prompt for the password to be entered for manager mode.

- Entering the manager password will log onto the switch in manager mode.

Telnet Example: 4:

```
Operator password: Set       Manager password: Set
Operator username: Not set    Manager username: Not set
```

- Both Operator and Manager Passwords are the same

- Entering the common password will log onto the switch in manager mode

Figure 72: *User login screen*


```
Copyright (C) 1991-2013 Hewlett-Packard Development Company, L.P.

                    RESTRICTED RIGHTS LEGEND
 Confidential computer software.  Valid license from HP required for possession,
 use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer
 Software, Computer Software Documentation, and Technical Data for Commercial
 Items are licensed to the U.S. Government under vendor's standard commercial
 license.
                    HEWLETT-PACKARD DEVELOPMENT COMPANY, L.P.
                    20555 State Highway 249, Houston, TX 77070


We'd like to keep you up to date about:
   * Software feature updates
   * New product announcements
   * Special events
Please register your products now at:  www.hp.com/networking/register








Username: █
```

Switch behavior with SSH

There is no username prompt in SSH, (for example: `ssh any_username@IP-address`). The following configuration examples are related to SSH:

SSH configuration Example: 1:

In the following configuration, the password is entered without the username. Once the enable command is entered, the username prompt will be bypassed.

```
Operator password: Set            Manager password: Not set
Operator username: Not set        Manager username: Not set
```

- Entering the operator password will log onto the switch in operator mode.

- Entering the `enable` command in the operator mode will prompt for the password to be entered for manager mode.

SSH configuration Example: 2:

```
Operator password: Not set        Manager password: Set
Operator username: Not set        Manager username: Not set
```

- Executing `any key` will log onto the switch in operator mode.

- Entering the `enable` command in the operator mode will prompt for the password to be entered for manager mode

- The manager password will log onto the manager mode.

SSH configuration Example: 3:

```
Operator password: Set            Manager password: Set
Operator username: Not set        Manager username: Not set
```

- Entering the operator password will log onto the switch in operator mode.

- Entering the `enable` command in the operator mode will prompt for the password to be entered for manager mode.

SSH configuration Example: 4:

```
Operator password: Set            Manager password: Set
Operator username: Not set        Manager username: Not set
```

Both Operator and Manager passwords are the same

Entering the password logs onto the switch in the manager mode.

Figure 73: *SSH configuration screen*

Switch behavior with WebUI

The user is prompted for both username and password. Any entry including blank in the username field is allowed. This condition is true for all the below configuration.

WebUI configuration Example: 1:

```
Operator password: Set          Manager password: Not set
Operator username: Not set      Manager username: Not set
```

Entering only the operator password, logs onto the switch in manager mode because manager password has not been set.

WebUI configuration Example: 2:

```
Operator password: Not set       Manager password: Set
Operator username: Not set       Manager username: Not set
```

• Selecting the login link on the WebUI will prompt for log in.

• Entering only the manager password, logs onto the switch in manager mode.

WebUI configuration Example: 3:

```
Operator password: Set          Manager password: Set
Operator username: Not set      Manager username: Not set
```

- Logging in using the operator password logs onto the switch in operator mode.

- Entering only the manager password, logs onto the switch in manager mode.

WebUI configuration Example: 4:

```
Operator password: Set              Manager password: Set
Operator username: Not set          Manager username: Not set
```

Both Operator and Manager passwords are the same

Entering only the password, logs onto the switch in manager mode.

Figure 74: *WebUI logon screen*

NOTE: This feature does not capture solutions for other applications that mandate the username and does not capture the solution based on the type of management interface used.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

SSH client public-key authentication notes

When configured for SSH operation, the switch automatically attempts to use its own host public key to authenticate itself to SSH clients. To provide the optional, opposite service—client public-key authentication to the switch—you can configure the switch to store up to 10 public keys for authenticating clients. This requires storing an ASCII version of each client's public key (without babble conversion, or fingerprint conversion) in a client public-key file that you create and TFTP-copy to the switch. In this case, only clients with a private key corresponding to one of the stored public keys can gain access to the switch using SSH. If you use this feature, only the clients whose public keys are in the client public-key file you store on the switch will have SSH access to the switch over the network. If you do not allow secondary SSH login (operator) access via local password the switch will refuse other SSH clients.

SSH clients that support client public-key authentication normally provide a utility to generate a key pair. The private key is usually stored in a password-protected file on the local host; the public key is stored in another file and is not protected.

> **NOTE**
> Without using client public-key authentication you can still require authentication from whoever attempts to access the switch from an SSH client— by employing the local username/password, TACACS+, or RADIUS features. See **5** on page 357.

If you enable client public-key authentication, the following events occur when a client tries to access the switch using SSH:

1. The client sends its public key to the switch with a request for authentication.

2. The switch compares the client's public key to those stored in the switch client public-key file. As a prerequisite, use the switch `copy tftp` command to download this file to flash.

3. If there is no match, and you have not configured the switch to accept a login password as a secondary authentication method, the switch denies SSH access to the client.

4. If there is a match, the switch:

   a. Generates a random sequence of bytes.

   b. Uses the client's public key to encrypt this sequence.

   c. Send these encrypted bytes to the client.

5. The client uses its private key to decrypt the byte sequence.

6. The client then:

   a. Combines the decrypted byte sequence with specific session data.

   b. Uses a secure hash algorithm to create a hash version of this information.

   c. Returns the hash version to the switch.

7. The switch computes its own hash version of the data from step 6 and compares it to the client's hash version. If they match, the client is authenticated. Otherwise, the client is denied access.

Using client public-key authentication

Procedure

1. Generate a public/private key pair for each client you want to have SSH access to the switch. This can be a separate key for each client or the same key copied to several clients.

2. Copy the public key for each client into a client public-key text file.

3. Use `copy tftp` to copy the client public-key file into the switch. Note that the switch can hold 100 keys. The new key is appended to the client public-key file

4. Use the `aaa authentication ssh` command to enable client public-key authentication.

Creating a client public-key text file

These steps describe how to copy client public keys into the switch for challenge-response authentication, and require an understanding of how to use your SSH client application.

Figure 75: *A client public key*

> **NOTE**
>
> Comments in public-key files may appear in a SSH client application's generated public key. While such comments may help to distinguish one key from another, they do not pose any restriction on the use of a key by multiple clients and/or users.
>
> Public key illustrations such as the key shown in **Figure 75: A client public key** on page 374
>
> usually include line breaks as a method for showing the whole key. However, in practice, line breaks in a public key will cause errors resulting in authentication failure.

Procedure

1. Use your SSH client application to create a public/private key pair, and see the documentation provided with your SSH client application for details. The switch supports the following client public-key properties:

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Property	Supported value	Comments
Key format	ASCII	See **Providing the switch public key to clients**. The key must be one unbroken ASCII string. If you add more than one client public key to a file, terminate each key (except the last one) with a <CR><LF>. Spaces are allowed within the key to delimit the key's components. Note that, unlike the use of the switch public key in an SSH client application, the format of a client public key used by the switch does not include the client's IP address.
Key type	RSA or DSA	You can choose either RSA or DSA key types when using the `crypto key generate ssh` command. The `cert` and `autorun` parameters only use RSA key types.
Maximum supported public-key length	3072 bits	Shorter key lengths allow faster operation, but also mean diminished security.
Maximum host key sizes in bits	RSA:1024, 2048, 3072DSA:1024	Includes the bit size, public index, modulus, any comments,<CR>, <LF>, and all blank spaces. If necessary, you can use an editor application to verify the size of a key. For example, placing a client public key into a Word for Windows text file and clicking on `File|Properties|Statistics` lets you view the number of characters in the file, including spaces.

2. Copy the client's public key into a text file (*filename.txt*). For example, use the Notepad editor included with the Microsoft® Windows® software. If you want several clients to use client public-key authentication, copy a public key for each of these clients (10) into the file. Each key should be separated from the preceding key by a **<CR><LF>**.

3. Copy the client public-key file into a TFTP server accessible to the switch.

   Copying a client public key into the switch requires the following:

   - One or more client-generated public keys. See the documentation provided with your SSH client application.

   - A copy of each client public key stored in a single text file or individually on a TFTP server to which the switch has access. Terminate all client public keys in the file except the last one with a **<CR><LF>**.

> **NOTE**
> The actual content of a public-key entry in a public-key file is determined by the SSH client application generating the key. Although you can manually add or edit any comments the client application adds to the end of the key.

Syntax:

```
copy <tftp|sftp> pub-key-file <ip-address> <filename>[<append|manager|operator>]
```

Copies a public-key file from a TFTP server into flash memory in the switch.

The `append` option adds the keys for operator access.

The `manager` option replaces the keys for manager access; follow with the 'append' option to add the keys.

The `operator` option replaces the keys for operator access (default); follow with the 'append' option to add the keys.

```
show crypto client-public-key [<manager|operator>][keylist-str][babble|fingerprint]
```

Displays the client public keys in the switch current client public-key file.

See **SSH client public-key authentication** on page 60 for information about public keys saved in a configuration file.

The `babble` option converts the key data to phonetic hashes that are easier for visual comparisons.

The `fingerprint` option converts the key data to hexadecimal hashes that are for the same purpose.

The `keylist-str` selects keys to display (comma-delimited list).

The `manager` option allows you to select manager public keys

The `operator` option allows you to select operator public keys.

NOTE	SFTP can also be used to copy a public key file to the switch.

Copying and displaying a client public-key file containing two different client public keys for the same client

To copy a client public-key file named `clientkeys.txt` from a TFTP server at 10.38.252.195 and then display the file contents:

```
HP Switch(config)# copy tftp pub-key-file 10.33.252.195 Clientkeys.txt
HP Switch(config)# show crypto client-public-key
0 "Maden name [1024-bit rsa, Jamie_wilson@Jamiewilson, Thu Nov 07 2002 21:25:41
2]" ssh-rsa AAAAB3NzaClyc2EAAAADAOABAAAAq0Cz9oNfqxllIUFEC6SulSa4UhlEFznlh0qmqP2
91DYp6l\lR/lO0UmACtrFUODllE14HQIN/XVZH/kIxTdEo5exFX/Sl0toRaFYzI9UjK80dBllqvBCKB
IYVELJCVwlqdAqbkaEX3d/Na?S2xArLCFHslZhnCVQTZDOC-ABlfrlou==
1 "[768-bit rsa, Jamie_wilson@Jamiewilson, Mon Dec 16 2002 23:01:51]" ssh-rsa
AAAAB3NzaClyoZEAAAADAOABAAAAYQDOtmzA32JBgeuFJNOiXI3bfooPKZ09KCPQoCEVk7NKf9HOX
vnmfF\lEpu/fpqhlvsE66n8FDu7U/B2tK1I130LFOtXCiVcNCLN0/Do5AuEvmEnclCu/LgAH9daM=

Key Index Number
```

Replacing or clearing the public-key file

The client public-key file remains in the switch flash memory even if you erase the startup-config file, reset the switch, or reboot the switch.

Remove the existing client public-key file or specific keys by executing the `clear crypto public-key` command. This clears the public keys from both management modules. The module that is not active must be in standby mode.

Syntax:

```
clear crypto public-key
```

Deletes the client public-key file from the switch.

Syntax:

```
clear crypto public-key 3
```

Deletes the entry with an index of 3 from the client public-key file on the switch.

Enabling client public-key authentication

After you TFTP a client public-key file into the switch, configure the switch to allow the following:

If an SSH client's public key matches the switch client public-key file, allow that client access to the switch. If there is not a public-key match, then deny access to that client.

Syntax:

```
aaa authentication ssh login public-key none
```

Allows SSH client access only if the switch detects a match between the client's public key and an entry in the client-public- key file most recently copied into the switch.

> **CAUTION**
>
> To enable client public-key authentication to block SSH clients whose public keys are not in the client public-key file copied into the switch, you must configure the Login Secondary as `none`. Otherwise, the switch allows such clients to attempt access using the switch operator password.

SSH client and secure sessions

SSH Client provides a method for establishing a secure session from one switch to another. In addition to providing secure sessions, SFTP is enhanced to allow bidirectional secure copying of files between a switch and an SFTP server, initiated from the switch with the copy command. The SFTP server can be another switch or a workstation/server with a running SSH server that supports SFTP.

Each switch with the SSH Client feature will have a known hosts file that can contain the public key from switches and servers that have been determined to be genuine. New public keys can be added to the known hosts file when new SSH servers are contacted, up to a maximum of 10 entries (if memory allows). The known hosts file can also be copied to another switch or to a server where it can be edited.

> **NOTE**
>
> You must be in manager context to use this SSH and SFTP feature.

Opening a secure session to an HPE switch

To initiate an SSH client session to another network device, use the following command, executed in the manager context.

Syntax:

```
ssh [user <username|username@>] <hostname|IPv4|IPv6> [port <1-65535>]
```

Enables an SSH client to open a secure session to an switch. Opening secure sessions to devices other than switches is not supported.

```
[user <username|username@>]
```

Optional, the username on the destination (remote) system. Usernames for operator and manager must be configured.

If `<username@>` is specified, enter the remote system information of either `hostname`, IPv4 or IPv6 address.

If no `username` is configured, the `username` of the current login is used. There will be a prompt for a password if needed.

`hostname`: Hostname of the remote system.

`IPv4`: IPv4 address of the remote system.

`IPv6`: IPv6 address of the remote system.

`[port <1-65535>]`: The TCP port running the SSH server on the remote system. If no port number is specified, the default port 22 is used.

Default: Port 22

General operating rules and notes

- The SSH server may challenge the client to authenticate itself depending on the authentication methods configured on the destination SSH server. The client first tries the "none" method of authentication; if that is unsuccessful, it examines the list of supported authentication methods from the server, if provided. If the server does not provide such a list, all methods of authentication will be tried in the following order until the session is successfully opened or rejected by the server:

  ◦ Authentication method "publickey", if a private key has been loaded onto the switch.

  ◦ Authentication method "password".

- During "public-key" authentication, the client must use its private key to authenticate itself to the server. There can be only one key pair on the switch for the manager.

- The private key should be passphrase protected for highest security; the user is prompted to enter the passphrase.

- The private key can be configured by copying it to the SSH client switch, using the `copy` command.

- If the public-key authentication fails or the client has not been configured with a key pair, the "password" method of authentication is used and the user is prompted for a password.

- Successful TACACS or RADIUS logins will give the user either operator or manager privileges. This is important if there are chained SSH sessions.

Copying client key files

Only one ssh client key for authenticating the manager is allowed on a switch. The `copy` command allows you to copy the client key files using `sftp`, `tftp`, and `usb` or `xmodem`, allowing encryption and authentication through SSH. There is no way to generate the private key on the switch; it must be copied onto the switch.

To load the client's private key onto the switch, use one of the following commands:

Syntax:

```
copy sftp ssh-client-key [user <username>|<username@>] <hostname|IPv4|IPv6> <private-key-filename> [port|<tcp-port-num>]

copy tftp ssh-client-key <hostname|IPv4|IPv6> <private-key-filename>

copy usb ssh-client-key <private-key-filename>

copy xmodem ssh-client-key
```

Copies the client key file <private-key-filename>onto the switch.

`ssh-client-key`: The client key file being copied to the switch. The file must contain an RSA or DSA key.

`[user <username|username@>]`: Optional, there must be configured usernames for operator and manager.

If no `username` is specified, the client's current `username` is used. There will be a prompt for a password if needed.

`hostname`: Specifies the hostname of the SFTP or TFTP server.

`IPv4`: Specifies the SFTP or TFTP server's IPv4 address.

`IPv6`: Specifies the SFTP or TFTP server's IPv6 address.

`<private-key-filename>` : The remote filename containing the key.

`[port <tcp-port-num>]` : TCP port of the SSH server on the remote system.

The copied private key is stored in the ssh directory of the switch file system and is persistent across switch reboots. After the initial copying is complete, the client key can be overwritten by repeating the `copy` command. No verification of the validity of the key is done when executing the `copy` command.

Copying the ssh-client-known-hosts file

Replacing or appending the ssh-client-known-hosts file

There is one SSH client known hosts file per switch, stored in the ssh directory of the switch file system. The SSH client known hosts file can be overwritten or appended using the commands below. If the file already exists at that location, you are prompted for confirmation before overwriting the existing file:

```
Warning: The existing known hosts file will be
overwritten, continue [y/n] ?
```

Syntax:

```
copy sftp ssh-client-known-hosts [user <username|username@>] <hostname|IPv4|IPv6> <filename> [append]
```

```
copy tftp ssh-client-known-hosts <hostname|IPv4|IPv6> <filename> [append]
```

```
copy usb ssh-client-known-hosts <filename> [append]
```

```
copy xmodem ssh-client-known-hosts [append]
```

Copies the SSH client known hosts file to the switch.

`ssh-client-known-hosts`: The known hosts file.

`[user <username|username@>]`: Optional, there must be configured usernames for operator and manager.

If no `username` is specified, the client's current `username` is used. There will be a prompt for a password if needed.

`hostname`: Specifies the hostname of the TFTP server.

`IPv4`: Specifies the TFTP server's IPv4 address.

`IPv6`: Specifies the TFTP server's IPv6 address.

`<filename>` : The remote filename containing the key.

`[append]`: Append hosts to the existing `ssh-client-known hosts` file.

Default: Replace the existing known hosts file.

Copying the SSH client known hosts file to another location

The SSH client known hosts file can be copied to a location on a remote device using the commands below. If the file already exists on the remote device, you are prompted for confirmation before overwriting the file:

```
Warning: The remote file will be overwritten,
continue [y/n] ?
```

Syntax:

```
copy ssh-client-known-hosts sftp [user <username|username@>] <hostname|IPv4|IPv6> <filename>

copy ssh-client-known-hosts tftp <hostname|IPv4|IPv6> <filename>

copy ssh-client-known-hosts <filename> usb

copy ssh-client-known-hosts xmodem
```

Copies the SSH client known hosts file to another location.

`ssh-client-known-hosts`: The known hosts file.

`[user <username|username@>]`: Optional, there must be configured usernames for operator and manager.

If no `username` is specified, the client's current `username` is used. There will be a prompt for a password if needed.

`hostname`: Specifies the hostname of the TFTP server.

`IPv4`: Specifies the TFTP server's IPv4 address.

`IPv6`: Specifies the TFTP server's IPv6 address.

`<filename>` : The remote filename containing the key.

Copying the host public key

The following `copy` commands can be used to manage public keys in a known hosts file. The public key can only be coped from the switch to a server or other media, not to another switch.

Syntax:

```
copy ssh-server-pub-key sftp [user <username|username@>] <hostname|IPv4|IPv6> <filename>

copy ssh-server-pub-key tftp <hostname|IPv4|IPv6> <filename>

copy ssh-server-pub-key usb

copy ssh-server-pub-key xmodem
```

Copies the switch SSH server public key to a server or other media.

`[user <username|username@>]`: Optional, there must be configured usernames for operator and manager.

If no `username` is specified, the client's current `username` is used. There will be a prompt for a password if needed.

`hostname`: Specifies the hostname of the TFTP server.

`IPv4`: Specifies the TFTP server's IPv4 address.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

`IPv6`: Specifies the TFTP server's IPv6 address.

`<filename>` : The name of the file to be copied.

Removing the SSH client key pair

To remove the SSH client key pair file, use this command.

Syntax:

```
crypto key zeroize ssh-client-key
```

Deletes the SSH client key pair.

You are prompted with the message:

```
Warning: The manager key pair will be deleted,
continue [y/n] ?
```

Removing the SSH client known hosts file

To remove the SSH client known hosts file, use this command.

Syntax:

```
crypto key zeroize ssh-client-known-hosts
```

Deletes the SSH client known hosts file.

You are prompted with the message:

```
Warning: The SSH client known hosts file
will be deleted, continue [y/n] ?
```

Displaying open sessions

Syntax:

```
show session-list
```

Displays the active incoming and outgoing sessions.

Open sessions listing

```
HP Switch(config)# show session-list

Session Information

Source IP Selection: Outgoing Interface
--------------------------------------------------------
Session  : 1
Privilege: Superuser
From     : Console
To       :
--------------------------------------------------------
Session  : 2
Privilege: manager
From     : 172.22.16.3
To       : 10.1.248.198
--------------------------------------------------------
Session  : ** 3
Privilege: manager
```

```
From      : 10.1.248.179
To        :
```

Messages related to SSH operation

Message	Meaning
`00000K Peer unreachable.`	File transfer did not occur. Indicates an error in communicating with the tftp server or not finding the file to download. Causes include such factors as: • Incorrect IP configuration on the switch • Incorrect IP address in the command • Case (upper/lower) error in the filename used in the command • Incorrect configuration on the TFTP server • The file is not in the expected location. • Network misconfiguration • No cable connection to the network
`00000K Transport error.`	File transfer did not occur. Indicates the switch experienced a problem when trying to `copy tftp` the requested file. The file may not be in the expected directory, the filename may be misspelled in the command, or the file permissions may be wrong.
`Cannot bind reserved TCP port <port-number>.`	The ip SSH port command has attempted to configure a reserved TCP port. Use the default or select another port number.
`Client public-key file corrupt or not found. Use 'copy tftp pub-key-file <ip-addr> <filename>' to download new file.`	The client key does not exist in the switch. Use `copy tftp` to download the key from a TFTP server.

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Message	Meaning
`Download failed: overlength key in key file.Download failed: too many keys in key file.Download failed: one or more keys is not a valid public key.`	The public-key file you are trying to download has one of the following problems: • A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR><LF>. • There are more than 10 public keys in the key file and switch total. Delete some keys from the switch or file. The switch does not detect duplicate keys. • One or more keys in the file is corrupted or is not a valid public key. See **Creating a client public-key text file** on page 374 for information on client public-key properties.
`Error: Requested keyfile does not exist.`	The client key does not exist in the switch. Use `copy tftp` to download the key from a TFTP server.
`Generating new RSA host key. If the cache is depleted, this could take up to two minutes.`	After you execute the `generate ssh [dsa|rsa]` command, the switch displays this message while it is generating the key.
`Host RSA key file corrupt or not found. Use 'generate ssh [dsa|rsa]' to create new host key.`	The switch key is missing or corrupt. Use the `generate ssh [dsa|rsa]` command to generate a new key for the switch.

Logging messages

There are event log messages when a new key is generated and zeroized for the server:

```
ssh: New <num-bits> -bit [rsa|dsa] SSH host key installed
```

```
ssh: New SSH host key zeroized
```

There are also messages that indicate when a client public key is installed or removed:

```
ssh: <num-bits> -bit [rsa|dsa] client public key [installed|removed] ([manager|operator]access)(key_comment)
```

NOTE   Only up to 39 characters of the key comment are included in the event log message.

Debug logging

To add ssh messages to the debug log output, enter this command:

```
HP Switch# debug ssh LOGLEVEL
```

where LOGLEVEL is one of the following (in order of increasing verbosity):

- fatal
- error
- info
- verbose
- debug
- debug2
- debug3

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Overview

A common access card (CAC) is a United States Department of Defense (DoD) smart card for multifactor authentication. CACs are issued as standard identification for active-duty military personnel, reserve personnel, civilian employees, non-DoD government employees, state employees of the National Guard, and eligible contractor personnel. In addition to its use as an ID card, a CAC is required for access to government buildings and computer networks.

Part of the requirement necessary to satisfy the Federal Government Certification (JITC requirements) is two-factor authentication. Two-factor authentication is the redundant authentication of the CAC. For example, the CAC satisfies two-factor authentication by mandating that you have both the physical card and know the pin number associated with the card.

To provide support for CAC authentication, the requirement for the network is the establishment of SSH connections. Two-factor authentication constitutes authentication based on public key or certificate and username/password on the switch.

See also "Common access card (two-factor) authentication" in the basic operation guide for your switch.

Several commands are provided to configure two-factor authentication.

Two-factor authentication configuration commands

aaa authentication ssh

Syntax

```
aaa authentication ssh [enable | login]
```

Description

Configure authentication mechanism used to control SSH access to the switch.

> **NOTE**
> This command must be used before using any of the two-factor forms of the `aaa authentication ssh` command,

Options

enable

Configure access to the privileged mode commands.

login

Configure login access to the switch.

aaa authentication ssh two-factor

Syntax

```
aaa authentication ssh [enable | login]
 two-factor [local | none | authorized |
```

```
   server-group <server-group> | two-factor-type]
```

Description

Set two-factor authentication method as the primary authentication method.

Options

local

Use local switch user/password database.

none

Do not use backup authentication methods.

authorized

Allow access without authentication.

server-group

Specify the server group to use.

two-factor-type

Use the certificate or public key for the first authentication method and username/password for the second authentication method.

aaa authentication ssh two-factor two-factor-type

Syntax

```
aaa authentication ssh [enable | login]
 two-factor two-factor-type [publickey-password | certificate-password]
```

Description

Use the certificate or public key for the first authentication method and username/password for the second authentication method.

Opions

publickey-password

Use the public key for the first authentication method and username/password for the second authentication method.

certificate-password

Use the X.509v3 certificate for the first authentication method and username/password for the second authentication method.

aaa authentication ssh two-factor two-factor-type publickey-password

Syntax

```
aaa authentication ssh [enable | login] two-factor two-factor-type
 publickey-password [local | tacacs | radius]
```

Description

Use the public key for the first authentication method and username/password for the second authentication method.

Options

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

local

Use local switch user/password database.

tacacs

Use TACACS+ server.

radius

Use RADIUS server.

aaa authentication ssh two-factor two-factor-type certificate-password

Syntax

```
aaa authentication ssh [enable | login] two-factor two-factor-type
 certificate-password [local | tacacs | radius]
```

Description

Use the X.509v3 certificate for the first authentication method and username/password for the second authentication method.

Options

local

Use local switch user/password database.

tacacs

Use TACACS+ server.

radius

Use RADIUS server.

crypto enforce secure-rsa

Syntax

```
crypto enforce secure-rsa
```

Description

Enable generation of secure RSA key size. Only secure keys will be generated using the RSA key. The recommended secure RSA key size is 2048.

Two-factor authentication restrictions

- When an SSH client establishes a connection by choosing the user authentication method **password** or **public-key** and **password**, the switch will terminate the connection if two-factor authentication or password configuration-control is enabled.

- For successful authentication when Two-factor authentication is enabled, the user authentication method must be **public-key** and **keyboard interactive**.

- When password configuration-control alone is enabled, the user authentication method must include **keyboard interactive**.

Two-factor authentication validation rules

Validation	Error/Warning/Prompt
If the authentication method is being set to two-factor authentication, various messages display.	If both the public key and username/password are not configured: Public key and username/password should be configured for a successful two-factor authentication. If public key is configured and username is not configured: Username and password should be configured for a successful two-factor authentication. If the username is configured and public key is not configured: Public key should be configured for a successful two-factor authentication. If "ssh-server" certificate is not installed at the time of enabling certificate-password authentication: The "ssh-server" certificate should be installed for a successful two-factor authentication.
If the authentication method is set to two-factor while installing the public key, a message displays.	The client public keys without username will not be considered for the two-factor authentication for the SSH session.
If the username and the key installation user for that privilege do not match, a message displays and installation is not allowed. This will also happen when the authentication method is set for two-factor.	The username in the key being installed does not match the username configured on the switch.
If secondary authentication type for two-factor authentication chosen is not none, a message displays.	Not legal combination of authentication methods.
If the authentication method is anything other than two-factor and the two-factor authentication method options are set, a message displays.	Not legal combination of authentication methods.
If two-factor authentication is set and user tries to SSH into another system using `ssh <IP	HOSTNAME>` command, a message displays.	SSH client is not supported when the two-factor authentication is enabled.
When RSA secure key is configured.	**** CAUTION **** Enabling secure RSA key feature will only allow generation of 2048 bits or higher bit as 1024 bits or lower bit are deprecated. Continue (y/n)?

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Two-factor authentication event log messages

Event	Message
RMON_AUTH_TWO_FACTOR_AUTHEN_STATUS	
	W 01/01/15 18:24:03 03397: auth: %s.
	Examples:
	W 01/01/15 18:24:03 03397: auth: Public key and username/password should be configured for the successful two-factor authentication.
	W 01/01/15 18:24:03 03397: auth: Username and password should be configured for the successful two-factor authentication.
	W 01/01/15 18:24:03 03397: auth: Public key should be configured for the successful two-factor authentication.
	I 01/01/15 18:24:03 03397: auth: The validation of certificate of SSH user (user1) is successful.
RMON_SSH_KEY_TWO_FACTOR_EN	
	W 01/01/15 18:24:03 03399: ssh: %s.
	Examples:
	W 01/01/15 18:24:03 03399: ssh: The client public keys without username will not be considered for the two-factor authentication for SSH session.
	W 01/01/15 18:24:03 03399: ssh: The privilege level for the user with the SSH key conflicts with the user configured.
RMON_SSH_TWO_FACTOR_AUTH_FAIL	
RMON_CERT_KEY_DEFAULT_MODIFIED	W 01/01/15 18:24:03 03398: ssh: %s.
	Examples:
	W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in public key authentication.
	W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in username/password authentication.
	W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed due to the failure in validating the client certificate.
	W 01/01/15 18:24:03 03398: ssh: The two-factor authentication for SSH session failed as "ssh-server" certificate is not installed.
	W 05/22/13 20:39:20 03362 crypto:
	Default RSA key size changed to %s for %s application.

Overview

HPE Switches use Secure Sockets Layer Version 3 (SSLv3) and support for Transport Layer Security (TLSv1) to provide remote web access to the switches through encrypted paths between the switch and management station clients capable of SSL/TLS operation.

- Switches use SSL and TLS for all secure web transactions, and all references to SSL mean using one of these algorithms unless otherwise noted.

- SSL provides all the web functions but, unlike standard web access, SSL provides encrypted, authenticated transactions. The authentication type includes server certificate authentication with user password authentication.

SSL in the switches is based on the OpenSSL software toolkit. For more information on OpenSSL, visit **http://www.openssl.org**.

Server certificate authentication with user password authentication

This is a subset of full certificate authentication of the user and host, only available when the switch has SSL enabled. As in **Figure 76: Switch/user authentication** on page 390, the switch authenticates itself to SSL-enabled web browser. Users on SSL browser then authenticate themselves to the switch - operator and manager levels - by providing passwords stored locally on the switch or on a TACACS+ or RADIUS server. However, the client does not use a certificate to authenticate itself to the switch.

Figure 76: *Switch/user authentication*

SSL on the switches covered in this guide supports these data encryption methods:

- 3DES (168-bit, 112 Effective)

- DES (56-bit)

- RC4 (40-bit, 128-bit)

> **NOTE**
>
> Switches use RSA public-key algorithms and Diffie-Hellman, and all references to a key mean keys generated using these algorithms unless otherwise noted.

Prerequisite for using SSL

Before using the switch as an SSL server, you must install a publicly or commercially available SSL-enabled web browser application on the computer(s) you use for management access to the switch.

Steps for configuring and using SSL for switch and client authentication

Procedure

1. Client preparation:

    a. Install an SSL capable browser application on a management station you want to use for access to the switch. See the documentation provided with your browser for details.

    b. Assign a login (operator) and enable (manager) password on the switch. See **Assigning a local login (operator) and enabling (manager) passwords**.

2. Switch preparation:

    a. Generate a host certificate on the switch. See **Generating the switch's server host certificate**.

        • Generate certificate key pair

        • Generate host certificate

        You need to do this only once. The switch's own public/private certificate key pair and host certificate are stored in the switch flash memory and are not affected by reboots or the erase `startup-config` command. You can remove or replace this certificate, if necessary. The certificate key pair and the SSH key pair are independent of each other, which means a switch can have two keys pairs stored in flash.

3. Enable SSL on the switch. See **SSL client contact behavior**.

4. Use your SSL enabled browser to access the switch using the switch IP address or DNS name (if allowed by your browser). See the documentation provided with the browser application.

General operating rules and notes

• Once you generate a certificate on the switch avoid re-generating the certificate without a compelling reason. Otherwise, you will have to re-introduce the switch certificate on all management stations (clients) previously set up for SSL access to the switch. In some situations this can temporarily allow security breaches.

• The switch public/private certificate key pair and certificate are stored in the switch flash memory and are not affected by reboots or the `erase startup-config` command.

• The public/private certificate key pair is not be confused with the SSH public/private key pair. The certificate key pair and the SSH key pair are independent of each other, which means a switch can have two key pairs stored in flash.

Configuring the switch for SSL operation

Assigning a local login (operator) and enabling (manager) password

At a minimum, Hewlett Packard Enterprise recommends that you always assign at least a manager password to the switch. Otherwise, under some circumstances, anyone with Telnet, web, or serial port access could modify the switch's configuration.

Using the WebAgent to configure local passwords

You can configure both the operator and manager password in the WebAgent. To access the WebAgent, see "Using the HPE WebAgent" in the basic operation guide for your switch.

Generating the switch's server host certificate

You must generate a server certificate on the switch before enabling SSL. The switch uses this server certificate, along with a dynamically generated session key pair to negotiate an encryption method and session with a browser trying to connect via SSL to the switch. (The session key pair mentioned above is not visible on the switch. It is a temporary, internally generated pair used for a particular switch/client session, and then discarded.)

The server certificate is stored in the switch's flash memory. The server certificate should be added to your certificate folder on the SSL clients who you want to have access to the switch. Most browser applications automatically add the switch's host certificate to there certificate folder on the first use. This method does allow for a security breach on the first access to the switch. (See the documentation for your browser application.)

There are two types of certificated that can be used for the switch's host certificate. The first type is a self-signed certificate, which is generated and digitally signed by the switch. Since self-signed certificates are not signed by a third-party certificate authority, there is no audit trail to a root CA certificate and no fool-proof means of verifying authenticity of certificate. The second type is a certificate authority-signed certificate, which is digitally signed by a certificate authority, has an audit trail to a root CA certificate, and can be verified unequivocally.

> **NOTE**
>
> There is usually a fee associated with receiving a verified certificate and the valid dates are limited by the root certificate authority issuing the certificate.

When you generate a certificate key pair and/or certificate on the switch, the switch places the key pair and/or certificate in flash memory (and not in running config.) Also, the switch maintains the certificate across reboots, including power cycles. You should consider this certificate to be "permanent"; that is, avoid regenerating the certificate without a compelling reason. Otherwise, you will have to reintroduce the switch's host certificate on all management stations you have set up for SSL access to the switch using the earlier certificate.

Removing (zeroizing) the switch's certificate key pair or certificate render the switch unable to engage in SSL operation and automatically disables SSL on the switch. (To verify whether SSL is enabled, execute `show config`.)

To generate or erase the switch's server certificate with the CLI

Since the host certificate is stored in flash instead of the running-config file, it is not necessary to use `write memory` to save the certificate. Erasing the host certificate automatically disables SSL.

CLI commands used to generate a server host certificate

Syntax:

```
crypto key generate cert [rsa] bits <1024|2048>
```

Generates a key pair for use in the certificate.

```
crypto key zeroize cert
```

Erases the switch's certificate key and disables SSL operation.

To generate a host certificate from the CLI:

Generate a certificate key pair. This is done with the `crypto key generate cert` command. The default key size is 512.

> **NOTE**
> If a certificate key pair is already present in the switch, it is not necessary to generate a new key pair when generating a new certificate. The existing key pair may be reused and the `crypto key generate cert` command does not have to be executed.

Comments on certificate fields

There are a number of arguments used in the generation of a server certificate. The following table describes these arguments.

Table 28: *Certificate field descriptions*

Field Name	Description
Valid Start Date	This should be the date you desire to begin using the SSL functionality.
Valid End Date	This can be any future date, however good security practices would suggest a valid duration of about one year between updates of passwords and keys.
Common Name	This should be the IP address or domain name associated with the switch. Your web browser may warn you if this field does not match the URL entered into the web browser when accessing the switch
Organization	This is the name of the entity (e.g. company) where the switch is in service.
Organizational Unit	This is the name of the sub-entity (e.g. department) where the switch is in service.
City or Location	This is the name of the city where switch is in service
State Name	This is the name of the state or province where switch is in service
Country Code	This is the ISO two-letter country-code where switch is in service

Figure 77: *Self-signed server host certificate on the CLI for the switch*


```
HP Switch(config)# crypto host-cert generate self-signed          Generate New Certificate
Validity start date [05/17/2011]: 05/31/2011
Validity end date   [05/31/2012]: 06/15/2011
Common name     [10.115.134.111]: 10.255.255.255
Organizational unit  [Dept Name]: Hewlett-Packard          Enter certificate Arguments
Organization      [Company Name]: HP Networking
City or location           [City]: Roseville
State name                [State]: CA
Country code                 [US]: US
```

NOTE

"Zeroizing" the switch's server host certificate or key automatically disables SSL (sets `web-management ssl` to `No`). Thus, if you zeroize the server host certificate or key and then generate a new key and server certificate, you must also re-enable SSL with the web-management ssl command before the switch can resume SSL operation.

Generate a self-signed host certificate with the WebAgent

You can configure SSL from the WebAgent. For more information on how to access the WebAgent, see "Using the HPE WebAgent" in the management and configuration guide for your switch.

To generate a self signed host certificate from the WebAgent:

1. In the WebAgent navigation pane, click **Security**.

2. Click **SSL**.

3. In the Web Management box, enable SSL if it is not already checked.

4. Complete the fields in the SSL Certificate box and click **Create request**.

NOTE

When generating a self-signed host certificate, if no key is present and the current option is selected in the RSA key size box and error will be generated. New key generation can take up to two minutes if the key queue is empty.

Figure 78: *Self-signed certificate generation via SSL WebAgent*

Generate a CA-Signed server host certificate with the WebAgent

The installation of a CA-signed certificate involves interaction with other entities and consists of three phases. The first phase is the creation of the CA certificate request, which is then copied off from the switch for submission to the certificate authority. The second phase is the actual submission process that involves having the certificate authority verify the certificate request and then digitally signing the request to generate a certificate response (the usable server host certificate). The third phase is the download phase consisting of pasting to the switch web server the certificate response, which is then validated by the switch and put into use by enabling SSL.

To generate a certificate request from the WebAgent:

1. In the navigation pane, click **Security**.

2. Click **SSL**.

3. In the Web Management box, enable SSL if it is not already checked.

4. In the SSL Certificate box, fill out the fields and select **Create request**.

Figure 79: *CA certificate generation via SSL WebAgent screen*

To access the online help for SSL certificate generation, click "?" in the upper right corner of the screen.

Enabling SSL on the switch and anticipating SSL browser contact behavior

The `web-management ssl` command enables SSL on the switch and modifies parameters the switch uses for transactions with clients. After you enable SSL, the switch can authenticate itself to SSL enabled browsers. If you want to disable SSL on the switch, use the `no web-management ssl` command.

> **NOTE**
>
> Before enabling SSL on the switch you must generate the switch's host certificate and key. If you have not already done so, see **Generating the switch's server host certificate** on page 392.
>
> When configured for SSL, the switch uses its host certificate to authenticate itself to SSL clients, however unless you disable the standard HPE WebAgent with the `no web-management` command it will be still available for unsecured transactions.

SSL client contact behavior

At the first contact between the switch and an SSL client, if you have not copied the switch's host certificate into the browser's certificate folder, your browser's first connection to the switch will question the connection and, for security reasons, give you the option of accepting or refusing. If a CA-signed certificate is used on the switch, for which a root certificate exists on the client browser side, then the browser will NOT prompt the user to ensure the validity of the certificate. The browser will be able to verify the certificate chain of the switch server certificate up to the root certificate installed in the browser, thus authenticating the switch unequivocally. As long as you are confident that an unauthorized device is not using the switch's IP address in an attempt to gain access to your data or network, you can accept the connection.

NOTE

When an SSL client connects to the switch for the first time, it is possible for a "man-in-the-middle" attack; that is, for an unauthorized device to pose undetected as the switch, and learn the usernames and passwords controlling access to the switch. When using self-signed certificates with the switch, there is a possibility for a "man-in-the-middle" attack when connecting for the first time; that is, an unauthorized device could pose undetected as a switch, and learn the usernames and passwords controlling access to the switch. Use caution when connecting for the first time to a switch using self-signed certificates. Before accepting the certificate, closely verify the contents of the certificate (see browser documentation for additional information on viewing contents of certificate.) The security concern described above does not exist when using CA-signed certificates that have been generated by certificate authorities that the web browser already trusts.

Using the CLI interface to enable SSL

Syntax:

```
[no] web-management ssl
```

Enables or disables SSL on the switch

```
[port <1-65535|default:443>]
```

The TCP port number for SSL connections (default: 443).

```
show config
```

Shows status of the SSL server. When enabled `web-management ssl` will be present in the config list.

Procedure

1. To enable SSL on the switch:

2. Generate a Host certificate if you have not already done so. (See **Generating the switch's server host certificate** on page 392.)

3. Execute the `web-management ssl` command.

To disable SSL on the switch, do either of the following:

Execute `[no] web-management ssl`.

Zeroize the switch's host certificate or certificate key.

Using the WebAgent to enable SSL

CAUTION

SSL does not protect the switch from unauthorized access via the Telnet, SNMP, or the serial port. While Telnet access can be restricted by the use of passwords local to the switch, if you are unsure of the security this provides, you may want to disable Telnet access (`no Telnet`). If you need to increase SNMP security, use SNMP version 3 only for SNMP access. Another security measure is to use the authorized IP managers feature described in the switch's *Access Security Guide*. To protect against unauthorized access to the serial port (and the **Clear** button, which removes local password protection), keep physical access to the switch restricted to authorized personnel.

Procedure

- To enable SSL on the switch:

  1. In the navigation pane, click **Security**.

  2. Click **SSL**.

  3. Click **Change**.

  4. Select SSL **Enable** to enable SSL.

  5. Enter the TCP port you desire to connect on. It is recommended you use the default IP port number of 443.

     > **NOTE**
     >
     > Hewlett Packard Enterprise recommends using the default IP port number (443). However, you can use `web-management ssl tcp-port` to specify any TCP port for SSL connections except those reserved for other purposes. Examples of reserved IP ports are 23 (Telnet) and 80 (http). Some other reserved TCP ports on the switches are 49, 80, 1506, and 1513.

  6. Click **Save**.

- To disable SSL on the switch:

  1. In the navigation pane, click **Security**.

  2. Click **SSL**.

  3. Click **Change** button.

  4. Clear the SSL **Enable** box to disable SSL.

  5. Click **Save**.

Common errors in SSL setup

Error during	Possible cause
Generating host certificate on CLI	You have not generated a certificate key. See **CLI commands used to generate a server host certificate** on page 392.
Enabling SSL on the CLI or WebAgent	You have not generated a host certificate. See **Generate a self-signed host certificate with the WebAgent** on page 394.You may be using a reserved TCP port.
Unable to Connect with SSL	You may not have SSL enabled. See **Enabling SSL on the switch and anticipating SSL browser contact behavior** on page 395.Your browser may not support SSLv3 or TLSv1 or it may be disabled. See the documentation provided for your browser.

An Access Control List (ACL) is a list of one or more Access Control Entries (ACEs), where each ACE consists of a matching criteria and an action (permit or deny). The information below describes how to configure, apply, and edit static IPv4 ACLs in a network populated with switches, and how to monitor IPv4 ACL actions.

> **NOTE**
>
> This chapter describes ACLs for IPv4 configuration and operation. In this chapter, unless otherwise noted:
>
> • The term "ACL" refers to static IPv4 ACLs.
>
> • Descriptions of ACL operation apply only to static IPv4 ACLs.

IPv4 filtering with ACLs can help improve network performance and restrict network use by creating policies for:

Switch Management Access: Permits or denies in-band management access, including limiting and/or preventing the use of designated protocols that run on top of IPv4, such as TCP, UDP, IGMP, ICMP, and others. Also included are the use of precedence and ToS criteria, and control for application transactions based on source and destination IPv4 addresses and transport layer port numbers.

Application Access Security: Eliminates unwanted traffic in a path by filtering IPv4 packets where they enter or leave the switch on specific VLAN interfaces.

IPv4 ACLs can filter traffic to or from a host, a group of hosts, or entire subnets.

> **NOTE**
>
> IPv4 ACLs can enhance network security by blocking selected traffic, and can serve as part of your network security program. However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IPv4 packet transmissions, they should not be relied upon for a complete security solution.
>
> IPv4 ACLs on the switches covered by this manual do not filter non-IPv4 traffic such as IPv6, AppleTalk, and IPX packets.

Options for applying IPv4 ACLs on the switch

To apply IPv4 ACL filtering, assign a configured IPv4 ACL to the interface on which you want traffic filtering to occur. VLAN and routed IPv4 traffic ACLs can be applied statically using the switch configuration.

Port traffic ACLs can be applied either statically or dynamically (using a RADIUS server).

Static ACLs

Static ACLs are configured on the switch. To apply a static ACL, you must assign it to an interface (VLAN or port).

VLAN ACL (VACL)

A VACL is an ACL configured on a VLAN to filter traffic entering the switch on that VLAN interface and having a destination on the same VLAN.

Static port ACL

A static port ACL is an ACL configured on a port to filter traffic entering the switch on that port, regardless of whether the traffic is routed, switched, or addressed to a destination on the switch itself.

Dynamic port ACLs

A dynamic port ACL is configured on a RADIUS server for assignment to a given port when the server authenticates a specific client on that port. When the server authenticates a client associated with that ACL, the ACL is assigned to the port the client is using. The ACL then filters the IP traffic received inbound on that port from the authenticated client. When the client session ends, the ACL is removed from the port. The switch allows as many dynamic port ACLs on a port as it allows authenticated clients.

> **NOTE**
>
> The information provided here describes the IPv4 ACL applications you can statically configure on the switch. For information on dynamic port ACLs assigned by a RADIUS server, see .

Overview

Types of IPv4 ACLs

A permit or deny policy for IPv4 traffic you want to filter can be based on source address alone, or on source address plus other factors.

Standard ACL

Use a standard ACL when you need to permit or deny IPv4 traffic based on source address only. Standard ACLs are also useful when you need to quickly control a performance problem by limiting IPv4 traffic from a subnet, group of devices, or a single device. This can block all IPv4 traffic from the configured source, but does not hamper IPv4 traffic from other sources within the network.

A standard ACL uses an alphanumeric ID string or a numeric ID of 1 through 99. Specify a single host, a finite group of hosts, or any host.

Extended ACL

Use an extended ACL when simple IPv4 source address restrictions do not provide the sufficient traffic selection criteria needed on an interface. Extended ACLs allow use of the following criteria:

* source and destination IPv4 address combinations

* IPv4 protocol options

Extended, named ACLs also offer an option to permit or deny IPv4 connections using TCP for applications such as Telnet, http, ftp, and others.

ACL applications

ACL filtering is applied to IPv4 traffic as follows:

Static port ACL

Any inbound IPv4 traffic on that port.

Dynamic port ACL

On a port having an ACL assigned by a RADIUS server to filter an authenticated client's traffic, filters inbound IPv4 and IPv6 traffic from that client.

For information on RADIUS-assigned ACLs, see **RADIUS Services Support on HPE Switches** on page 274.

VLAN ACL (VACL)

On a VLAN configured with a VACL, inbound IP traffic, regardless of whether it is switched or routed. On a multinetted VLAN, this includes inbound IPv4 traffic from any subnet.

VACL applications

VACLs filter any IPv4 traffic entering the switch on a VLAN configured with the "VLAN" ACL option.

Syntax:

Within the vlan context.

```
ip access-group <identifier> vlan
```

For example, in **Figure 80: VACL filter application to IPv4 traffic entering the switch** on page 400, you would assign a VACL to VLAN 2 to filter all inbound switched or routed IPv4 traffic received from clients on the 10.28.20.0 network. In this instance, routed traffic received on VLAN 2 from VLANs 1 or 3 would not be filtered by the VACL on VLAN 2.

Figure 80: *VACL filter application to IPv4 traffic entering the switch*

> The switch allows one VACL assignment configured per VLAN. This is in addition to any other ACL applications assigned to the VLAN or to ports in the VLAN.

Static port ACL and RADIUS-assigned ACL applications

An IPv4 static port ACL filters any IPv4 traffic inbound on the designated port, regardless of whether the traffic is switched or routed.

RADIUS-assigned (dynamic) port ACL applications

Dynamic (RADIUS-assigned) port ACLs are configured on RADIUS servers and can be configured to filter IPv4 traffic inbound from clients authenticated by such servers.

802.1X user-based and port-based applications

User-Based 802.1X access control allows up to 32 individually authenticated clients on a given port. Port-Based access control does not set a client limit, and requires only one authenticated client to open a given port; it is recommended for applications where only one client at a time can connect to the port.

- If you configure 802.1X user-based security on a port and the RADIUS response includes a RADIUS-assigned ACL for at least one authenticated client, then the RADIUS response for all other clients authenticated on the port must also include a RADIUS-assigned ACL. Inbound IP traffic on the port from a client that authenticates without receiving a RADIUS-assigned ACL will be dropped and the client will be de-authenticated.

- Using 802.1X port-based security on a port where the RADIUS response includes a dynamic port ACL, only the first client to authenticate can use the port. Traffic from other clients will be dropped.

Multiple ACLs on an interface

The switch allows multiple ACL applications on an interface (subject to internal resource availability). This means that a port belonging to a given VLAN "X" can simultaneously be subject to all of the following:

- One static port ACL for any IPv4 or IPv6 traffic entering the switch on the port.

- One static VACL for IPv4 or IPv6 traffic for VLAN "X" entering the switch on the port.

- One dynamic (RADIUS-assigned) port ACL applied to inbound IPv4 traffic for each authenticated client on the port.

Features common to all ACL applications

- Any ACL can have multiple entries (ACEs).

- You can apply any one ACL to multiple interfaces.

- All ACEs in an ACL configured on the switch are automatically sequenced (numbered). For an existing ACL, entering an ACE without specifying a sequence number automatically places the ACE at the end of the list. Specifying a sequence number inserts the ACE into the list at the specified sequential location.

  - Automatic sequence numbering begins with "10" and increases in increments of 10. You can renumber the ACEs in an ACL and also change the sequence increment between ACEs.

  - The CLI `remark` command option allows you to enter a separate comment for each ACE.

- A source or destination IPv4 address and a mask, together, can define a single host, a range of hosts, or all hosts.

- Every ACL populated with one or more explicit ACEs includes an Implicit Deny as the last entry in the list. The switch applies this action to any packets that do not match other criteria in the ACL. For standard ACLs, the Implicit Deny is `deny any`. For extended ACLs, it is `deny ip any any`.

- In any ACL, you can apply an ACL log function to ACEs that have an explicit "deny" action. The logging occurs when there is a match on a "deny" ACE. The switch sends ACL logging output to Syslog, if configured, and, optionally, to a console session.

You can create ACLs for the switch configuration using either the CLI or a text editor. The text-editor method is recommended when you plan to create or modify an ACL that has more entries than you can easily enter or edit using the CLI alone.

General steps for planning and configuring ACLs

Procedure

1. Identify the ACL action to apply. As part of this step, determine the best points at which to apply specific ACL controls. For example, you can improve network performance by filtering unwanted IPv4 traffic at the edge of the network instead of in the core. Also, on the switch itself, you can improve performance by filtering unwanted IPv4 traffic where it is inbound to the switch instead of outbound.

> **CAUTION**
>
> **Regarding the use of source routing:** Source routing is enabled by default on the switch and can be used to override ACLs. For this reason, if you are using ACLs to enhance network security, the recommended action is to use the `no ip source-route` command to disable source routing on the switch. (If source routing is disabled in the running-config file, the `show running` command includes "`no ip source-route`" in the running-config file listing.)

Traffic source	ACL application
IPv4 or IPv6 traffic from a specific, authenticated client	dynamic port ACL (RADIUS-assigned ACL) for inbound IP traffic from an authenticated client on a port
IPv4 traffic entering the switch on a specific port	static port ACL (static-port assigned) for any inbound IPv4 traffic on a port from any source
switched or routed IPv4 traffic entering the switch on a specific VLAN	VACL (VLAN ACL)

2. Identify the traffic types to filter.

   - The SA and/or the DA of traffic you want to permit or deny. This can be a single host, a group of hosts, a subnet, or all hosts.

   - Traffic of a specific IPv4 protocol type (0-255)

   - Any TCP traffic (only) for a specific TCP port or range of ports, including optional control of connection traffic based on whether the initial request should be allowed

   - All UDP traffic or UDP traffic for a specific UDP port

   - All ICMP traffic or ICMP traffic of a specific type and code

   - All IGMP traffic or IGMP traffic of a specific type

   - Any of the above with specific precedence and/or ToS settings

3. Design the ACLs for the control points (interfaces) selected. When using explicit "deny" ACEs, optionally use the VACL logging feature for notification that the switch is denying unwanted packets.

4. Configure the ACLs on the selected switches.

5. Assign the ACLs to the interfaces you want to filter, using the ACL application (static port ACL) appropriate for each assignment.

6. Test for desired results.

IPv4 static ACL operation

Introduction

An ACL is a list of one or more Access Control Entries (ACEs), where each ACE consists of a matching criteria and an action (permit or deny). A static ACL applies only to the switch in which it is configured. ACLs operate on assigned interfaces, and offer these traffic filtering options:

IPv4 traffic inbound on a port.

The following table lists the range of interface options:

Interface	ACL Application	Application Point	Filter Action
Port	Static Port ACL (switch configured)Dynamic Port ACL	inbound on the switchinbound on the switch port used by authenticated client	inbound IPv4 trafficinbound IPv4 traffic from the authenticated client
VLAN	VACL	entering the switch on the VLAN	inbound IPv4 traffic

NOTE
After you assign an IPv4 ACL to an interface, the default action on the interface is to implicitly deny IPv4 traffic that is not specifically permitted by the ACL. (This applies only in the direction of traffic flow filtered by the ACL.)

The packet-filtering process

Sequential comparison and action

When an ACL filters a packet, it sequentially compares each ACE's filtering criteria to the corresponding data in the packet until it finds a match. The action indicated by the matching ACE (deny or permit) is then performed on the packet.

Implicit Deny

If a packet does not have a match with the criteria in any of the ACEs in the ACL, the ACL denies (drops) the packet. If you need to override the implicit deny so that a packet that does not have a match will be permitted, then you can use the "permit any" option as the last ACE in the ACL. This directs the ACL to permit (forward) packets that do not have a match with any earlier ACE listed in the ACL, and prevents these packets from being filtered by the implicit "deny any".

The ACL tries to apply the first ACE in the list. If there is not a match, it tries the second ACE, and so on. When a match is found, the ACL invokes the configured action for that entry (permit or drop the packet) and no further comparisons of the packet are made with the remaining ACEs in the list. This means that when an ACE whose criteria matches a packet is found, the action configured for that ACE is invoked, and any remaining ACEs in the

ACL are ignored. **Because of this sequential processing, successfully implementing an ACL depends in part on configuring ACEs in the correct order for the overall policy you want the ACL to enforce.**

Figure 81: *The packet-filtering process in an ACL with N entries (ACEs)*

The order in which an ACE occurs in an ACL is significant.

For example, if an ACL contains six ACEs, but the first ACE allows Permit Any forwarding, then the ACL permits all IPv4 traffic, and the remainingACEs in the list do not apply, even if they specify criteria that would make a match with any of the traffic permitted by the first ACE.

It is important to remember that all IPv4 ACLs configurable on the switch include an implicit `deny ip any`. That is, IPv4 packets that the ACL does not **explicitly** permit or deny will be **implicitly** denied, and therefore dropped instead of forwarded on the interface. If you want topreempt the implicit deny so that IPv4 packets not explicitly denied by other ACEs in the ACL will be permitted, insert an explicit "permit any" as the last ACE in the ACL. Doing so permits any packet not explicitly denied by earlier entries.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Planning an ACL application

Before creating and implementing ACLs, you need to define the policies you want your ACLs to enforce, and understand how the ACL assignments will impact your network users.

> **NOTE**
>
> All IPv4 traffic entering the switch on a given interface is filtered by all ACLs configured for inbound traffic on that interface. For this reason, an inbound IPv4 packet will be denied (dropped) if it has a match with either an implicit or explicit `deny` in any of the inbound ACLs applied to the interface.
>
> See **Multiple ACLs on an interface** on page 401 for more detail.

IPv4 traffic management and improved network performance

Use ACLs to block traffic from individual hosts, workgroups, or subnets, and to block access to VLANs, subnets, devices, and services.

Traffic criteria for ACLs include:

- Switched and/or routed traffic

- Any traffic of a specific IPv4 protocol type (0-255)

- Any TCP traffic (only) for a specific TCP port or range of ports, including optional control of connection traffic based on whether the initial request should be allowed

- Any UDP traffic or UDP traffic for a specific UDP port

- Any ICMP traffic or ICMP traffic of a specific type and code

- Any IGMP traffic or IGMP traffic of a specific type

- Any of the above with specific precedence and/or ToS settings

Answering the following questions can help you to design and properly position IPv4 ACLs for optimum network usage.

- What are the logical points for minimizing unwanted traffic, and what ACL application(s) should be used? In many cases it makes sense to prevent unwanted traffic from reaching the core of your network by configuring ACLs to drop the unwanted traffic at or close to the edge of the network. The earlier in the network path you can block unwanted traffic, the greater the benefit for network performance.

- From where is the traffic coming? The source and destination of traffic you want to filter determines the ACL application to use (static port ACL, VACL, and dynamic port ACL).

- What traffic should you explicitly block? Depending on your network size and the access requirements of individual hosts, this can involve creating a large number of ACEs in a given ACL (or a large number of ACLs), which increases the complexity of your solution.

- What traffic can you implicitly block by taking advantage of the implicit `deny ip any` to deny traffic that you have not explicitly permitted? This can reduce the number of entries needed in an ACL.

- What traffic should you permit? In some cases you will need to explicitly identify permitted traffic. In other cases, depending on your policies, you can insert an ACE with "permit any" forwarding at the end of an ACL. This means that all IPv4 traffic not specifically matched by earlier entries in the list will be permitted.

Security

ACLs can enhance security by blocking traffic carrying an unauthorized source IPv4 address (SA). This can include:

- Blocking access from specific devices or interfaces (port or VLAN)

- Blocking access to or from subnets in your network

- Blocking access to or from the internet

- Blocking access to sensitive data storage or restricted equipment

- Preventing specific IPv4, TCP, UDP, IGMP, and ICMP traffic types, including unauthorized access using functions such as Telnet, SSH, and web browser

You can also enhance switch management security by using ACLs to block IPv4 traffic that has the switch itself as the destination address (DA).

> **CAUTION**
>
> IPv4 ACLs can enhance network security by blocking selected traffic, and can serve as one aspect of maintaining network security. **However, because ACLs do not provide user or device authentication, or protection from malicious manipulation of data carried in IP packet transmissions, they should not be relied upon for a complete security solution.**

> **NOTE**
>
> Static ACLs for the switches covered by this guide do not filter non-IPv4 traffic such as IPv6, AppleTalk, and IPX. Dynamic port ACLs assigned by a RADIUS server can be configured on the server to filter IPv4 traffic, but do not filter non-IP traffic.

Guidelines for planning the structure of a static ACL

After determining the filtering type (standard or extended) and ACL application (VACL or static port ACL) to use at a particular point in your network, determine the order in which to apply individual ACEs to filter IPv4 traffic. For information on ACL applications, see **ACL applications** on page 399.

- The sequence of ACEs is significant. When the switch uses an ACL to determine whether to permit or deny a packet on a particular VLAN, it compares the packet to the criteria specified in the individual Access Control Entries (ACEs) in the ACL, beginning with the first ACE in the list and proceeding sequentially until a match is found. When a match is found, the switch applies the indicated action (permit or deny) to the packet.

- The first match in an ACL dictates the action on a packet. Subsequent matches in the same ACL areignored.However, if a packet is permitted by one ACL assigned to an interface, but denied by another ACL assigned to the same interface, the packet will be denied on the interface.

- On any ACL, the switch implicitly denies IPv4 packets that are not explicitly permitted or denied by the ACEs configured in the ACL. If you want the switch to forward a packet for which there is not a match in an ACL, append an ACE that enables Permit Any forwarding as the last ACE in the ACL. This ensures that no packets reach the Implicit Deny case for that ACL.

- Generally, you should list ACEs from the most specific (individual hosts) to the most general (subnets or groups of subnets) unless doing so permits traffic that you want dropped. For example, an ACE allowing a small group of workstations to use a specialized printer should occur earlier in an ACL than an entry used to block widespread access to the same printer.

IPv4 ACL configuration and operating rules

Static port ACLs

A static port ACL filters traffic entering the switch on the ports or trunks to which it is assigned.

Per switch ACL limits for all ACL types

At a minimum an ACL must have one, explicit "permit" or "deny" Access Control Entry. You can configure up to 512 ACLs.

For IPv4 ACLs, the maximums are as follows:

- Named (Extended or Standard) ACLs: Up to 512 (minus any numeric standard or extended ACL assignments)

- Numeric Standard ACLs: Up to 99; numeric range: 1-99

- Numeric Extended ACLs: Up to 100; numeric range: 100-199

- The maximum number of ACEs supported by the switch is up to 1024 for IPv4 ACEs. The maximum number of ACEs allowed on port depends on the concurrent resource usage by multiple configured features. For more information, use the `show <qos|access-list> resources` command and/or refer to **General ACL operating notes** on page 459.

Implicit deny

In any static IPv4 ACL, the switch automatically applies an implicit `deny ip any` that does not appear in show listings. This means that the ACL denies any IPv4 packet it encounters that does not have a match with an entry in the ACL. Thus, if you want an ACL to permit any packets that you have not expressly denied, you must enter a `permit any` or `permit ip any any` as the last ACE in an ACL. Because, for a given packet the switch sequentially applies the ACEs in an ACL until it finds a match, any packet that reaches the `permit any` or `permit ip any any` entry will be permitted, and will not encounter the `deny ip any` ACE the switch automatically includes at the end of the ACL.

For Implicit Deny operation in dynamic ACLs, see **RADIUS Services Support on HPE Switches** on page 274.

Explicitly permitting any IPv4 traffic

Entering a `permit any` or a permit `ip any any` ACE in an ACL permits all IPv4 traffic not previously permitted or denied by that ACL. Any ACEs listed after that point do not have any effect.

Explicitly denying any IPv4 traffic

Entering a `deny any` or a `deny ip any any` ACE in an ACL denies all IPv4 traffic not previously permitted or denied by that ACL. Any ACEs after that point have no effect.

Replacing one ACL with another using the same application

For a specific interface, the most recent ACL assignment using a given application replaces any previous ACL assignment using the same application on the same interface.

Static port ACLs:

These are applied per-port, per port-list, or per static trunk. Adding a port to a trunk applies the trunk's ACL configuration to the new member. If a port is configured with an ACL, the ACL must be removed before the port is added to the trunk. Also, removing a port from an ACL-configured trunk removes the ACL configuration from that port.

VACLs

These filter any IPv4 traffic entering the switch through any port belonging to the designated VLAN. VACLs do not filter traffic leaving the switch or being routed from another VLAN.

A VACL affects all physical ports in a static VLAN

A VACL assigned to a VLAN applies to all physical ports on the switch belonging to that VLAN, including ports that have dynamically joined the VLAN.

How an ACE uses a mask to screen packets for matches

When the switch applies an ACL to IPv4 traffic, each ACE in the ACL uses an IPv4 address and ACL mask to enforce a selection policy on the packets being screened. That is, the mask determines the range of IPv4 addresses (SA only or SA/DA) that constitute a match between the policy and a packet being screened.

What Is the difference between network (or subnet) masks and the masks used with ACLs?

In common IPv4 addressing, a network (or subnet) mask defines which part of the address to use for the network number and which part to use for the hosts on the network. For example:

Address	Mask	Network address	Host address
10.38.252.195	255.255.255.0	first three octets	The fourth octet.
10.38.252.195	255.255.248.0	first two octets and the left-most five bits of the third octet	The right most three bits of the third octet and all bits in the fourth octet.

Thus, the bits set to 1 in a network mask define the part of an IPv4 address to use for the network number, and the bits set to 0 in the mask define the part of the address to use for the host number.

In an ACL, IPv4 addresses and masks provide criteria for determining whether to deny or permit a packet, or to pass it to the next ACE in the list. If there is a match, the configured deny or permit action occurs. If there is not a match, the packet is compared with the next ACE in the ACL. Thus, where a standard network mask defines how to identify the network and host numbers in an IPv4 address, the mask used with ACEs defines which bits in a packet's SA or DA must match the corresponding bits in the SA or DA listed in an ACE, and which bits can be **wildcards**.

Rules for defining a match between a packet and an ACE

- For a given ACE, when the switch compares an IPv4 address and corresponding mask in the ACE to an IPv4 address carried in a packet:

  - **A mask-bit setting of 0 ("off")** requires that the corresponding bits in the packet's address and in the ACE's address must be the same. Thus, if a bit in the ACE's address is set to 1 ("on"), the same bit in the packet's address must also be 1.

  - **A mask-bit setting of 1 ("on")** means the corresponding bits in the packet's address and in the ACE's address do not have to be the same. Thus, if a bit in the ACE's address is set to 1, the same bit in the packet's address can be either 1 or 0 ("on" or "off").

- In any ACE, a mask of all ones means any IPv4 address is a match. Conversely, a mask of all zeros means the only match is an IPv4 address identical to the host address specified in the ACE.

- Depending on your network, a single ACE that allows a match with more than one source or destination IPv4 address may allow a match with multiple subnets. For example, in a network with a prefix of 31.30.240 and a subnet mask of 255.255.240.0 (the leftmost 20 bits), applying an ACL mask of 0.0.31.255 causes the subnet mask and the ACL mask to overlap one bit, which allows matches with hosts in two subnets: 31.30.224.0 and 31.30.240.0.

Bit Position in the Third Octet of Subnet Mask 255.255.240.0								
Bit Values	128	64	32	16	8	4	2	1
Subnet Mask Bits	1	1	1	1	n/a	n/a	n/a	n/a
Mask Bit Settings Affecting Subnet Addresses	0	0	0	1 or 0	n/a	n/a	n/a	n/a

This ACL supernetting technique can help to reduce the number of ACLs you need. You can apply it to a multinetted VLAN and to multiple VLANs. However, ensure that you exclude subnets that do not belong in the policy. If this creates a problem for your network, you can eliminate the unwanted match by making the ACEs in your ACL as specific as possible, and using multiple ACEs carefully ordered to eliminate unwanted matches.

- Every IPv4 address and mask pair (source or destination) used in an ACE creates one of the following policies:

  ◦ **Any IPv4 address fits the matching criteria.**

    In this case, the switch automatically enters the address and mask in the ACE. For example:

    Syntax: `access-list 1 deny any`

    Produces this policy in an ACL listing:

Address	Mask
0.0.0.0	255.255.255.255

    This policy states that every bit in every octet of a packet's SA is a wildcard, which covers any IPv4 address.

  ◦ **One IPv4 address fits the matching criteria.**

    In this case, you provide the address and the switch provides the mask. For example:

    Syntax: `access-list 1 permit host 10.28.100.15`Produces this policy in an ACL listing:

Address	Mask
10.28.100.15	0.0.0.0

    This policy states that every bit in every octet of a packet's SA must be the same as the corresponding bit in the SA defined in the ACE.

  ◦ **A group of IPv4 addresses fits the matching criteria.**

    In this case you provide both the address and the mask. For example:

    Syntax: `access-list 1 permit 10.28.32.1 0.0.0.31`

Address	Mask
10.28.32.1	0.0.0.31

This policy states that:

- In the first three octets of a packet's SA, every bit must be set the same as the corresponding bit in the SA defined in the ACE.

- In the last octet of a packet's SA, the first three bits must be the same as in the ACE, but the last five bits are wildcards and can be any value.

- Unlike subnet masks, the wildcardbits in an ACL mask need not be contiguous. For example, 0.0.7.31 is a valid ACL mask. However, a subnet mask of 255.255.248.224 is not a valid subnet mask.

Example of how the mask bit settings define a match

Assume an ACE where the second octet of the mask for an SA is 7 (the rightmost three bits are "on", or "1") and the second octet of the corresponding SA in the ACE is 31 (the rightmost five bits). In this case, a match occurs when the second octet of the SA in a packet being filtered has a value in the range of 24 to 31.

Table 29: *How the mask defines a match*

Location of octet	Bit position in the octet						
	128	64	32	16	8	4	2
SA in ACE	0	0	0	1	1	1	1
Mask for SA	0	0	0	0	0	1	1
Corresponding Octet of a Packet's SA	0[*]	0[*]	0[*]	1[*]	1[*]	0/1	0/1

[*] Indicates that the bits in the packet must exactly match the bits in the source address in the ACE. Wherever the mask bits are zeros, the bit in the packet must exactly match must exactly match the corresponding bit in the source address in the ACE. Wherever the mask bits are ones (wildcards), the corresponding address bits in the packet can be any value.

NOTE

This example covers only one octet of an IPv4 address. An actual ACE applies this method to all four octets of the address.

Example of allowing only one IPv4 address ("host" option)

Suppose, for example, that you have configured the ACL in **Figure 82: An ACL with an ACE that allows only one source address** on page 411 to filter inbound packets on VLAN 20. Because the mask is all zeros, the ACE

policy dictates that a match occurs only when the source address on such packets is identical to the address configured in the ACE.

Figure 82: *An ACL with an ACE that allows only one source address*

Examples allowing multiple IPv4 addresses

The following tables provide examples of how to apply masks to meet various filtering requirements.

Table 30: *Using an IP address and mask in an ACE*

Address in the ACE	Mask	Policy for a match between a packet and the ACE	Allowed addresses
A: 10.38.252.195	0.0.0.255	Exact match in first three octets only.	10.38.252.<0-255> (See row A in the following table.)
B: 10.38.252.195	0.0.7.255	Exact match in the first two octets and the leftmost five bits (248) of the third octet.	10.38.<248-255>.<0-255> (In the third octet, only the rightmost three bits are wildcard bits. The leftmost five bits must be a match, and in the ACE, these bits are all set to 1. See row B in the following table.)

Table Continued

Chapter 20 IPv4 Access Control Lists (ACLs)

Address in the ACE	Mask	Policy for a match between a packet and the ACE	Allowed addresses
C: 10.38.252.195	0.0.0.0	Exact match in all octets.	10.38.252.195 (There are no wildcard bits in any of the octets. See row C in the following table.)
D: 10.38.252.195	0.15.255.255	Exact match in the first octet and the leftmost four bits of the second octet.	10.<32-47> .<0-255> .<0-255> (In the second octet, the rightmost four bits are wildcard bits. See row D in the following table.)

Table 31: *Mask effect on selected octets of the IPv4 addresses in Using an IP address and mask in an ACE*

Addr	Octet	Mask	Octet range	128	64	32	16	8	4	2	1
A	3	0 all bits	252	1	1	1	1	1	1	0	0
B	3	7 last 3 bits	248-255	1	1	1	1	1	0 or 1	0 or 1	0 or 1
C	4	0 all bits	195	1	1	0	0	0	0	1	1
D	2	15 last 4 bits	32-47	0	0	1	0	0 or 1	0 or 1	0 or 1	0 or 1
All bit settings in this table must be an exact match.											

If there is a match between the policy in the ACE and the IPv4 address in a packet, the packet is either permitted or denied according to how the ACE is configured. If there is no match, the next ACE in the ACL is applied to the packet. The same operation applies to a destination IPv4 address used in an extended ACE.

Where an ACE includes both source and destination addresses, there is one address/ACL-mask pair for the source address, and another address/ACL-mask pair for the destination address.

Configuring and assigning an IPv4 ACL

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

General steps for implementing ACLs

Procedure

1. Configure one or more ACLs.

2. This creates and stores the ACL(s) in the switch configuration.

3. Assign an ACL.

4. This step uses one of the following applications to assign the ACL to an interface:

   a. VACL (any IPv4 traffic entering the switch on a given VLAN)

   b. Static Port ACL (any IPv4 traffic entering the switch on a given port, port list, or static trunk)

Options for permit/deny policies

The permit or deny policy for IPv4 traffic you want to filter can be based on source address alone, or on source address plus other IPv4 factors.

- **Standard ACL:** Uses only a packet's source IPv4 address as a criterion for permitting or denying the packet. For a standard ACL ID, use either a unique numeric string in the range of 1-99 or a unique name string of up to 64 alphanumeric characters.

- **Extended ACL:** Offers the following criteria as options for permitting or denying a packet:

  ◦ source IPv4 address

  ◦ destination IPv4 address

  ◦ IPv4 protocol options:

    – Any IPv4 traffic

    – Any traffic of a specific IPv4 protocol type (0-255)

    – Any TCP traffic (only) for a specific TCP port or range of ports, including optional use of TCP control bits or control of connection (established) traffic based on whether the initial request should be allowed

    – Any UDP traffic (only) or UDP traffic for a specific UDP port

    – Any ICMP traffic (only) or ICMP traffic of a specific type and code

    – Any IGMP traffic (only) or IGMP traffic of a specific type

  For an extended ACL ID, use either a unique number in the range of 100-199 or a unique name string of up to 64 alphanumeric characters.

Carefully plan ACL applications before configuring specific ACLs.

ACL configuration structure

After you enter an ACL command, you may want to inspect the resulting configuration. This is especially true where you are entering multiple ACEs into an ACL. Also, it is helpful to understand the configuration structure when using the following information.

The basic ACL structure includes four elements:

Procedure

1. ACL identity and type: identifies the ACL as `standard` or `extended` and shows the ACL name or number.

2. Optional `remark` entries.

3. One or more deny/permit list entries (ACEs): One entry per line.

Element	Notes
Type	Standard or Extended
Identifier	• Alphanumeric; Up to 64 Characters, Including Spaces • Numeric: 1-99 (Standard) or 100-199 (Extended)
Remark	Allows up to 100 alphanumeric characters, including blank spaces. (If any spaces are used, the remark must be enclosed in a pair of single or double quotes.) A remark is associated with a particular ACE and will have the same sequence number as the ACE. (One remark is allowed per ACE.)
Maximum ACEs per Switch	The upper limit on ACEs supported by the switch depends on the concurrent resource usage by configured ACL, QoS, Mirroring, and other features.

4. Implicit Deny: Where an ACL is in use, it denies any packets that do not have a match with the ACEs explicitly configured in the list. The Implicit Deny does not appear in ACL configuration listings, but always functions when the switch uses an ACL to filter packets. (You cannot delete the Implicit Deny, but you can supersede it with a `permit any` or `permit ip any any` statement.)

Standard ACL structure

Individual ACEs in a standard ACL include only a permit/deny statement, the source addressing, and an optional `log` command (available with "deny" statements).

Figure 83: *The general structure for a standard ACL*


```
ip access-list standard < identifier >"
  [[ seq-# ] remark < remark-str >]
  < permit | deny > < SA > [log]
  .
  .
  .
  < Implicit Deny >
exit
```

Note: The optional **log** function is available only for explicit "deny" ACEs.

For example, **Figure 84: A displayed standard ACL configuration with two ACEs** on page 415 shows how to interpret the entries in a standard ACL.

Figure 84: *A displayed standard ACL configuration with two ACEs*

Extended ACL configuration structure

Individual ACEs in an extended ACL include:

- A permit/deny statement

- Source and destination IPv4 addressing

- Choice of IPv4 criteria, including optional precedence and ToS

Optional ACL `log` command (for `deny` entries)

Optional remark statements

Figure 85: *Example of general structure options for an extended ACL*

```
ip access-list extended < identifier >
[[ seq-# ] remark < remark-str >]
< permit | deny > < ipv4-protocol-type > < SA > < src-acl-mask > < DA > <dest-acl-mask > [log]
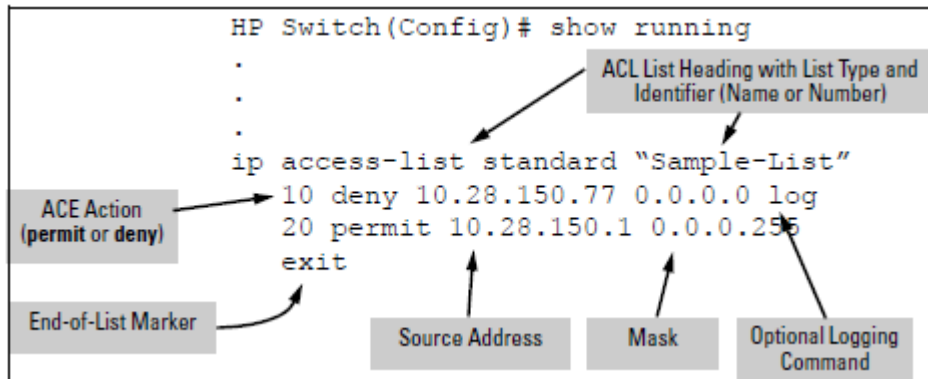
        < permit | deny > tcp
            < SA > < src-acl-mask > [< operator > < port-id >]
            < DA > < desti-acl-mask > [< operator > < port-id >]  [log]
            [ established ]

        < permit | deny > udp
            < SA > < src-acl-mask > [< operator > < port-id >]
            < DA > < dest-acl-mask > [< operator > < port-id >]  [log]

        < permit | deny > icmp
            < SA > < src-acl-mask > < DA > < dest-acl-mask > [ icmp-type ]  [log]

        < permit | deny > igmp
            < SA > < SA-mask > < DA > < dest-acl-mask > [ igmp-type ]  [log]

        [ precedence < priority >]
        [ tos < tos-setting >]
        . . .
        < Implicit Deny >
        exit
```

Note: The optional **log** function is available only for "deny" ACEs.

For example, **Figure 86: Displayed extended ACL configuration** on page 417 shows how to interpret the entries in an extended ACL.

Figure 86: *Displayed extended ACL configuration*

ACL configuration factors

The sequence of entries in an ACL is significant

When the switch uses an ACL to determine whether to permit or deny a packet, it compares the packet to the criteria specified in the individual ACEs in the ACL, beginning with the first ACE in the list and proceeding sequentially until a match is found. When a match is found, the switch applies the indicated action (permit or deny) to the packet. This is significant because, once a match is found for a packet, subsequent ACEs in the same ACL will not be applied to that packet, regardless of whether they match the packet.

For example, suppose that you have applied the ACL shown in to inbound IPv4 traffic on VLAN 1 (the default VLAN):

Figure 87: *A standard ACL that permits all IPv4 traffic not implicitly denied*

Table 32: *Effect of the above ACL on inbound IPv4 traffic in the assigned VLAN*

Line #	Action
n/a	Shows type (extended) and ID (Sample-List-2).
10	A packet from SA 10.28.235.10 will be denied (dropped). This ACE filters out all packets received from 10.28.235.10. As a result, IPv4 traffic from that device will not be allowed and packets from that device will not be compared against any later entries in the list.
20	A packet from SA 10.28.245.89 will be denied (dropped). This ACE filters out all packets received from 10.28.245.89. As the result, IPv4 traffic from that device will not be allowed and packets from that device will not be compared against any later entries in the list.
30	A TCP packet from SA 10.28.18.100 with a DA of 10.28.237.1 will be permitted (forwarded). Since no earlier ACEs in the list have filtered TCP packets from 10.28.18.100 and destined for 10.28.237.1, the switch will use this ACE to evaluate such packets. Any packets that meet this criteria will be forwarded. (Any packets that do not meet this TCP source-destination criteria are not affected by this ACE.)
40	A TCP packet from source address 10.28.18.100 to any destination address will be denied (dropped). Since, in this example, the intent is to block TCP traffic from 10.28.18.100 to any destination except the destination stated in the ACE at line 30, this ACE must follow the ACE at line 30. (If their relative positions were exchanged, all TCP traffic from 10.28.18.100 would be dropped, including the traffic for the 10.28.18.1 destination.)
50	Any packet from any IPv4 SA to any IPv4 DA will be permitted (forwarded). The only traffic to reach this ACE will be IPv4 packets not specifically permitted or denied by the earlier ACEs.

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Line #	Action
n/a	The Implicit Deny is a function the switch automatically adds as the last action in all ACLs. It denies (drops) any IPv4 traffic from any source to any destination that has not found a match with earlier entries in the ACL. In this example, the ACE at line 50 permits (forwards) any IPv4 traffic not already permitted or denied by the earlier entries in the list, so there is no traffic remaining for action by the Implicit Deny function.
exit	Marks the end of the ACL.

Allowing for the Implied Deny function

In any ACL having one or more ACEs there will always be a packet match. This is because the switch automatically applies an Implicit Deny as the last ACE in any ACL. This function is not visible in ACL listings, but is always present, see **Figure 87: A standard ACL that permits all IPv4 traffic not implicitly denied** on page 418. This means that if you configure the switch to use an ACL for filtering either inbound or outbound IPv4 traffic on a VLAN, any packets not specifically permitted or denied by the explicit entries you create will be denied by the Implicit Deny action. If you want to preempt the Implicit Deny (so that IPv4 traffic not specifically addressed by earlier ACEs in a given ACL will be permitted), insert an explicit `permit any` (for standard ACLs) or `permit ip any any` (for extended ACLs) as the last explicit ACE in the ACL.

A configured ACL has no effect until you apply it to an interface

The switch stores ACLs in the configuration file. Thus, until you actually assign an ACL to an interface, it is present in the configuration, but not used (and does not use any of the monitored resources, see "Monitored Resources" in the management and configuration guide for your switch.)

You can assign an ACL name or number to an interface even if the ACL does not exist in the switch configuration

In this case, if you subsequently create an ACL with that name or number, the switch automatically applies each ACE as soon as you enter it in the running-config file. Similarly, if you modify an existing ACE in an ACL you already applied to an interface, the switch automatically implements the new ACE as soon as you enter it. (See **General ACL operating notes** on page 459.) The switch allows a maximum of 512 ACLs (IPv4), and determines the total from the number of unique ACL names in the configuration.

Using the CLI to create an ACL

You can use either the switch CLI or an offline text editor to create an ACL. This section describes the CLI method, which is recommended for creating short ACLs.

Inserting or adding an ACE to an ACL

These rules apply to all IPv4 ACEs you create or edit using the CLI:

Named IPv4 ACLs:

Add an ACE to the end of a named ACE by using the `ip access-list` command to enter the Named ACL (`nacl`) context and entering the ACE without the sequence number.For example, if you wanted to add a "permit" ACL at the end of a list named "List-1" to allow traffic from the device at 10.10.10.100:

```
HP Switch(config)# ip access-list standard List-1
HP Switch(config-std-nacl)# permit host 10.10.10.100
```

Insert an ACE anywhere in a named ACL by specifying a sequence number. For example, if you wanted to insert a new ACE as line 15 between lines 10 and 20 in an existing ACL named "List-2" to deny IPv4 traffic from the device at 10.10.10.77:

```
HP Switch(config)# ip access-list standard List-2

HP Switch(config-std-nacl)# 15 deny host 10.10.10.77
```

Numbered IPv4 ACLs:

Add an ACE to the end of a numbered ACL by using the `access-list <1-99|100-199>` command. For example, if you wanted to add a "permit" ACE at the end of a list identified with the number "11" to allow IPv4 traffic from the device at 10.10.10.100:

```
HP Switch(config)# access-list 11 permit host 10.10.10.100
```

To insert an ACE anywhere in a numbered ACL, use the same process as described above for inserting an ACE anywhere in a named ACL. For example, to insert an ACE denying IPv4 traffic from the host at 10.10.10.77 as line 52 in an existing ACL identified (named) with the number 11:

```
HP Switch(config)# ip access-list standard 99

HP Switch(config-std-nacl)# 52 deny host 10.10.10.77
```

> **NOTE**
>
> After a numbered ACL has been created (using `access-list <1-99|100-199>`), it can be managed as either a named or numbered ACL.

- Deleting an ACE: Enter the ACL context and delete the sequence number for the unwanted ACE. (To view the sequence numbers of the ACEs in a list, use `show access-list <acl-name-str> config` .)

- Duplicate ACEs are not allowed in the same ACL. Attempting to enter a duplicate ACE displays the `Duplicate access control entry` message.

Using CIDR notation to enter the IPv4 ACL mask

Use CIDR notation to enter ACL masks. The switch interprets the bits specified with CIDR notation as the address bits in an ACL and the corresponding address bits in a packet that must match. The switch then converts the mask to inverse notation for ACL use.

Table 33: *Examples of CIDR notation for masks*

Address used in an ACL with CIDR notation	Resulting ACL mask	Meaning
10.38.240.125/15	0.1.255.255	The leftmost 15 bits must match; the remaining bits are wildcards.
10.38.240.125/20	0.0.15.255	The leftmost 20 bits must match; the remaining bits are wildcards.
10.38.240.125/21	0.0.7.255	The leftmost 21 bits must match; the remaining bits are wildcards.

Table Continued

Address used in an ACL with CIDR notation	Resulting ACL mask	Meaning
10.38.240.125/24	0.0.0.255	The leftmost 24 bits must match; the remaining bits are wildcards.
18.38.240.125/32	0.0.0.0	All bits must match.

Configuring standard ACLs

A standard ACL uses only source IPv4 addresses in its ACEs. This type of ACE is useful when you need to:

- Permit or deny any IPv4 traffic based on source address only.

- Quickly control the IPv4 traffic from a specific address. This allows you to isolate IPv4 traffic problems generated by a specific device, group of devices, or a subnet threatening to degrade network performance. This gives you an opportunity to troubleshoot without sacrificing performance for users outside of the problem area.

A **named**, standard ACL is identified by an alphanumeric string of up to 64 characters and is created by entering the Named ACL (`nacl`) context. A numbered, standard ACL is identified by a number in the range of 1-99 and is created without having to leave the global config context. Note that the CLI Command syntax for creating a named ACL differs from the Command syntax for creating a numbered ACL. For example, the first pair of entries below illustrate how to create (or enter) a named, standard ACL and enter an ACE. The next entry illustrates creating a numbered, standard ACL with the same ACE.

```
HP Switch(config)# ip access-list standard Test-List
HP Switch(config-std-nacl)# permit host 10.10.10.147

HP Switch(config)# access-list 1 permit host 10.10.10.147
```

Note that once a numbered ACL has been created, it can be accessed using the named ACL method. This is useful if it becomes necessary to edit a numbered ACL by inserting or removing individual ACEs. Inserting or deleting an ACE is done by sequence number, and requires the Named ACL (`nacl`) context.

Configuring named, standard ACLs

This section describes the commands for performing the following:

- creating and/or entering the context of a named, standard ACL

- appending an ACE to the end of an existing list or entering the first ACE in a new list

Entering the IPv4 named ACL context

This command is a prerequisite to entering or editing ACEs in a named ACL.

Syntax:

```
ip access-list standard <name-str>
```

Places the CLI in the "Named ACL" (`nacl`) context specified by the <name-str> alphanumeric identifier. This enables entry of individual ACEs in the specified ACL. If the ACL does not already exist, this command creates it.

<name-str>: Specifies an identifier for the ACL. Consists of an alphanumeric string of up to 64 case-sensitive characters. Including spaces in the string requires that you enclose the string in single or double quotes.

Configuring ACEs in a named, standard ACL

Configuring ACEs is done after using the `ip access-list standard <name-str>` command to enter the "Named ACL" (`nacl`) context of an access list. See **Configuring ACEs in a named, standard ACL** on page 422.

Syntax:

`<deny|permit>`

`<any|host <SA>|<SA> <mask>|<SA>/<mask-length> > [log]`

Executing this command appends the ACE to the end of the list of ACEs in the current ACL. In the default ACL configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using `resequence` (See **Resequencing the ACEs in an ACL** on page 444.)

> **NOTE**
>
> To insert a new ACE between two existing ACEs, precede `deny` or `permit` with an appropriate sequence number. See **Inserting an ACE in an existing ACL** on page 442.

`<deny|permit>`

For named ACLs, used in the "Named ACL" (`nacl`) context to configure an ACE. Specifies whether the ACE denies or permits a packet matching the criteria in the ACE, as described below.

`<any|host <SA>|<SA> <mask>|<SA>/<mask-length> > [log]`

Defines the source IPv4 address (SA) a packet must carry for a match with the ACE.

- `any`: Allows IPv4 packets from any SA.

- `host <SA>`: Specifies only packets having **<SA>** as the source. Use this criterion when you want to match the IPv4 packets from a single source address.

- `<SA> <mask>` or `<SA>/<mask-length>`:

  Specifies packets received from either a subnet or a group of IPv4 addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits).

- Mask Application: The mask is applied to the IPv4 address in the ACE to define which bits in a packet's SA must exactly match the SA configured in the ACE and which bits need not match. For example: `10.10.10.1/24` and `10.10.10.1 0.0.0.255` both define any address in the range of 10.10.10.(1-255).

  > **NOTE**
  >
  > Specifying a group of contiguous addresses may require more than one ACE.

- `[log]`: This option generates an ACL log message if:

  ◦ There is a match.

  ◦ ACL logging is enabled on the switch.

  ◦ Use the debug command to direct ACL logging output to the current console session and/or to a Syslog server. Note that you must also use the `logging <ip-addr>` command to specify the addresses of Syslog servers to which you want log messages sent.

Example:

This example creates an ACL that:

- permits IPv4 traffic from a host with the address of 10.10.10.104

- creates another ACE that blocks all other IPv4 traffic from the same subnet

- allows all other IPv4 traffic

Figure 88: *Commands used to create a standard, named ACL*

```
HP Switch(config)# ip access-list standard Sample-List        Creates the "Sample-List"
HP Switch(config-std-nacl)# permit host 10.10.10.104          ACL and enters the "Named
HP Switch(config-std-nacl)# deny 10.10.10.1/24 log            ACL" context for this list.
HP Switch(config-std-nacl)# permit any                        Appends three ACEs to the
HP Switch(config-std-nacl)# exit                              list in the order shown.
HP Switch(config)# _                                          Exits from the nacl context.
```

Screen output listing the sample-list ACL content

```
HP Switch(config)# show access-list Sample-List
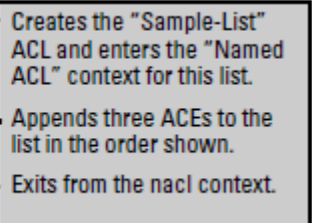
Access Control Lists

  Name: Sample-List
  Type: Standard
  Applied: No

 SEQ  Entry
 ---------------------------------------------------------
 10    Action: permit
       IP    : 10.10.10.104      Mask: 0.0.0.0

 20    Action: deny (log)
       IP    : 10.10.10.1        Mask: 0.0.0.255

 30    Action: permit
       IP    : 0.0.0.0           Mask: 255.255.255.255
```

Note that each ACE is automatically assigned a sequence number.

Creating numbered, standard ACLs

Use the following steps when creating or adding to a numbered, standard ACL:

Procedure

1. Create a numbered, standard ACL by entering the first ACE in the list.

2. Append a new ACE to the end of an existing, standard ACL.

The following describes the commands for performing these steps.

This command is an alternative to using `ip access-list standard <name-str>` and does not use the "Named ACL"(`nacl`) context.

Syntax:

```
access-list <1-99> <deny|permit>
```

```
<any|host <SA>|<SA> <mask>|<SA>/<mask-length> > [log]
```

Appends an ACE to the end of the list of ACEs in the current IPv4 standard, numbered ACL. If the ACL does not already exist, creates both the ACL and its first ACE. In the default configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using `resequence`, see **Resequencing the ACEs in an ACL** on page 444.

> **NOTE**
>
> To insert a new ACE between two existing ACEs in a standard, numbered ACL:
>
> • Use `ip access list extended <1-99>` to open the ACL as a named ACL.
>
> • Enter the desired sequence number along with the ACE keywords and variables you want.
>
> (After a numbered ACL has been created, it can be managed as either a named or numbered ACL.)

```
<1-99>
```

Specifies the ACL identifier as a number. The switch interprets an ACL with a value in this range as a standard ACL (which filters all IPv4 traffic on the basis of SA). To create a standard access list with an alphanumeric name (`name-str`) instead of a number, see **Configuring named, standard ACLs** on page 421.

```
<deny|permit>
```

Specifies whether the ACE denies or permits a packet matching the criteria in the ACE, as described next.

```
<any|host <SA>|<SA> <mask>|<SA>/<mask-length> > [log]
```

Defines the source IPv4 address (SA) a packet must carry for a match with the ACE.

• `any` – Allows IPv4 packets from any SA.

• `host <SA>` – Specifies only packets having **<SA>** as the source. Use this criterion when you want to match only the IPv4 packets from a single SA.

`<SA> <mask>` or `<SA>/<mask-length>` – Specifies packets received from an SA, where the SA is either a subnet or a group of IPv4 addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits). `SA Mask` application: The mask is applied to the SA in the ACE to define which bits in a packet's SA must exactly match the SA configured in the ACL and which bits need not match.

Example:

10.10.10.1/24 and **10.10.10.1 0.0.0.255** both define any address in the range of 10.10.10.(1–255).

> **NOTE**
>
> Specifying a group of contiguous addresses may require more than one ACE.

```
[log]
```

This option can be used after the SA to generate an Event Log message if:

- The action is deny or permit.

- There is a match.

- ACL logging is enabled.

Creating and viewing a standard ACL

This example creates a standard, numbered ACL with the same ACE content as shown in **Configuring ACEs in a named, standard ACL**.

Note that each ACE is automatically assigned a sequence number.

```
HP Switch(config)# access-list 17 permit host 10.10.10.104
HP Switch(config)# access-list 17 deny 10.10.10.1/24 log
HP Switch(config)# access-list 17 permit any
HP Switch(config)# show access-list 17

Access Control Lists

  Name: 17
  Type: Standard
  Applied: No

 SEQ   Entry
 -------------------------------------------------------
 10    Action: permit
       IP   : 10.10.10.104        Mask: 0.0.0.0

 20    Action: deny (log)
       IP   : 10.10.10.1          Mask: 0.0.0.255

 30    Action: permit
       IP   : 0.0.0.0             Mask: 255.255.255.255
```

Configuring extended ACLs

Standard ACLs use only source IPv4 addresses for filtering criteria, extended ACLs use multiple filtering criteria. This enables you to more closely define your IPv4 packet-filtering.

Extended ACLs enable filtering on source and destination IPv4 addresses (required), in one of the following options:

- Source and destination IPv4 addresses for filtering criteria, extended ACLs use multiple filtering criteria. This enables you to more closely define your IPv4 packet filtering. Extended ACLs enable filtering on the following:

  ◦ specific host

  ◦ subnet or group of addresses

  ◦ any address

- choice of any IPv4 protocol

- optional packet-type criteria for IGMP and ICMP traffic

- optional source and/or destination TCP or UDP port

- filtering for TCP traffic based on either TCP control bits or whether the subject traffic is initiating a connection ("established" option)

- optional IP precedence and ToS criteria

HP Switches allow up to 512 ACLs, and determines the total from the number of unique identifiers in the configuration. For example, configuring two ACLs results in an ACL total of two, even if neither is assigned to an interface. If you then assign a nonexistent ACL to an interface, the new ACL total is three, because the switch now has three unique ACL names in its configuration.

Configuring named, extended ACLs

For a match to occur with an ACE in an extended ACL, a packet must have the source and destination address criteria specified by the ACE, as well as any IPv4 protocol-specific criteria included in the command.

Use the following general steps to create or add to a named, extended ACL:

Procedure

1. Create and/or enter the context of a named, extended ACL.

2. Enter the first ACE in a new, extended ACL or append an ACE to the end of an existing, extended ACL.

The following command is a prerequisite to entering or editing ACEs in a named, extended ACL.

Syntax:

```
ip access-list extended <name-str>
```

Places the CLI in the "Named ACL" (`nacl`) context specified by the `<name-str>` alphanumeric identifier. This enables entry of individual ACEs in the specified ACL. If the ACL does not already exist, this command creates it.

```
<name-str>
```

Specifies an alphanumeric identifier for the ACL. Consists of an alphanumeric string of up to 64 case-sensitive characters. Including spaces in the string requires that you enclose the string in single or double quotes. For example: `accounting ACL`. You can also use this command to access an existing, numbered ACL. See **Using the CLI to edit ACLs**.

Configuring ACEs in named, extended ACLs

Configuring ACEs is done after using the `ip access-list standard <name-str>` command described.

See the section "Standard ACL structure" for filtering criteria, extended ACLs use multiple filtering criteria. This enables you to more closely define your IPv4 packet-filtering.

Syntax: (nacl context)

```
<deny|permit> <ip|ip-protocol|ip-protocol-nbr>
<any|host> <SA>|SA|mask-length|SA <mask>>
<any|host> <DA>|DA|mask-length|DA <mask>>
[precedence] [tos] [log]
```

Appends an ACE to the end of the list of ACEs in the current ACL. In the default configuration, ACEs are automatically assigned consecutive sequence numbers in increments of 10 and can be renumbered using `resequence`, see **Resequencing the ACEs in an ACL** on page 444).

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

To insert a new ACE between two existing ACEs in an extended, named ACL, precede `deny` or `permit` with an appropriate sequence number along with the ACE keywords and variables you want. See **Inserting an ACE in an existing ACL** on page 442.

For a match to occur, a packet must have the source and destination addressing criteria specified in the ACE, as well as:

- the protocol-specific criteria configured in the ACE, including any included, optional elements (described later in this section)

- any (optional) precedence and/or ToS settings configured in the ACE

`<deny|permit>`

For named ACLs, these keywords are used in the "Named ACL" (`nacl`) context to specify whether the ACE denies or permits a packet matching the criteria in the ACE, as described below.

`<ip|ip-protocol|ip-protocol-nbr>`

Used after `deny` or `permit` to specify the packet protocol type required for a match. An extended ACL must include one of the following:

- `ip`

  – any IPv4 packet.

- `ip-protocol` – any one of the following IPv4 protocol names:

  ◦ `ip-in-ip`

  ◦ `ipv6-in-ip`

  ◦ `gre`

  ◦ `esp`

  ◦ `ah`

  ◦ `ospf`

  ◦ `pim`

  ◦ `vrrp`

  ◦ `sctp`

  ◦ `tcp*`

  ◦ `udp*`

  ◦ `icmp*`

  ◦ `igmp*`

*For TCP, UDP, ICMP, and IGMP, additional criteria can be specified.

- `ip-protocol-nbr` – the protocol number of an IPv4 packet type, such as "8" for Exterior Gateway Protocol or 121 for Simple Message Protocol. (For a listing of IPv4 protocol numbers and their corresponding protocol names, see theIANA "Protocol Number Assignment Services" at **http://www.iana.com**. (Range: 0–255)

`<any|host <SA>|SA <mask>|SA/mask-length`

This is the first instance of IPv4 addressing in an extended ACE. It follows the protocol specifier and defines the source address (SA) a packet must carry for a match with the ACE.

- `any` Allows IPv4 packets from any SA.

- `host <SA>` Specifies only packets having a single address as the SA. Use this criterion when you want to match only the IPv4 packets from a single SA.

- `SA <mask>` or `SA/mask-length` Specifies packets received from an SA, where the SA is either a subnet or a group of addresses. The mask can be in either dotted-decimal format or CIDR format (number of significant bits).

- `SA Mask` application. The mask is applied to the SA in the ACL to define which bits in a packet's SA must exactly match the SA configured in the ACL and which bits need not match.

  Example:10.10.10.1/24 and 10.10.10.1 0.0.0.255 both define any address in the range of 10.10.10.(1 - 255). Note: Specifying a group of contiguous addresses may require more than one ACE.

`<any|host <DA>|DA/mask-length|DA <mask>>`

This is the second instance of IPv4 addressing in an extended ACE. It follows the first (SA) instance, described earlier, and defines the destination address (DA) that a packet must carry in order to have a match with the ACE.

- `any` Allows routed IPv4 packets to any DA.

- `host <DA>` Specifies only packets having `DA`as the destination address. Use this criterion when you want to match only the IPv4 packets for a single DA.

- `DA/mask-length` or `DA <mask>` Specifies packets intended for a destination address, where the address is either a subnet or a group of addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits).

- `DA Mask` application The mask is applied to the DA in the ACL to define which bits in a packet's DA must exactly match the DA configured in the ACL and which bits need not match.

`[precedence <0-7|precedence-name>]`

This option can be used after the DA to cause the ACE to match packets with the specified IP precedence value. Values can be entered as the following IP precedence numbers or alphanumeric names:

0 **or** routine

1 " priority

2 " immediate

3 " flash

4 " flash-override

5 " critical

6 " internet (for internetwork control)

7 " network (for network control)

> **NOTE** The precedence criteria described in this section are applied in addition to any other selection criteria configured in the same ACE.

`[tos <`*`tos-setting`*`>]`

This option can be used after the DA to cause the ACE to match packets with the specified Type-of-Service (ToS) setting. ToS values can be entered as the following numeric settings or, in the case of 0, 2, 4, and 8, as alphanumeric names:

0 **or** normal

2 " max-reliability

4 " max-throughput

6

8 " minimize-delay

10

12

14

> **NOTE** The ToS criteria in this section are applied in addition to any other criteria configured in the same ACE.

`[log]`

This option can be used after the DA to generate an Event Log message if:

- The action is `deny`. Not applicable to `permit`.

- There is a match.

- ACL logging is enabled.

Including options for TCP and UDP traffic in extended ACLs

An ACE designed to permit or deny TCP or UDP traffic can optionally include port number criteria for either the source or destination, or both. Use of TCP criteria also allows the `established` option for controlling TCP connection traffic.

Syntax:

```
<deny|permit> <tcp|udp>
```

```
<SA> [comparison-operator <tcp/udp-src-port>]
```

```
<DA> [comparison-operator <tcp-dest-port>][established]
```

```
[comparison-operator <udp-dest-port>]
```

In an extended ACL using either `tcp` or `udp` as the packet protocol type, you can optionally use TCP or UDP source and/or destination port numbers or ranges of numbers to further define the criteria for a match. For example:

```
#deny tcp host 10.20.10.17 eq 23 host 10.20.10.155 established
#permit tcp host 10.10.10.100 host 10.20.10.17 eq telnet
#deny udp 10.30.10.1/24 host 10.20.10.17 range 161 162
```

[*comparison-operator <tcp/udp-src-port>*]

To specify a TCP or UDP source port number in an ACE:

(1) Select a comparison operator from the following list

and

(2) Enter the port number or a well-known port name.

Comparison operators

- `eq <tcp/udp-port-nbr>`

  "Equal To"; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be equal to `<tcp/udp-port-nbr>` .

- `gt <tcp/udp-port-nbr>`

  "Greater Than"; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be greater than `<tcp/udp-port-nbr>` .

- `lt <tcp/udp-port-nbr>`

  "Less Than"; to have a match with the ACE entry, the TCP or UDP source port number in a packet must be less than `<tcp/udp-port-nbr>` .

- `neq <tcp/udp-port-nbr>`

  "Not Equal"; to have a match with the ACE entry, the TCP or UDP source port number in a packet must not be equal to `<tcp/udp-port-nbr>` .

- `range <start-port-nbr> <end-port-nbr>`

  For a match with the ACE entry, the TCP or UDP source-port number in a packet must be in the range `<start-port-nbr> <end-port-nbr>` .

Port number or well-known port name:

Use the TCP or UDP port number required by your application.

The switch also accepts these well-known TCP or UDP port names as an alternative to their port numbers:

- TCP – bgp, dns, ftp, http, imap4, ldap, nntp, pop2, pop3, smtp, ssl, telnet
- UDP – bootpc, bootps, dns, ntp, radius, radius-old, rip, snmp, snmp-trap, tftp

To list the above names, press the **[Shift] [?]** key combination after entering an operator. For a comprehensive listing of port numbers, visit **http://www.iana.org/assignments/port-numbers**.

[*comparison-operator <tcp-dest-port>*][established]

[*comparison-operator <udp-dest-port>*]

This option, if used, is entered immediately after the `<DA>` entry.

To specify a TCP or UDP port number;

1. select a comparison operator

2. enter the port number or a well-known port name

Comparison operators and well-known port names:

These are the same as are used with the TCP/UDP source-port options, and are listed earlier in this command description.

```
[established]
```

This option applies only where TCP is the configured protocol type. It blocks the synchronizing packet associated with establishing a TCP connection in one direction on a VLAN while allowing all other IPv4 traffic for the same type of connection in the opposite direction. For example, a Telnet connect requires TCP traffic to move both ways between a host and the target device. Simply applying a `deny` to inbound Telnet traffic on a VLAN would prevent Telnet sessions in either direction because responses to outbound requests would be blocked. However, by using the `established` option, inbound Telnet traffic arriving in response to outbound Telnet requests would be permitted, but inbound Telnet traffic trying to establish a connection would be denied.

Options for ICMP traffic in extended ACLs

This option is useful where it is necessary to permit some types of ICMP traffic and deny other types, instead of simply permitting or denying all types of ICMP traffic. That is, an ACE designed to permit or deny ICMP traffic can optionally include an ICMP type and code value to permit or deny an individual type of ICMP packet while not addressing other ICMP traffic types in the same ACE. As an optional alternative, the ACE can include the name of an ICMP packet type.

Syntax:

```
<deny|permit> icmp <SA> <DA> [icmp-type [icmp-code]
```

```
<deny|permit> icmp <SA> <DA> [icmp-type-name][]|]
```

In an extended ACL using `icmp` as the packet protocol type (see above), you can optionally specify an individual ICMP packet type or packet type/code pair to further define the criteria for a match. This option, if used, is entered immediately after the destination address (DA) entry. The following example shows two ACEs entered in a Named ACL context:

```
#permit icmp any any host-unknown
```

```
#permit icmp any any 3 7
```

```
[icmp-type [icmp-code]
```

This option identifies an individual ICMP packet type as criteria for permitting or denying that type of ICMP traffic in an ACE.

- `icmp-type` — This value is in the range of 0 - 255 and corresponds to an ICMP packet type.

- `icmp-code`— This value is in the range of 0 - 255 and corresponds to an ICMP code for an ICMP packet type.For more information on ICMP type names, visit the Internet Assigned Numbers Authority (IANA) website

at **http://www.iana.com**, click "Protocol Number Assignment Services", and then go to the selections under "Internet Control Message Protocol (ICMP) Parameters".

[*icmp-type-name*]

These name options are an alternative to the `[icmp-type [icmp-code]]` methodology described above. For more information, visit the IANA website cited above.

- `administratively-prohibited`

- `alternate-address`

- `conversion-error`

- `dod-host-prohibited`

- `dod-net-prohibited`

- `echo`

- `echo-reply`

- `general-parameter-problem`

- `host-isolated`

- `host-precedence-unreachable`

- `host-redirect`

- `host-tos-redirect`

- `host-tos-unreachable`

- `host-unknown`

- `host-unreachable`

- `information-reply`

- `information-request`

- `mask-reply`

- `mask-request`

- `mobile-redirect`

- `net-redirect`

- `net-tos-redirect`

- `net-tos-unreachable`

- `net-unreachable`

- `network-unknown`

- `no-room-for-option`

- `option-missing`

- `packet-too-big`

- parameter-problem

- port-unreachable

- precedence-unreachable

- protocol-unreachable

- reassembly-timeout

- redirect

- router-advertisement

- router-solicitation

- source-quench

- source-route-failed

- time-exceeded

- timestamp-reply

- timestamp-request

- traceroute

- ttl-exceeded

- unreachable

Option for IGMP in extended ACLs

This option is useful where it is necessary to permit some types of IGMP traffic and deny other types instead of simply permitting or denying all types of IGMP traffic. That is, an ACE designed to permit or deny IGMP traffic can optionally include an IGMP packet type to permit or deny an individual type of IGMP packet while not addressing other IGMP traffic types in the same ACE.

Syntax:

```
<permit|deny> igmp <SA> <DA> [icmp-type]
```

In an extended ACL using `igmp` as the packet protocol type, you can optionally specify an individual IGMP packet type to further define the criteria for a match. This option, if used, is entered immediately after the destination addressing entry. The following example shows an IGMP ACE entered in the Named ACL context:

```
HP Switch(config-ext-nacl)# permit igmp any
any host-query
```

[*igmp-type*]

The complete list of IGMP packet-type options includes:

- dvmrp

- host-query

- host-report

- pim

- `trace`

- `v2-host-report`

- `v2-host-leave`

- `mtrace-reply`

- `mtrace-request`

- `v3-host-report`

For more information on IGMP packet types, visit the Internet Assigned Numbers Authority (IANA) website at **http://www.iana.com**, click "Protocol Number Assignment Services", and then go to the selections under "Internet Group Management Protocol (IGMP) Type Numbers".

Configuring numbered, extended ACLs

This section describes the commands for performing the following in a numbered, extended ACL:

- Creating the ACL by entering the first ACE in the list

- Appending a new ACE to the end of an existing ACL

Creating or adding to an extended, numbered ACL

Procedure

Use the `access-list` command as an alternative to using `ip access-list extended <name-str>`. This command does not use the `nacl` context. See **access-list** .

`access-list` command for creating a numbered, extended ACL

Syntax

```
access-list <100-199> <deny|permit> <ip|ip-protocol|ip-protocol-nbr>
<any|host <SA>|SA/mask-length|SA <mask>>
<any|host <DA>|DA/mask-length|DA <mask>>
[precedence <0-7|precedence-name>]
[tos <tos-bit-setting>]
[log]
```

Description

If the ACL does not already exist, this command creates the specified ACL and its first ACE. If the ACL already exists, the new ACE is appended to the end of the configured list of explicit ACEs. In the default configuration, the ACEs in an ACL will automatically be assigned consecutive sequence numbers in increments of 10 and can be renumbered with `resequence` see **Resequencing the ACEs in an ACL** on page 444.

NOTE

To insert a new ACE between two existing ACEs in an extended, numbered ACL:

1. Use `ip access list extended <100-199>` to open the ACL as a named ACL.

2. Enter the desired sequence number along with the ACE statement you want.

For a match to occur, a packet must have the source and destination addressing criteria specified in the ACE, as well as:

- The protocol-specific criteria configured in the ACE, including any included, optional elements (described later in this section.)

- Any (optional) precedence and/or ToS settings configured in the ACE.

Options

`<100-199>`

Specifies the ACL ID number. The switch interprets a numeric ACL with a value in this range as an extended ACL.

`<deny|permit>`

Specifies whether to deny (`drop`) or permit (forward) a packet that matches the criteria specified in the ACE, as described below.

`<ip|`*ip-protocol*`|`*ip-protocol-nbr*`>`

Specifies the packet protocol type required for a match. An extended ACL must include one of the following:

- `ip` – any IPv4 packet.

- `ip-protocol` – any one of the following IPv4 protocol names:

  - `ip-in-ip`

  - `ipv6-in-ip`

  - `gre`

  - `esp`

  - `ah`

  - `ospf`

  - `pim`

  - `vrrp`

  - `sctp`

  - `tcp*`

  - `udp*`

  - `icmp*`

  - `igmp*`

* For TCP, UDP, ICMP, and IGMP, additional criteria can be specified, as described later in this section.

- `ip-protocol-nbr` – the protocol number of an IPv4 packet type, such as "8" for Exterior Gateway Protocol or 121 for Simple Message Protocol. (For a listing of IPv4 protocol numbers and their corresponding protocol names, see the IANA "Protocol Number Assignment Services" at **http://www.iana.com**.) (Range: 0-255).

`<any|host <SA>|SA/mask-length|SA <mask>>`

In an extended ACL, this parameter defines the source address (SA) that a packet must carry in order to have a match with the ACE.

- `any` – Specifies all inbound IPv4 packets.

- `host <SA>` – Specifies only inbound IPv4 packets from a single address. Use this option when you want to match only the IPv4 packets from a single source address.

- `SA/mask-length` or `SA <mask>` – Specifies packets received from an SA, where the SA is either a subnet or a group of IPv4 addresses. The mask can be in either dotted-decimal format or CIDR format with the number of significant bits.

`SA mask application`

The mask is applied to the SA in the ACL to define which bits in a packet's source SA must exactly match the address configured in the ACL and which bits need not match.

Example:

10.10.10.1/24 and 10.10.10.1 0.0.0.255 both define any IPv4 address in the range of 10.10.10. (1-255).

NOTE

Specifying a group of contiguous IPv4 addresses may require more than one ACE.

This is the second instance of addressing in an extended ACE. It follows the first (SA) instance, described earlier, and defines the destination address (DA) that a packet must carry in order to have a match with the ACE.

Syntax:

`<any|host <DA>|DA/mask-length>>`

The options are the same as shown for `<SA>`.

- `any` – Allows routed IPv4 packets to any DA.

- `host <DA>` – Specifies only the packets having `DA` as the destination address. Use this criterion when you want to match only the IPv4 packets for a single DA.

- `DA/mask-length` or `DA <mask>` – Specifies packets intended for a destination address, where the address is either a subnet or a group of IPv4 addresses. The mask format can be in either dotted-decimal format or CIDR format (number of significant bits).

`DA Mask application`

The mask is applied to the DA in the ACL to define which bits in a packet's DA must exactly match the DA configured in the ACL and which bits need not match. See also the above example and note.

`[precedence <0-7|precedence-name>]`

This option causes the ACE to match packets with the specified IP precedence value. Values can be entered as the following IP precedence numbers or alphanumeric names:

0 **or** routine

1 " priority

2 " immediate

3 " flash

4 " flash-override

5 " critical

6 " internet (for internetwork control)

7 " network (for network control)

NOTE

The precedence criteria described in this section are applied in addition to any other selection criteria configured in the same ACE.

`[tos]`

This option can be used after the DA to cause the ACE to match packets with the specified Type-of-Service (ToS) setting. ToS values can be entered as the following numeric settings or, in the case of 0, 2, 4, and 8, as alphanumeric names:

0 **or** normal

2 " max-reliability

4 " max-throughput

6

8 " minimize-delay

10

12

14

NOTE

The ToS criteria in this section are applied in addition to any other criteria configured in the same ACE.

`[log]`

Optional; generates an Event Log message if:

- The action is `deny`. This option is not configurable for Permit.

- There is a match.

- ACL logging is enabled on the switch.

Controlling TCP and UDP traffic flow

An ACE designed to permit or deny TCP or UDP traffic can optionally include port number criteria for either the source or destination, or both. Use of TCP criteria also allows the `established` option for controlling TCP connection traffic. For a summary of the extended ACL syntax options, see **Including options for TCP and UDP traffic in extended ACLs** on page 429.

Syntax:

```
access-list <100-199> <deny|permit> <tcp|udp>

<SA> [comparison-operator <tcp/udp-src-port>]

<DA> [comparison-operator <tcp-dest-port>][established]

<DA> [comparison-operator <udp-dest-port>]
```

This source-port and destination-port TCP/UDP criteria is identical to the criteria described for TCP/UDP use in named, extended ACLs. See **Including options for TCP and UDP traffic in extended ACLs** on page 429.

Controlling ICMP traffic flow

This command is useful where it is necessary to permit some types of ICMP traffic and deny other types, instead of simply permitting or denying all types of ICMP traffic. That is, an ACE designed to permit or deny ICMP traffic can optionally include an ICMP type and code value to permit or deny an individual type of ICMP packet while not addressing other ICMP traffic types in the same ACE. As an optional alternative, the ACE can include the name of an ICMP packet type.

Syntax:

```
access-list <100-199> <deny|permit> icmp <SA> <DA>

[[icmp-type [icmp-code]]|[icmp-type-name]]
```

The ICMP "type" and "code" criteria are identical to the criteria described for ICMP in named, extended ACLs.

Controlling IGMP traffic flow

This command is useful where it is necessary to permit some types of IGMP traffic and deny other types, instead of simply permitting or denying all types of IGMP traffic. That is, an ACE designed to permit or deny IGMP traffic can optionally include an IGMP packet type to permit or deny an individual type of IGMP packet while not addressing other IGMP traffic types in the same ACE. As an optional alternative, the ACE can include the name of an ICMP packet type.

Syntax:

```
access-list <100-199>

<deny|permit> igmp <src-ip> <dest-ip> [igmp-type]
```

The IGMP "type" criteria is identical to the criteria described for IGMP in named, extended ACLs.

Adding or removing an ACL assignment on an interface

Filtering IPv4 traffic inbound on a VLAN

For a given VLAN interface on a switch configured for routing, you can assign an ACL as a VACL to filter inbound IPv4 traffic entering the switch on that VLAN. You can also use the same ACL for assignment to multiple VLANs.

Syntax:

Within the vlan context.

```
[no] vlan <vid> ip access-group <ACL> vlan
```

where: *<ACL>* = either a ACL name or an ACL ID number.

Assigns an ACL as a VACL to a VLAN to filter routed IPv4 traffic entering or leaving the switch on that VLAN. You can use either the global configuration level or the VLAN context level to assign or remove an VACL.

> **NOTE**
>
> The switch allows you to assign a nonexistent ACL name or number to a VLAN. In this case, if you subsequently configure an ACL with that name or number, it automatically becomes active on the assigned interface. Also, if you delete an assigned ACL from the switch without subsequently using the "**no**" form of this command to remove the assignment to an interface, the ACL assignment remains and will automatically activate any new ACL you create with the same identifier (name or number).

Figure 89: *Methods for enabling and disabling VACLs*

Filtering inbound IPv4 traffic per port

For a given port, port list, or static port trunk, you can assign an ACL as a static port ACL to filter any IPv4 traffic entering the switch on that interface. You can also use the same ACL for assignment to multiple interfaces. For limits and operating rules.

Syntax:

```
[no] interface <port-list|Trkx> ip access-group <identifier> in
```

where: *<identifier>* = either a ACL name or an ACL ID number.

Assigns an ACL as a static port ACL to a port, port list, or static trunk to filter any IPv4 traffic entering the switch on that interface. You can use either the global configuration level or the interface context level to assign or remove a static port ACL.

> **NOTE**
> The switch allows for assigning a nonexistent ACL name or number to a VLAN. In this case, if you subsequently configure an ACL with that name or number, it automatically becomes active on the assigned VLAN. Also, if deleting an assigned ACL from the switch without subsequently using the "**no**" form of this command to remove its assignment to a VLAN, the ACL assignment remains and automatically activates any new ACL created with the same identifier (name or number).

Figure 90: *Methods for enabling and disabling ACLs*

```
HP Switch(config)# interface 10 ip access-group My-List in    ◀── Enables a static port ACL
                                                                   from the Global
                                                                   Configuration level.
HP Switch(config)# interface 10
HP Switch(eth-10)# ip access-group 155 in        ◀────────────    Enables a static port ACL
HP Switch(eth-10)# exit                                            from a port context.

HP Switch(config)# no interface 10 ip access-group My-List in ◀── Disables a static port ACL
                                                                   from the Global
                                                                   Configuration level.
HP Switch(config)# interface 10
HP Switch(eth-10)# no ip access-group 155 in    ◀────────────     Uses a VLAN context to
HP Switch(eth-10)# exit                                           disable a static port ACL.
```

Deleting an ACL

Syntax:

```
no ip access-list standard <name-str|1-99>
```

```
no ip access-list extended <name-str|100-199>
```

```
no access-list <1-99|100-199>
```

Removes the specified ACL from the switch running-config file.

> **NOTE**
> Deleting an ACL does not delete any assignment of that ACL's identifier on a specific interface. Creating a new ACL using an identifier that is already configured on an interface causes the switch to automatically activate that ACL.

Editing an existing ACL

The CLI provides the capability for editing in the switch by using sequence numbers to insert or delete individual ACEs. An offline method is also available. This section describes using the CLI for editing ACLs.

Using the CLI to edit ACLs

You can use the CLI to delete individual ACEs from anywhere in an ACL, append new ACEs to the end of an ACL, and insert new ACEs anywhere within an ACL.

General editing rules

- Named ACLs:

◦ When you enter a new ACE in a named ACL without specifying a sequence number, the switch inserts the ACE as the last entry in the ACL.

◦ When you enter a new ACE in a named ACL and include a sequence number, the switch inserts the ACE according to the position of the sequence number in the current list of ACEs.

- Numbered ACLs: When using the `access-list <1-99|100-199>` command to create or add ACEs to a numbered ACL, each new ACE you enter is added to the end of the current list. (This command does not offer a `<seq-#>` option for including a sequence number to enable inserting an ACE at other points in the list.) Note, however, that once a numbered list has been created, you have the option of accessing it in the same way as a named list by using the `ip access-list <standard|extended>` command. This enables you to edit a numbered list in the same way that you would edit a named list. (See the next item in this list.)

- You can delete any ACE from any ACL (named or numbered) by using the `ip access-list` command to enter the ACL's context, and then using the `no <seq-#>` command, see **Deleting an ACE from an existing ACL** on page 443.

- Deleting the last ACE from an ACL leaves the ACL in memory. In this case, the ACL is "empty" and cannot perform any filtering tasks. (In any ACL the Implicit Deny does not apply unless the ACL includes at least one explicit ACE.)

Sequence numbering in ACLs

The ACEs in any ACL are sequentially numbered. In the default state, the sequence number of the first ACE in a list is "10" and subsequent ACEs are numbered in increments of 10. For example, the following `show run` output lists three ACEs with default numbering in a list named "My-List":

Figure 91: *The default sequential numbering for ACEs*

```
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   20 permit 10.20.10.117 0.0.0.0
   30 deny 10.20.10.1 0.0.0.255
   exit
```

You can add an ACE to the end of a named or numbered ACL by using either `access-list` for numbered ACLs or `ip access-list` for named ACLs:

Figure 92: *Adding an ACE to the end of numbered or named ACLs*

For example, to append a fourth ACE to the end of the ACL in **Figure 91: The default sequential numbering for ACEs** on page 441:

Figure 93: *Appending an ACE to an existing list*

```
HP Switch(config)# ip access-list standard My-List
HP Switch(config-std-nacl)# permit any
HP Switch(config-std-nacl)# show run
.
.
.
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   20 permit 10.20.10.117 0.0.0.0
   30 deny 10.20.10.1 0.0.0.255
   40 permit 0.0.0.0 255.255.255.255
   exit
```

NOTE

When using the `access-list <1-99|100-199> <permit|deny> <SA>` command to create an ACE for a numbered ACL, the ACE is always added to the end of the current list and given the appropriate sequence number. However, once a numbered list has been created, you can use the `ip access-list` command to open it as a named ACL and specify a nondefault sequence number, as described in the next section.

Inserting an ACE in an existing ACL

This action uses a sequence number to specify where to insert a new ACE into an existing sequence of ACLs.

Syntax:

```
ip access-list <standard|extended> <name-str|1-99|100-199>
```

```
<1-2147483647> <permit|deny> <standard-acl-ip-criteria> [log]
```

```
<1-2147483647> <permit|deny> <extended-acl-ip-criteria> [options]
```

The first command enters the "Named-ACL" context for the specified ACL. The remaining two commands insert a new ACE in a standard or extended ACL, respectively.

To insert a new ACE between existing ACEs in a list:

Procedure

1. Use `ip access-list` to enter the "Named-ACL" (`nacl`) context of the ACE. This applies regardless of whether the ACE was originally created as a numbered ACL or a named ACL.

2. Begin the ACE command with a sequence number that identifies the position you want the ACE to occupy. (The sequence number range is 1-2147483647).

3. Complete the ACE with the Command syntax appropriate for thetype of ACL you are editing.

For example, inserting a new ACE between the ACEs numbered 10 and 20 in the following figure requires a sequence number in the range of 11-19 for the new ACE.

Figure 94: *Inserting an ACE in an existing ACL*

```
HP Switch(config)# ip access-list standard My-List
HP Switch(config-std-nacl)# 15 deny 10.10.10.1/24
HP Switch(config-std-nacl)# show run
.
.
.
                                                    Enters the "Named-ACL
                                                    context for "My-List".
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   15 deny 10.10.10.1 0.0.0.255
   20 permit 10.20.10.117 0.0.0.0      Inserts the new ACE.
   30 deny 10.20.10.1 0.0.0.255
   40 permit 0.0.0.0 255.255.255.255
   exit
```

In the following example, the first two ACEs entered become lines 10 and 20 in the list. The third ACE entered is configured with a sequence number of 15 and is inserted between lines 10 and 20.

Figure 95: *Inserting an ACE into an existing sequence*

```
HP Switch(config)# ip access-list standard List-01
HP Switch(config-std-nacl)# permit 10.10.10.1/24   Becomes Line 10
HP Switch(config-std-nacl)# deny 10.10.1.1/16       Becomes Line 20
HP Switch(config-std-nacl)# 15 permit 10.10.20.1/24
HP Switch(config-std-nacl)# show run
                                                    Lines 10 and 20
Running configuration:                              were automatically
                                                    numbered accord-
. . .                                               ing to their order of
ip access-list standard "List-01"                   entry in the list. Line
   10 permit 10.10.10.1 0.0.0.255                   15 was explicitly
   15 permit 10.10.20.1 0.0.0.255                   numbered by the 15
   20 deny 10.10.1.1 0.0.255.255                    permit command
   exit                                             and was inserted in
                                                    its proper place in
                                                    the list.
```

Deleting an ACE from an existing ACL

This action uses ACL sequence numbers to delete ACEs from an ACL.

Syntax:

```
ip access-list <standard|extended> <name-str|1-99|100-199>

no <seq-#>
```

The first command enters the "Named-ACL" context for the specified ACL. The `no` command deletes the ACE corresponding to the sequence number entered.

Range: 1–2147483647

Procedure

1. To find the sequence number of the ACE you want to delete, use `show run` or `show access-list <name-str|1-99|100-199>` to view the ACL.

2. Use `ip access-list` to enter the "Named-ACL" (`nacl`) context of the ACE. This applies regardless of whether the ACE was originally created as a numbered ACL or a named ACL.

3. In the "Named-ACL" context, type `no` and enter the sequence number of the ACE you want to delete.

Figure 96: *Deleting an ACE from any ACL*

```
HP Switch(config)# show run
. . .
ACL Before Deleting an ACE
ip access-list standard "My-List"
    10 permit 10.10.10.25 0.0.0.0
    15 deny 10.10.10.1 0.0.0.255
    20 permit 10.20.10.117 0.0.0.0
    30 deny 10.20.10.1 0.0.0.255            This command enters the
    40 permit 0.0.0.0 255.255.255.255       "Named-ACL" (nacl)
    exit                                    context for "My-List".
HP Switch(config)# ip access-list standard My-List
HP Switch(config-std-nacl)# no 20          This command deletes the
HP Switch(config-std-nacl)# show run       ACE at line 20.
. . .
ACL After Deleting the ACE at Line 20
ip access-list standard "My-List"
    10 permit 10.10.10.25 0.0.0.0
    15 deny 10.10.10.1 0.0.0.255            The ACE at line 20 has been
    30 deny 10.20.10.1 0.0.0.255            removed.
    40 permit 0.0.0.0 255.255.255.255
    exit
```

Resequencing the ACEs in an ACL

This action reconfigures the starting sequence number for ACEs in an ACL, and resets the numeric interval between sequence numbers for ACEs configured in the ACL.

Syntax:

`ip access-list resequence <name-str|1-99|100-199>`

`<starting-seq-#> <interval>`

Resets the sequence numbers for all ACEs in the ACL.

`<starting-seq-#>`

Specifies the sequence number for the first ACE in the list. (Default: 10; Range: 1 – 2147483647)

`<interval>`

Specifies the interval between sequence numbers for the ACEs in the list. (Default: 10; Range: 1–2147483647)

Procedure

1. To view the current sequence numbering in an ACE, use `show run` or `show access-list <`*`name-str|`*
   `1-99|100-199> .`

2. Use the Command syntax (above) to change the sequence numbering.

This example resequences the "My-List" ACL at the bottom of figure so that the list begins with line 100 and uses a sequence interval of 100.

Figure 97: *Viewing and resequencing an ACL*

```
HP Switch(config)# show run
. . .
ip access-list standard "My-List"
   10 permit 10.10.10.25 0.0.0.0
   15 deny 10.10.10.1 0.0.0.255
   30 deny 10.20.10.1 0.0.0.255
   40 permit 0.0.0.0 255.255.255.255
   exit
. . .
HP Switch(config)# ip access-list resequence My-List 100 100
HP Switch(config)# show run
. . .
ip access-list standard "My-List"
   100 permit 10.10.10.25 0.0.0.0
   200 deny 10.10.10.1 0.0.0.255
   300 deny 10.20.10.1 0.0.0.255
   400 permit 0.0.0.0 255.255.255.255
  exit
```

Attaching a remark to an ACE

A remark is numbered in the same way as an ACE, and uses the same sequence number as the ACE to which it refers. This operation requires that the remark for a given ACE be entered prior to entering the ACE itself.

Syntax:

`access-list <1-99|100-199> remark <`*`remark-str`*`>`

This syntax appends a remark to the end of a numbered ACL and automatically assigns a sequence number to the remark. The next command entry should be the ACE to which the remark belongs. (The new ACE will automatically be numbered with the same sequence number as was used for the preceding remark.)

Syntax:

`ip access-list <standard|extended> <`*`name-str`*`|1-99|100-199> [seq-#]remark <`*`remark-str`*`> no <`*`seq-#`*`> remark`

This syntax applies to both named and numbered ACLs. Without an optional sequence number, the remark is appended to the end of the list and automatically assigned a sequence number. When entered with an optional sequence number, the remark is inserted in the list according to the numeric precedence of the sequence number. The `no` form of the command deletes the indicated remark, but does not affect the related ACE.

To associate a remark with a specific ACE, enter the remark first, and then enter the ACE.

- Entering a remark without a sequence number and then entering an ACE without a sequence number results in the two entries being automatically paired with the same sequence number and appended to the end of the current ACL.

- Entering a remark with a sequence number and then entering an ACE with the same sequence number results in the two entries being paired together and positioned in the list according to the sequence number they share.

NOTE

After a numbered ACL has been created (using `access-list <1-99|100-199>`), it can be managed as either a named or numbered ACL. For example, in an existing ACL with a numeric identifier of "115", either of the following command sets adds an ACE denying IPv4 traffic from any source to a host at 10.10.10.100:

```
HP Switch(config)# access-list 115 deny ip host 10.10.10.100

HP Switch(config)# ip access-list extended 115

HP Switch(config-ext-nacl)# deny ip any 10.10.10.100
```

Appending remarks and related ACEs to the end of an ACL

To include a remark for an ACE that will be appended to the end of the current ACL, enter the remark first, then enter the related ACE. This results in the remark and the subsequent ACE having the same sequence number. For example, to add remarks using the "Named-ACL" (`nacl`) context:

Figure 98: *Appending a remark and its related ACE to the end of an ACL*


```
HP Switch(config)# ip access-list standard My-List
HP Switch(config-std-nacl)# permit host 10.10.10.15
HP Switch(config-std-nacl)# deny 10.10.10.1/24
HP Switch(config-std-nacl)# remark HOST-10.20.10.34
HP Switch(config-std-nacl)# permit host 10.20.10.34
HP Switch(config-std-nacl)# show run
. . .
hostname "Networking"
ip access-list standard "My-List"
    10 permit 10.10.10.15 0.0.0.0
    20 deny 10.10.10.1 0.0.0.255
    30 remark "HOST-10.20.10.34"
    30 permit 10.20.10.34 0.0.0.0
    exit
```

The remark is assigned the same number that the immediately following ACE ("30" in this example) is assigned when it is automatically appended to the end of the list. This operation applies where new remarks and ACEs are appended to the end of the ACL and are automatically assigned a sequence number.

You can also perform the operation illustrated in **Figure 98: Appending a remark and its related ACE to the end of an ACL** on page 446 by using the numbered, `access-list <1-99|100-199>` syntax shown at the beginning of this section.

Inserting remarks and related ACEs within an existing list

To insert an ACE with a remark within an ACL by specifying a sequence number, insert the numbered remark first, then, using the same sequence number, insert the ACE. This operation applies only to ACLs accessed using the "Named-ACL" (`nacl`) context.

Figure 99: *Inserting remarks*

```
HP Switch(config-std-nacl)# 15 remark "HOST 10.10.10.21"
HP Switch(config-std-nacl)# 15 permit host 10.10.10.21
HP Switch(config-std-nacl)# show run
ip access-list standard "My-List"
    10 permit 10.10.10.15 0.0.0.0
    15 remark "HOST 10.10.10.21"
    15 permit 10.10.10.21 0.0.0.0
    20 deny 10.10.10.1 0.0.0.255
    30 remark "HOST-10.20.10.34"
    30 permit 10.20.10.34 0.0.0.0
    exit
```

Inserting a remark/ACE pair with the same sequence number requires that the remark (with the desired sequence number) be inserted **before** the ACE with the same number.

Inserting a remark for an ACE that already exists in an ACL

If a sequence number is already assigned to an ACE in a list, you cannot insert a remark by assigning it to the same number. (To configure a remark with the same number as a given ACE, the remark must be configured first.) To assign a remark to the same number as an existing ACE:

Procedure

1. Delete the ACE.

2. Configure the remark with the number you want assigned to the pair.

3. Re-Enter the deleted ACE with the number used to enter the remark.

Removing a remark from an existing ACE

If you want to remove a remark, but want to retain the ACE, do the following:

Procedure

1. Use the Named ACL context to enter the ACL.

2. Using `show run` or `show access-list <list-name> config` , note the sequence number and content of the ACE having a remark you want to remove.

3. Delete the ACE.

4. Using the same sequence number, re-enter the ACE.

Operating notes for remarks

- The `resequence` command ignores "orphan" remarks that do not have an ACE counterpart with the same sequence number. For example, if:

- a remark numbered "55" exists in an ACE

- there is no ACE numbered "55" in the same ACL

- `resequence`

  is executed on an ACL

then the remark retains "55" as its sequence number and will be placed in the renumbered version of the ACL according to that sequence number.

- Entering an unnumbered remark followed by a numbered ACE, or the reverse, creates an "orphan" remark. The unnumbered entry will be assigned a sequence number that is an increment from the last ACE in the list. The numbered entry will then be placed sequentially in the list according to the sequence number used.

- Configuring two remarks without either sequence numbers or an intervening, unnumbered ACE results in the second remark overwriting the first.

Figure 100: *Overwriting one remark with another*

```
HP Switch(config)# ip access-list standard Accounting
HP Switch(config-std-nacl)# permit host 10.10.10.115
HP Switch(config-std-nacl)# deny 10.10.10.1/24
HP Switch(config-std-nacl)# remark Marketing
HP Switch(config-std-nacl)# remark Channel_Mktg
HP Switch(config-std-nacl)# show run
.
.
.
ip access-list standard "Accounting"
    10 permit 10.10.10.115 0.0.0.0
    20 deny 10.10.10.1 0.0.0.255
    30 remark "Channel_Mktg"
    exit
```

Where multiple remarks are sequentially entered for automatic inclusion at the end of an ACL, each successive remark replaces the previous one until an ACE is configured for automatic inclusion at the end of the list.

Viewing ACL configuration data

The show commands in this section apply to both IPv4 and IPv6 ACLs. For information on IPv6 ACL operation, see the chapter titled "IPv6 Access Control Lists" in the IPv6 configuration guide for your switch.

Viewing an ACL summary

This command lists the configured IPv4 and IPv6 ACLs, regardless of whether they are assigned to any VLANs.

Syntax:

```
show access-list
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

List a summary table of the name, type, and application status of IPv4 and IPv6 ACLs configured on the switch.

Figure 101: *Summary table of access list*

```
HP Switch(config)# show access-list

Access Control Lists

 Type  Appl  Name
 ----  ----  ----------------------------------------
   ext   yes   101
   std   yes   55
   ext   yes   Marketing
   ipv6   no   Accounting
   ipv6   no   List-01-Inbound
   ipv6  yes   List-02-Outbound
   ipv6  yes   Test-1
```

IPv4 ACLs

These ACLs exist in the configuration but are not applied to any interfaces and thus do not affect traffic.

Term	Meaning
Type	Shows whether the listed ACL is an IPv4 `std` ACL, an IPv4 `ext`ACL, or an IPv6 ACL.
Appl	Shows whether the listed ACL has been applied to an interface (`yes`/`no`).
Name	Shows the identifier (name or number) assigned to each ACL configured in the switch.

Viewing the content of all ACLs on the switch

This command lists the configuration details for the IPv4 and IPv6 ACLs in the running-config file, regardless of whether any are actually assigned to filter IPv4 traffic on specific VLANs.

Syntax:

```
show access-list config
```

List the configured syntax for all IPv4 ACLs currently configured on the switch.

> **NOTE**
>
> Notice that you can use the output from this command for input to an offline text file in which you can edit, add, or delete ACL commands.
>
> This information also appears in the `show running` display. If you executed `write memory` after configuring an ACL, it appears in the `show config` display.

For example, with two ACLs configured in the switch, you will see results similar to the following:

An ACL configured syntax listing

```
HP Switch(config)# show access-list config

ip access-list standard "List-43"
   10 deny 10.28.236.77 0.0.0.0
   20 deny 10.29.140.107 0.0.0.0
   30 permit 0.0.0.0 255.255.255.255
   exit
ip access-list extended "111"
```

```
10 permit tcp 10.30.133.27 0.0.0.0 0.0.0.0 255.255.255.255
20 permit tcp 10.30.155.101 0.0.0.0 0.0.0.0 255.255.255.255
30 deny ip 10.30.133.1 0.0.0.0 0.0.0.0 255.255.255.255 log
40 deny ip 10.30.155.1 0.0.0.255 0.0.0.0 255.255.255.255
exit
```

Viewing the VACL assignments for a VLAN

This command briefly lists the identification and types of IPv4 and IPv6 VACLs currently assigned to a particular VLAN in the running-config file. For IPv6, the switch supports, per-VLAN, one VACL assignment.

Syntax:

```
show access-list vlan <vid>
```

Lists the current ACL assignments to the specified VLAN (in the running config file).

NOTE	This information also appears in the `show running` display. If you execute `write memory` after configuring an ACL, it also appears in the `show config` display.

Viewing static port (and trunk) ACL assignments

This command lists the identification and types of current static port ACL assignments to individual switch ports and trunks, as configured in the running-config file. The switch allows one static port ACL assignment per port.

Syntax:

```
show access-list ports <all|interface>
```

Lists the current static port ACL assignments for ports and trunks in the running config file.

NOTE	This information also appears in the `show running` output. If you execute `write memory` after configuring an ACL, it also appears in the `show config` output.

For example, if you assigned a standard ACL with an ACL-ID of "Port-10" to filter inbound IP traffic on switch ports B10-B11 and trunk trk1, you could verify these assignments as shown in **Figure 102: Listing the ACL assignments for ports and trunks** on page 451.

Figure 102: *Listing the ACL assignments for ports and trunks*

```
HP Switch(config)# show access-list ports all

Access Lists for Port B10
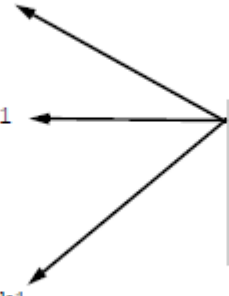
  Inbound  : 15
  Type     : Standard


Access Lists for Port B11

  Inbound  : 15
  Type     : Standard


Access Lists for Port Trk1

  Inbound  : 15
  Type     : Standard
```

Indicates that a standard ACL with the ID of "15" is assigned to filter traffic entering the switch on ports B10 and B11, and on trunk Trk1.

Viewing the content of a specific ACL

This command displays a specific ACL configured in the running config file in an easy-to-read tabular format.

NOTE

This information also appears in the `show running` display. If you execute `write memory` after configuring an ACL, it also appears in the `show config` display.

Syntax:

`show access-list <acl-id>`

Displays detailed information on the content of a specific ACL configured in the running-config file.

For example, suppose you configured the following two ACLs in the switch:

ACL ID	Type	Desired Action
1	Standard	• Deny IP traffic from 18.28.236.77 and 18.29.140.107. • Permit IP traffic from all other sources.
105	Extended	• Permit any TCP traffic from 18.30.133.27 to any destination. • Deny any other IP traffic from 18.30.133. (1-255). • Permit all other IP traffic from any source to any destination.

Inspect the ACLs as follows:

Figure 103: *Example of a listing a standard ACL*

```
HP Switch(config)# show access-list 1

Access Control Lists

  Name: 1
  Type: Standard          Indicates whether the ACL
  Applied: Yes            is applied to an interface.

 SEQ  Entry
 ----------------------------------------------------------------------------
  10   Action: deny (log)
       IP    : 10.28.236.77      Mask: 0.0.0.0

  20   Action: deny
       IP    : 10.29.140.107     Mask: 0.0.0.0

  30   Action: permit
       IP    : 0.0.0.0           Mask: 255.255.255.255
```

Figure 104: *Examples of listings showing the content of standard and extended ACLs*

```
HP Switch(config)# show access-list List-120

Access Control Lists

  Name: List-120
  Type: Extended          Indicates whether the ACL
  Applied: No             is applied to an interface.

 SEQ  Entry
 ----------------------------------   Indicates source and destination entries in the ACL  -------
  10   Action: permit
       Remark: Telnet Allowed
       Src IP: 10.30.133.27     Mask: 0.0.0.0           Port(s): eq 23
       Dst IP: 0.0.0.0          Mask: 255.255.255.255   Port(s):
       Proto : TCP (Established)
       TOS   : -               Precedence: routine
                                                        Empty field indicates that
                                                        the destination TCP port
  20   Action: deny (log)                               can be any value.
       Src IP: 10.30.133.1      Mask: 0.0.0.255         Port(s)
       Dst IP: 0.0.0.0          Mask: 255.255.255.255   Port(s):
       Proto : IP
       TOS   : -               Precedence: -

  30   Action: permit
       Src IP: 0.0.0.0          Mask: 255.255.255.255   Port(s):
       Dst IP: 0.0.0.0          Mask: 255.255.255.255   Port(s):
       Proto : IP
       TOS   : -               Precedence: -
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Table 34: *Descriptions of data types included in show access-list <acl-id> output*

Field	Description
Name	The ACL identifier. Can be a number from 1 to 199, or a name.
Type	Standard or Extended. The former uses only source IPv4 addressing. The latter uses both source and destination IPv4 addressing and also allows TCP or UDP port specifiers.
Applied	"Yes" means the ACL has been applied to a port or VLAN interface. "No" means the ACL exists in the switch configuration, but has not been applied to any interface, and is therefore not in use.
SEQ	The sequential number of the Access Control Entry (ACE) in the specified ACL.
Entry	Lists the content of the ACEs in the selected ACL.
Action	Permit (forward) or deny (drop) a packet when it is compared to the criteria in the applicable ACE and found to match. Includes the optional log option, if used, in deny actions.
Remark	Displays any optional remark text configured for the selected ACE.
IP	Used for Standard ACLs: The source IPv4 address to which the configured mask is applied to determine whether there is a match with a packet.
Src IP	Used for Extended ACLs: Same as above.
Dst IP	Used for Extended ACLs: The source and destination IPv4 addresses to which the corresponding configured masks are applied to determine whether there is a match with a packet.
Mask	The mask configured in an ACE and applied to the corresponding IPv4 address in the ACE to determine whether a packet matches the filtering criteria.
Proto	Used only in extended ACLs to specify the packet protocol type to filter. Must be either IPv4, TCP, or UDP. For TCP protocol selections, includes the `established` option, if configured.
Port(s)	Used only in extended ACLs to show any TCP or UDP operator and port number(s) included in the ACE.
TOS	Used only in extended ACLs to indicate Type-of-Service setting, if any.
Precedence	Used only in extended ACLs to indicate the IP precedence setting, if any.

Viewing all ACLs and their assignments in the switch startup-config and running-config files

The `show config` and `show running` commands include in their listings any configured ACLs and any ACL assignments to VLANs. See **ACL configuration factors**. Remember that `show config` lists the startup-config file and `show running` lists the running-config file.

Creating or editing an ACL offline

The section titled **Editing an existing ACL** on page 440 describes how to use the CLI to edit an ACL, and is most applicable in cases where the ACL is short or there is only a minor editing task to perform. The offline method provides an alternative to using the CLI for creating or extensively editing a large ACL. This section describes how to:

- Move an existing ACL to a TFTP server

- Use a text (.txt) file format to create a new ACL or edit an existing ACL offline

- Use TFTP to load an offline ACL into the switch's running-config

For longer ACLs that may be difficult or time-consuming to accurately create or edit in the CLI, you can use the offline method described in this section.

> **NOTE**
>
> `Copy` commands that used either `tftp` or `xmodem` also include an option to use `usb` as a source or destination device for file transfers. So although the following example highlights tftp, bear in mind that `xmodem` or `usb` can also be used to transfer ACLs to and from the switch.

Procedure

1. Begin by doing one of the following:

   - To edit one or more existing ACLs, use `copy command-output tftp` to copy the current version of the ACL configuration to a file in your TFTP server. For example, to copy the ACL configuration to a file named acl-02.txt in the TFTP directory on a server at 10.28.227.2:

     ```
     HP Switch# copy command-output 'show access-list
     config' tftp 10.28.227.2 acl02.txt pc
     ```

   - To create a new ACL, open a text (.txt) file in the appropriate directory on a TFTP server accessible to the switch.

2. Use a text editor to create or edit the ACL(s) in the `*.txt` ASCII file format.

   If you are replacing an ACL on the switch with a new ACL that uses the same number or name syntax, begin the command file with a `no ip access-list` command to remove the earlier version of the ACL from the switch running-config file. Otherwise, the switch will append the new ACEs in the ACL you download to the existing ACL.

For example, if you planned to use the `copy` command to replace ACL "List-120", place this command at the beginning of the edited file:

```
no ip access-list extended List-120
```

Figure 105: *An offline ACL file designed to replace an existing ACL*

```
no ip access-list extended List-120
ip access-list extended "List-120"
  10 remark "THIS ACE ALLOWS TELNET"
  10 permit tcp 10.30.133.27 0.0.0.0 eq 23 0.0.0.0 255.255.255..
  20 deny ip 10.30.133.1 0.0.0.255 0.0.0.0 255.255.255.255
  30 deny ip 10.30.155.1 0.0.0.255 0.0.0.0 255.255.255.255
  40 remark "THIS IS THE FINAL ACE IN THE LIST"
  40 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
  exit
```

Removes an existing ACL and replaces it with a new version with the same identity. To append new ACEs to an existing ACL instead of replacing it, you would omit the first line and ensure that the sequence numbering for the new ACEs begin with a number greater than the highest number in the existing list.

3. Use `copy tftp command-file` to download the file as a list of commands to the switch.

4. Use the `show run` or `show access-list config` command to inspect the switch configuration to ensure that the ACL was properly downloaded.

NOTE

If a transport error occurs, the switch does not execute the command and the ACL is not configured.

Monitoring static ACL performance

ACL statistics counters provide a means for monitoring ACL performance by using counters to display the current number of matches the switch has detected for each ACE in an ACL assigned to a switch interface. This can help in determining whether a particular traffic type is being filtered by the intended ACE in an assigned list, or if traffic from a particular device or network is being filtered as intended.

NOTE

This section describes the command for monitoring static ACL performance. To monitor RADIUS-assigned ACL performance, use either of the following commands:

```
show access-list radius <all|port-list>
```

```
show port-access <authenticator|mac-based|web-based> clients <port-list>
detailed
```

See **Displaying the current RADIUS-assigned ACL activity on the switch** on page 283.

Syntax:

```
<show|clear> statistics
```

```
aclv4 <acl-name-str> port <port-#>
```

```
aclv4 <acl-name-str> vlan <vid> <in|out|vlan>
```

```
aclv6 <acl-name-str> port <port-#>
```

```
aclv6 <acl-name-str> vlan <vid> <in|out|vlan>
```

`show`: Displays the current match (hit) count per ACE for the specified IPv6 or IPv4 static ACL assignment on a specific interface:

`clear`: Resets ACE hit counters to zero for the specified IPv6 or IPv4 static ACL assignment on a specific interface.

Total: This column lists the running total of the matches the switch has detected for the ACEs in an applied ACL since the ACL's counters were last reset to 0 (zero)

Figure 106: *IPv6 and IPv4 ACL statistics*

```
HP Switch# show statistics aclv6 IPV6-ACL vlan 20 vlan

 HitCounts for ACL IPV6-ACL

  Total

( 12) 10 permit icmp ::/0 fe80::20:2/128 128
( 6) 20 deny tcp ::/0 fe80::20:2/128 eq 23 log
( 41) 30 permit ipv6 ::/0 ::/0

HP Switch# show statistics aclv4 102 vlan 20 vlan

 HitCounts for ACL 102

  Total
(  4) 10 permit icmp 10.10.20.3 0.0.0.0 10.10.20.2 0.0.0.0 8
(  8) 20 deny icmp 0.0.0.0 255.255.255.255 10.10.20.2 0.0.0.0 8
(  2) 30 permit tcp 10.10.20.3 0.0.0.255 10.10.20.2 0.0.0.255 eq 23
(  2) 55 deny tcp 0.0.0.0 255.255.255.255 10.10.20.2 0.0.0.0 8
( 125) 60 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
```

ACE Counter Operation: For a given ACE in an assigned ACL, the counter increments by 1 each time the switch detects a packet that matches the criteria in that ACE, and maintains a running total of the matches since the last counter reset.

For example, in ACL line 10 below, there has been a total of 37 matches on the ACE since the last time the ACL's counters were reset.

```
Total

( 37) 10 permit icmp 10.10.20.3
```

> **NOTE**
>
> This ACL monitoring feature does not include hits on the "implicit deny" that is included at the end of all ACLs.

Resetting ACE Hit Counters to Zero:

- Removing an ACL from an interface zeros the ACL's ACE counters for that interface only.

- For a given ACL, either of the following actions clear the ACE counters to zero for all interfaces to which the ACL is assigned.

  ◦ adding or removing a permit or deny ACE in the ACL

  ◦ rebooting the switch

Below is an example of performance monitoring output for an IPv6 ACL assigned as a VACL.

Figure 107: *IPv6 ACL performance monitoring output*

```
HP Switch# show statistics aclv6 V6-02 vlan 20 vlan

 HitCounts for ACL V6-02

  Total

(   5) 10 permit icmp ::/0 fe80::20:2/128 128
(   4) 20 permit icmp ::/0 fe80::20:3/128 128
( 136) 30 permit tcp fe80::20:1/128 ::/0 eq 23
(   2) 40 deny icmp ::/0 fe80::20:1/128 128
(  10) 50 deny tcp ::/0 ::/0 eq 23
(   8) 60 deny icmp ::/0 ::/0 133
( 155) 70 permit ipv6 ::/0 ::/0
```

Below is an example of performance monitoring output for an IPv4 ACL assigned as a VACL.

Figure 108: *IPv4 ACL performance monitoring output*

```
HP Switch# show statistics aclv4 102 vlan 20 vlan

 HitCounts for ACL 102

  Total

(  1) 10 permit icmp 10.10.20.3 0.0.0.0 10.10.20.2 0.0.0.0 8
(  2) 20 deny icmp 10.10.20.3 0.0.0.0 10.10.20.1 0.0.0.0 8 log
(  2) 30 deny icmp 10.10.20.2 0.0.0.0 10.10.20.3 0.0.0.0 8 log
(  1) 40 deny icmp 10.10.20.2 0.0.0.0 10.10.20.1 0.0.0.0 8 log
( 10) 50 deny tcp 10.10.20.2 0.0.0.255 10.10.20.3 0.0.0.255 eq 23 log
( 27) 60 permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
```

The following example demonstrates using `clear statistics` to reset the counters to zero.

Figure 109: *IPv6 ACL performance monitoring output*

```
HP Switch# show statistics aclv6 V6-02 vlan 20 vlan

 HitCounts for ACL V6-02

  Total

(   5) 10 permit icmp ::/0 fe80::20:2/128 128
(   4) 20 permit icmp ::/0 fe80::20:3/128 128
( 136) 30 permit tcp fe80::20:1/128 ::/0 eq 23
(   2) 40 deny icmp ::/0 fe80::20:1/128 128
(  10) 50 deny tcp ::/0 ::/0 eq 23
(   8) 60 deny icmp ::/0 ::/0 133
( 155) 70 permit ipv6 ::/0 ::/0
HP Switch# clear statistics aclv6 V6-02 vlan 20 vlan
HP Switch# show statistics aclv6 V6-02 vlan 20 vlan

 HitCounts for ACL V6-02

  Total

( 0) 10 permit icmp ::/0 fe80::20:2/128 128
( 0) 20 permit icmp ::/0 fe80::20:3/128 128
( 0) 30 permit tcp fe80::20:1/128 ::/0 eq 23
( 0) 40 deny icmp ::/0 fe80::20:1/128 128
( 0) 50 deny tcp ::/0 ::/0 eq 23
( 0) 60 deny icmp ::/0 ::/0 133
( 0) 70 permit ipv6 ::/0 ::/0
```

General ACL operating notes

- **ACLs do not affect serial port access.**

  ACLs do not apply to the switch's serial port.

- **ACL screening of IPv4 traffic generated by the switch.**

  ACLs applied on the switch screen IP traffic when other devices generate it. Similarly, ACL applications can screen responses from other devices to unscreened IP traffic the switch generates.

- **ACL logging.**

  ○ The ACL logging feature generates a message only when packets are explicitly denied as the result of a match, and not when implicitly denied. To help test ACL logging, configure the last entry in an ACL as an explicit deny statement with a log statement included.

  ○ Logging enables you to selectively test specific devices or groups. However, excessive logging can affect switch performance. For this reason, Hewlett Packard Enterprise recommends that you remove the logging option from ACEs for which you do not have a current need. Also, avoid configuring logging where it does not serve an immediate purpose. (Note that ACL logging is not designed to function as an accounting method.) See also "Apparent Failure To Log All 'Deny' Matches" in the section titled "ACL Problems", found in "Troubleshooting" of the management and configuration guide for your switch.

  ○ When configuring logging, you can reduce excessive resource use by configuring the appropriate ACEs to match with specific hosts instead of entire subnets.

- **Minimum number of ACEs in an ACL.**

  Any ACL must include at least one ACE to enable IP traffic screening. A numbered ACL cannot be created without at least one ACE. A named ACL can be created "empty"; that is, without any ACEs. However in an empty ACL applied to an interface, the Implicit Deny function does not operate, and the ACL has no effect on traffic.

- **Monitoring shared resources.**

  Applied ACLs share internal switch resources with several other features. The switch provides ample resources for all features. However, if the internal resources become fully subscribed, additional ACLs cannot be applied until the necessary resources are released from other applications. For information on determining current resource availability and usage, see "Monitoring Resources" in the management and configuration guide for your switch.

- **Protocol Support.**

  IP ACL criteria does not include use of MAC information or QoS.

- **Replacing or adding to an active ACL policy.**

  If you assign an ACL to an interface and subsequently add or replace ACEs in that ACL, each new ACE becomes active when you enter it. If the ACL is configured on multiple interfaces when the change occurs, then the switch resources must accommodate all applications of the ACL. If there are insufficient resources to accommodate one of several ACL applications affected by the change, then the change is not applied to any of the interfaces and the previous version of the ACL remains in effect.

Introduction

As your network expands to include an increasing number of mobile devices, continuous Internet access, and new classes of users (such as partners, temporary employees, and visitors), additional protection from attacks launched from both inside and outside your internal network is often necessary.

Advanced threat protection can detect port scans and hackers who try to access a port or the switch itself. The following software features provide advanced threat protection and are described here:

- DHCP snooping: Protects your network from common DHCP attacks, such as:

  ◦ Address spoofing in which an invalid IP address or network gateway address is assigned by a rogue DHCP server

  ◦ Address exhaustion of available addresses in the network DHCP server, caused by repeated attacker access to the network and numerous IP address requests

- Dynamic ARP protection: Protects your network from ARP cache poisoning as in the following cases:

  ◦ An unauthorized device forges an illegitimate ARP response and network devices use the response to update their ARP caches

  ◦ A denial-of-service (DoS) attack from unsolicited ARP responses changes the network gateway IP address so that outgoing traffic is prevented from leaving the network and overwhelms network devices

- Instrumentation monitor: Protects your network from a variety of other common attacks besides DHCP and ARP attacks, including:

  ◦ Attempts at a port scan to expose a vulnerability in the switch, indicated by an excessive number of packets sent to closed TCP/UDP ports.

  ◦ Attempts to fill all IP address entries in the switch's forwarding table and cause legitimate traffic to be dropped, indicated by an increased number of learned IP destination addresses.

  ◦ Attempts to spread viruses, indicated by an increased number of ARP request packets

  ◦ Attempts to exhaust system resources so that sufficient resources are not available to transmit legitimate traffic, indicated by an unusually high use of specific system resources

  ◦ Attempts to attack the switch's CPU and introduce delay in system response time to new network events

  ◦ Attempts by hackers to access the switch, indicated by an excessive number of failed logins or port authentication failures

  ◦ Attempts to deny switch service by filling the forwarding table, indicated by an increased number of learned MAC addresses or a high number of MAC address moves from one port to another

  ◦ Attempts to exhaust available CPU resources, indicated by an increased number of learned MAC address events being discarded

DHCP snooping

You can use DHCP snooping to help avoid the Denial of Service attacks that result from unauthorized users adding a DHCP server to the network that then provides invalid configuration data to other DHCP clients on the network. DHCP snooping accomplishes this by allowing you to distinguish between trusted ports connected to a DHCP server or switch and untrusted ports connected to end-users. DHCP packets are forwarded between trusted ports without inspection. DHCP packets received on other switch ports are inspected before being forwarded. Packets from untrusted sources are dropped. Conditions for dropping packets are shown below.

Table 35: *Condition for dropping a packet*

Condition for Dropping a Packet	Packet Types
A packet from a DHCP server received on an untrusted port	DHCPOFFER, DHCPACK, DHCPNACK
If the switch is configured with a list of authorized DHCP server addresses and a packet is received from a DHCP server on a trusted port with a source IP address that is not in the list of authorized DHCP server addresses.	DHCPOFFER, DHCPACK, DHCPNACK
Unless configured to not perform this check, a DHCP packet received on an untrusted port where the DHCP client hardware address field does not match the source MAC address in the packet	N/A
Unless configured to not perform this check, a DHCP packet containing DHCP relay information (option 82) received from an untrusted port	N/A
A broadcast packet that has a MAC address in the DHCP binding database, but the port in the DHCP binding database is different from the port on which the packet is received	DHCPRELEASE, DHCPDECLINE

Enabling DHCP snooping

DHCP snooping is enabled globally by entering this command:

```
HP Switch(config)# dhcp-snooping
```

Use the `no` form of the command to disable DHCP snooping.

Syntax:

```
[no] dhcp-snooping [authorized-server|database|option|trust|verify|vlan]
```

authorized server

Enter the IP address of a trusted DHCP server. If no authorized servers are configured, all DHCP server addresses are considered valid. Maximum: 20 authorized servers.

database

To configure a location for the lease database, enter a URL in the format `tftp://ip-addr/ascii-string`. The maximum number of characters for the URL is 63.

---

option

> Add relay information option (Option 82) to DHCP client packets that are being forwarded out trusted ports. The default is `yes`, add relay information.

trust

> Configure trusted ports. Only server packets received on trusted ports are forwarded. Default: `untrusted`.

verify

> Enables DHCP packet validation. The DHCP client hardware address field and the source MAC address must be the same for packets received on untrusted ports or the packet is dropped. Default: `Yes`.

vlan

> Enable DHCP snooping on a vlan. DHCP snooping must be enabled already. Default: `No`.

To display the DHCP snooping configuration, enter this command:

```
HP Switch(config)# show dhcp-snooping
```

Output for the show dhcp-snooping command

```
HP Switch(config)# show dhcp-snooping
 DHCP Snooping Information
  DHCP Snooping              : Yes
  Enabled Vlans              :
  Verify MAC                 : Yes
  Option 82 untrusted policy : drop
  Option 82 Insertion        : Yes
  Option 82 remote-id        : mac
  Store lease database : Not configured
  Port Trust
  ----- -----
  5     No
  6     No
```

To display statistics about the DHCP snooping process, enter this command:

```
HP Switch(config)# show dhcp-snooping stats
```

An example of the output is shown below.

Output for the show DHCP snooping statistics command

```
HP Switch(config)# show dhcp-snooping stats

Packet type  Action   Reason                        Count
-----------  -------  ----------------------------  -----
server       forward  from trusted port             8
client       forward  to trusted port               8
server       drop     received on untrusted port    2
server       drop     unauthorized server           0
client       drop     destination on untrusted port 0
client       drop     untrusted option 82 field     0
client       drop     bad DHCP release request      0
client       drop     failed verify MAC check       0
```

Enabling DHCP snooping on VLANs

DHCP snooping on VLANs is disabled by default. To enable DHCP snooping on a VLAN or range of VLANs enter this command:

```
HP Switch(config)# dhcp-snooping vlan <vlan-id-range>
```

You can also use this command in the vlan context, in which case you cannot enter a range of VLANs for snooping. Below is an example of DHCP snooping enabled on VLAN 4.

DHCP snooping on a VLAN

```
HP Switch(config)# dhcp-snooping vlan 4
HP Switch(config)# show dhcp-snooping

 DHCP Snooping Information

  DHCP Snooping             : Yes
  Enabled Vlans             : 4
  Verify MAC                : Yes
  Option 82 untrusted policy : drop
  Option 82 Insertion       : Yes
  Option 82 remote-id       : mac
```

Clearing DHCP snooping table overview

Commands are available to dynamically learn DHCPv4 snooping bindings and clear DHCP snooping tables dynamic binding table entries.

These commands will:

- Allow the administrator to clear all entries in DHCP snooping binding tables.

- Allow the administrator to selectively clear DHCP snooping bindings by filtering entries by IP, VLAN, or port number.

Limitations

- When a dynamically learned binding is cleared, DHCPv4 snooping drops any further DHCP messages (`DHCP REQUEST`, `DHCP INFORM` and `DHCP RELEASE`) from that client.

- Clearing of a dynamically learned DHCPv4 snooping binding would not be synced to the DHCP server and the DHCP client.

  The administrator will receive a warning about these limitations before proceeding with the clear command execution.

  Execution of this command results in clearing of dynamically learnt DHCP Snooping entries from the binding table on this switch. Since the DHCP Server and the DHCP Clients would not be aware of this change, this can have side effects.

- The client will not be allowed to renew the IP address until the client sends a new `DHCP DISCOVER` packet. This process will allow the IP access to renew through to the server.

- The DHCP Server will consider the IP as assigned for the entire duration of the lease time even if the client attempts to release or renew the IP.

- If Dynamic IP Lockdown is enabled, the clearing of a dynamic DHCP snooping table binding will result in removal of the Dynamic IP Lockdown entry for that client. This action results in data packets being dropped until a new DHCP snooping table binding is added.

- If Dynamic Address Resolution Protocol Protection (DARPP) is enabled, clearing of a dynamic DHCP snooping table binding results in ARP packets being dropped from the client by DARPP until a new DHCP snooping table binding is added.

One or more dynamic bindings from DHCP snooping binding table may have to be removed in following cases:

- Change of the DHCPv4 Server.

- Change to the configuration of the server.

- Changes in the network topology causing clients to be moved to a different VLAN or port.

- Clearing existing binding entries of inactive clients.

Restrictions

Clearing DHCP snooping table is not available in these circumstances:

- Support is unavailable for clearing static bindings.

- Support is unavailable for DHCPv6 snooping.

clear dhcp-snooping binding

Syntax

```
clear dhcp-snooping binding [all | ip <IP-ADDR> | port <PORT-NUM> | vlan <VLAN-ID>]
```

Description

Clear all DHCP snooping binding entries or the binding entries on the specified IP address, port, or VLAN.

Parameters

binding

Clear all the DHCP snooping binding entries as specified by the options.

Options

all

Clear all the DHCP snooping binding entries.

IP *<IP-ADDR>*

Clear the DHCP snooping binding entry on the specified IP address.

port *<PORT-NUM>*

Clear the DHCP snooping binding entries on the specified port by entering the port number.

vlan *<VLAN-ID>*

Clear the DHCP snooping binding entries on the specified VLAN by entering the VLAN identifier.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Usage

```
clear dhcp-snooping binding all
clear dhcp-snooping binding ip <IP-ADDR>r


clear dhcp-snooping binding VLAN <VLAN-ID>r
```

Example

```
HP-Switch-5406Rzl2(config)# clear dhcp-snooping binding all

Warning: Execution of this command results in clearing of dynamically learnt
DHCP Snooping entries from the binding table on this switch. Since the
DHCP Server and the DHCP Clients would not be aware of this change, this can
have side effects.

Do you want to continue (y/n)? y

HP-Switch-5406Rzl2(config)# sh dhcp-snooping binding
  MacAddress     IP              VLAN Interface Time Left
  ------------- --------------- ---- --------- ---------


HP-Switch-5406Rzl2(config)# show log -r
 Keys:   W=Warning   I=Information
         M=Major     D=Debug E=Error
----  Reverse event Log listing: Events Since Boot  ----
W 12/12/16 14:27:03 05357 dhcp-snoop: AM1: All the dynamic binding entries were cleared.
I 12/12/16 14:23:24 03125 mgr: AM1: Startup configuration changed by CLI.  New seq. number 7
```

Example

```
HP-Switch-5406Rzl2(config)# clear dhcp-snooping binding all

Warning: Execution of this command results in clearing of dynamically learnt
DHCP Snooping entries from the binding table on this switch. Since the
DHCP Server and the DHCP Clients would not be aware of this change, this can
have side effects.

Do you want to continue (y/n)? y

HP-Switch-5406Rzl2(config)# sh dhcp-snooping binding

  MacAddress     IP              VLAN Interface Time Left
  ------------- --------------- ---- --------- ---------

HP-Switch-5406Rzl2(config)# show log -r
 Keys:   W=Warning   I=Information
         M=Major     D=Debug E=Error
----  Reverse event Log listing: Events Since Boot  ----
W 12/12/16 14:27:03 05357 dhcp-snoop: AM1: All the dynamic binding entries were cleared.
I 12/12/16 14:23:24 03125 mgr: AM1: Startup configuration changed by CLI. New seq. number 7

HP-Switch-5406Rzl2(config)# show dhcp-snooping

 DHCP Snooping Information

  DHCP Snooping             : Yes
  Enabled VLANs             : 20 40
  Verify MAC address        : Yes
  Option 82 untrusted policy : drop
  Option 82 insertion       : Yes
  Option 82 remote-id       : mac
  Store lease database      : Not configured

                Max      Current Bindings
```

```
  Port  Trust  Bindings  Static   Dynamic
  -----  -----  --------  ----------------
   A1    Yes      -         -          -
   A2    No       -         -          10

  Ports A3-A24,B1-B24,C1-C24,F1-F24 are untrusted
```
_____

```
HP-Switch-5406Rzl2(config)# show dhcp-snooping binding

  MacAddress       IP               VLAN Interface Time Left
  -------------  ---------------  ----  ---------  ---------
  000000-c97c83 192.168.101.21   20    A2         3579
  000000-c97c84 192.168.101.30   20    A2         3582
  000000-c97c85 192.168.101.28   20    A2         3579
  000000-c97c86 192.168.101.25   20    A2         3577
  000000-c97c87 192.168.101.23   20    A2         3574
  000000-c97c88 192.168.101.26   20    A2         3579
  000000-c97c89 192.168.101.24   20    A2         3580
  000000-c97c8a 192.168.101.27   20    A2         3582
  000000-c97c8b 192.168.101.29   20    A2         3577
  000000-c97c8c 192.168.101.22   20    A2         3580
```
_____

```
HP-Switch-5406Rzl2(config)# show dhcp-snooping binding

  MacAddress       IP               VLAN Interface Time Left
  -------------  ---------------  ----  ---------  ---------

HP-Switch-5406Rzl2(config)# show log -r
 Keys:    W=Warning    I=Information
          M=Major      D=Debug E=Error
---- Reverse event Log listing: Events Since Boot  ----
W 12/12/16 14:32:28 05359 dhcp-snoop: AM1: Dynamic binding entries on the port A2 were cleared.
W 12/12/16 14:29:58 05360 dhcp-snoop: AM1: Dynamic binding entries on the VLAN 20 were cleared.
W 12/12/16 14:29:36 05360 dhcp-snoop: AM1: Dynamic binding entries on the VLAN 40 were cleared.
W 12/12/16 14:28:29 00861 dhcp-snoop: AM1: backplane: Ceasing bad release logs for 5m
W 12/12/16 14:28:29 00860 dhcp-snoop: AM1: backplane: Attempt to release address
            192.168.101.30 leased to port A2 detected on port A2
W 12/12/16 14:28:29 00860 dhcp-snoop: AM1: backplane: Attempt to release address
            192.168.101.21 leased to port A2 detected on port A2
W 12/12/16 14:27:03 05357 dhcp-snoop: AM1: All the dynamic binding entries were cleared.
I 12/12/16 14:23:24 03125 mgr: AM1: Startup configuration changed by CLI. New seq. number 7
W 12/12/16 14:22:53 00853 dhcp-snoop: AM1: backplane: Ceasing untrusted port
            destination logs for 5m
W 12/12/16 14:22:53 00852 dhcp-snoop: AM1: backplane: Client packet destined to
          untrusted port dropped
```

clear dhcp-snooping statistics

Syntax

```
clear dhcp-snooping statistics
```

Description

Reset all DHCP snooping statistics.

Usage

```
clear dhcp-snooping statistics
```

Validation rules

Validation	Error/Warning/Prompt
IP address of a nonexisting dynamic binding.	`DHCP Snooping dynamic binding entry is not found for IP address: <IP-ADDR>.`
The interactive warning that will be display on execution of the command. User will have to enter y/n after this message to execute the command.	`Warning: Execution of this command results in clearing of dynamically learnt DHCP Snooping entries from the binding table on this switch. Since the DHCP Server and the DHCP Clients would not be aware of this change, this can have side effects.Do you want to continue (y/n)?`
When the show DHCP snooping binding command is executed and the clear DHCP snooping binding operation is still going on.	`Clear operation on DHCP Snooping dynamic binding entries is in progress; please try again later.`

Error Log

Event	Message
When dynamic DHCP snooping binding is cleared for the `IP Address <IP-ADDR>`	`Dynamic binding entry with the IP address <IP-ADDR> was cleared.`
When all dynamic DHCP snooping bindings are cleared on the `Port <PORT-NUM>`	`Dynamic binding entries on the port <PORT-NUM> were cleared.`
When all dynamic DHCP snooping bindings are cleared on VLAN `<VLAN-ID>`	`Dynamic binding entries on the VLAN <VLAN-ID> were cleared.`

RMON table

Error	Message	Description
RMON_DSNOOP_CLEAR_BIND_ALL		
	`W 05/06/16 07:15:46 02688 dhcp-snoop: All the dynamic binding entries were cleared.`	All of the dynamic DHCP snooping bindings are cleared.
RMON_DSNOOP_CLEAR_BIND_IP		
	`W 05/06/16 07:15:46 02688 dhcp-snoop: Dynamic binding entry for IP address <IP-ADDR> was cleared.`	When the dynamic DHCP snooping binding is cleared for the `IP Address <IP-ADDR>` , this message will appear.
RMON_DSNOOP_CLEAR_BIND_PORT		

Table Continued

Error	Message	Description
	`W 05/06/16 07:15:46 02688 dhcp-snoop: Dynamic binding entries on the port <PORT-NUM> were cleared.`	When all the dynamic DHCP snooping bindings are cleared on the `Port` *`<PORT-NUM>`* , this message will appear.
RMON_DSNOOP_CLEAR_BIND_VLAN		
	`W 05/06/16 07:15:46 02688 dhcp-snoop: Dynamic binding entries on the VLAN <VLAN-ID> were cleared.`	When all the dynamic DHCP snooping bindings are cleared on VLAN *`<VLAN-ID>`* , this message will appear.

Configuring DHCP snooping trusted ports

HPE Networking switches support DHCPv4 and DHCPv6 snooping. Configuring both versions helps protect your entire network by blocking unintended or rogue DHCPv4 and DHCPv6 servers. By default, all ports are untrusted. Once configured, DHCP server packets are forwarded only if received on a trusted port. DHCP server packets received on an untrusted port are dropped.

For DHCPv4 servers

To configure a port or range of ports as trusted, enter this command:

```
HP Switch(config)# dhcp-snooping trust <port-list>
```

You can also use this command in the interface context, in which case you are not able to enter a list of ports.

Setting trusted ports

```
HP Switch(config)# dhcp-snooping trust 5-6
HP Switch(config)# show dhcp-snooping

 DHCP Snooping Information

  DHCP Snooping              : Yes
  Enabled Vlans              : 4
  Verify MAC                 : Yes
  Option 82 untrusted policy : drop
  Option 82 Insertion        : Yes
  Option 82 remote-id        : mac

  Store lease database : Not configured

  Port  Trust
  ----- -----
  5     Yes
  6     Yes
  7     No
```

Use the `no` form of the command to remove the trusted configuration from a port.

For DHCPv6 servers

To configure a port or range of ports as trusted, enter this command:

```
HP Switch(config)# dhcpv6-snooping trust <port-list>
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

You can also use this command in the interface context, in which case you are not able to enter a list of ports.

Use the `no` form of the command to remove the trusted configuration from a port.

Configuring authorized server addresses

If authorized server addresses are configured, a packet from a DHCP server must be received on a trusted port AND have a source address in the authorized server list in order to be considered valid. If no authorized servers are configured, all servers are considered valid. You can configure a maximum of 20 authorized servers.

To configure a DHCP authorized server address, enter this command in the global configuration context:

```
HP Switch(config)# dhcp-snooping authorized-server <ip-address>
```

Authorized servers for DHCP snooping

```
HP Switch(config)# show dhcp-snooping

 DHCP Snooping Information

  DHCP Snooping              : Yes
  Verify MAC                 : No
  Option 82 untrusted policy : drop
  Option 82 Insertion        : Yes
  Option 82 remote-id        : subnet-ip

Authorized Servers
--------------------
111.222.3.4
10.0.0.11
```

Using DHCP snooping with option 82

DHCP adds Option 82 (relay information option) to DHCP request packets received on untrusted ports by default. (See "Configuring DHCP Relay" in the management and configuration guide for more information on Option 82.)

When DHCP is enabled globally and also enabled on a VLAN, and the switch is acting as a DHCP relay, the settings for the DHCP relay Option 82 command are ignored when snooping is controlling Option 82 insertion. Option 82 inserted in this manner allows the association of the client's lease with the correct port, even when another device is acting as a DHCP relay or when the server is on the same subnet as the client.

> **NOTE**
> DHCP snooping only overrides the Option 82 settings on a VLAN that has snooping enabled, not on VLANs without snooping enabled.

If DHCP snooping is enabled on a switch where an edge switch is also using DHCP snooping, it is desirable to have the packets forwarded so the DHCP bindings are learned. To configure the policy for DHCP packets from untrusted ports that already have Option 82 present, enter this command in the global configuration context.

Syntax:

```
[no] dhcp-snooping option 82 [remote-id <mac|subnet-ip|mgmt-ip>][untrusted-policy <drop|keep|replace>]
```

Enables DHCP Option 82 insertion in the packet

remote-id

Set the value used for the `remote-id` field of the relay information option.

mac

The switch mac address is used for the remote-id. This is the default.

subnet-ip

The IP address of the VLAN the packet was received on is used for the remote-id. If `subnet-ip` is specified but the value is not set, the MAC address is used.

mgmt-ip

The management VLAN IP address is used as the remote-id. If mgmt-ip is specified but the value is not set, the MAC address is used.

untrusted-policy

Configures DHCP snooping behavior when forwarding a DHCP packet from an untrusted port that already contains DHCP relay information (Option 82). The default is `drop`.

drop

The packet is dropped.

keep

The packet is forwarded without replacing the option information.

replace

The existing option is replaced with a new Option 82 generated by the switch.

> **NOTE**
>
> The default `drop` policy should remain in effect if there are any untrusted nodes, such as clients, directly connected to this switch.

Changing the remote-id from a MAC to an IP address

By default, DHCP snooping uses the MAC address of the switch as the remoteid in Option 82 additions. The IP address of the VLAN the packet was received on or the IP address of the management VLAN can be used instead by entering this command with the associated parameter:

```
HP Switch(config)# dhcp-snooping option 82 remote-id <mac|subnet-ip|mgmt-ip>
```

DHCP snooping option 82 using the VLAN IP address

```
HP Switch(config)# dhcp-snooping option 82 remote-id subnetip
HP Switch(config)# show dhcp-snooping

 DHCP Snooping Information

  DHCP Snooping              : Yes
  Enabled Vlans              : 4
  Verify MAC                 : Yes
  Option 82 untrusted policy : drop
  Option 82 Insertion        : Yes
  Option 82 remote-id        : subnet-ip
```

Disabling the MAC address check

DHCP snooping drops DHCP packets received on untrusted ports when the check address (chaddr) field in the DHCP header does not match the source MAC address of the packet (default behavior). To disable this checking, use the `no` form of this command.

```
HP Switch(config)# dhcp-snooping verify mac
```

Showing the DHCP snooping verify MAC setting

```
HP Switch(config)# dhcp-snooping verify mac
HP Switch(config)# show dhcp-snooping

 DHCP Snooping Information

  DHCP Snooping              : Yes
  Enabled Vlans              : 4
  Verify MAC                 : yes
  Option 82 untrusted policy : drop
  Option 82 Insertion        : Yes
  Option 82 remote-id        : subnet-ip
```

DHCP binding database

DHCP snooping maintains a database of up to 8192 DHCP bindings on untrusted ports. Each binding consists of:

- Client MAC address

- Port number

- VLAN identifier

- Leased IP address

- Lease time

The switch can be configured to store the bindings at a specific URL so they will not be lost if the switch is rebooted. If the switch is rebooted, it will read its binding database from the specified location. To configure this location use this command.

Syntax:

```
[no] dhcp-snooping database [file <tftp://<ip-address>/<ascii-string>>][delay <15-86400>][timeout <0-86400>]
```

file

Must be in Uniform Resource Locator (URL) format – "tftp://ip-address/ascii-string". The maximum filename length is 63 characters.

delay

Number of seconds to wait before writing to the database. Default = 300 seconds.

timeout

Number of seconds to wait for the database file transfer to finish before returning an error. A value of zero (0) means retry indefinitely. Default = 300 seconds.

A message is logged in the system event log if the DHCP binding database fails to update. To display the contents of the DHCP snooping binding database, enter this command.

Syntax:

```
show dhcp-snooping binding
```

DHCP snooping binding database contents

```
HP Switch(config)# show dhcp-snooping binding

 MacAddress          IP              VLAN Interface Time left
```

```
-------------   --------------- ---- --------- ---------
22.22.22.22.22.22  10.0.0.1         4    6         1600
```

> **NOTE**
> If a lease database is configured, the switch drops all DHCP packets until the lease database is read. This only occurs when the switch reboots and is completed quickly. If the switch is unable to read the lease database from the tftp server, it waits until that operation times out and then begins forwarding DHCP packets.

DHCPv4 snooping max-binding

DHCP snooping max-binding prevents binding entries from getting exhausted. This feature is on a per-port basis. It restricts the maximum number of bindings allowed on a port/interface. It applies to untrusted interfaces only. The maximum bindings for a particular port includes both statically configured and dynamically learned. The number of bindings on a per port basis is maintained i.e., incremented upon a lease offer and decremented upon a lease expiry or release.

DHCP snooping max-binding can be configured in configuration context or in an interface context for an untrusted interface. In case of configuration context, a port or a list of ports is selected for which max-binding is to be configured. Then the corresponding max-binding value is provided within a range of <1-8192>. For the interface context, after selecting the interface on which max-binding is to be configured, the max-binding value is provided within a range of <1-8192>. The max-binding configuration for a port can be removed using the no option of the command. max-binding cannot be set on trusted ports and ports for which the associated VLAN is not DHCP-snooping enabled. Once the max-bindings limit on an interface is reached, packets for DHCP clients which do not have a binding entry are dropped.

Syntax:

```
(config)# dhcp-snooping max-bindings [PORT-LIST][MAX-BINDING-NUM]
```

Configure the maximum number of bindings on specified ports. The maximum number of bindings default value is 8192. The allowed range on a port is 1 to 8192.

Syntax:

```
(interface)# dhcp-snooping <trust|max-bindings>[1-8192]
```

Configures the maximum binding value on a port. Only this number of clients are allowed on a port. By specifying [no] the max-binding is removed from the configuration and set to the default value of 8192.

Syntax:

```
(config)# show dhcp-snooping
```

Show all available dhcp-snooping information.

Example:

```
DHCP Snooping Information
DHCP Snooping : Yes
                Max         Current  Bindings
Port    Trust   Bindings    Static   Dynamic

_____   _____   _____    _____  _____
 1      Yes        -           -        -
 2      No        200          10       3
 3      No         3*          3        6
 4      No         5*          23       0
 5      No         -           -        -
 6      No         -           -        -
 7      No         -           -        -
```

```
 8      No          -         -        -
 9      No          -         -        -
10      No          -         -        -
11      Yes         -         -        -
12      Yes         -         -        -
13      No          -         -        -
14      No          -         -        -
15      No          -         -        -
16      No          -         2        8
17      No          21        12       24
18      Yes         -         -        -
19      No          -         -        -
20      No          -         -        -
21      No          -         -        -
22      No          -         -        -
23      No          -         -        -
24      Yes         -         -        -
```

Syntax:

```
(config)# show dhcp-snooping stats
```

Shows the dhcp-snooping statistics.

```
Packet type  Action   Reason                         Count
-----------  -------  -----------------------------  ---------
server       forward  from trusted port              0
client       forward  to trusted port                0
server       drop     received on untrusted port     0
server       drop     unauthorized server            0
client       drop     destination on untrusted port  0
client       drop     untrusted option 82 field      0
client       drop     bad DHCP release request       0
client       drop     failed verify MAC check        0
client       drop     failed on max-binding limit    0
```

Enabling debug logging

To enable debug logging for DHCP snooping, use this command.

Syntax:

```
[no] debug security dhcp-snooping [agent|event|packet]
```

agent

 Displays DHCP snooping agent messages.

event

 Displays DHCP snooping event messages.

packet

 Displays DHCP snooping packet messages.

DHCP operational notes

• DHCP is not configurable from the WebAgent or menu interface.

• If packets are received at too high a rate, some may be dropped and need to be re-transmitted.

- Hewlett Packard Enterprise recommends running a time synchronization protocol such as SNTP in order to track lease times accurately.

- A remote server must be used to save lease information or there may be a loss of connectivity after a switch reboot.

Log messages

Server <ip-address> packet received on untrusted port <port-number> dropped.

Indicates a DHCP server on an untrusted port is attempting to transmit a packet. This event is recognized by the reception of a DHCP server packet on a port that is configured as untrusted.

Ceasing untrusted server logs for %s.

More than one packet was received from a DHCP server on an untrusted port. To avoid filling the log file with repeated attempts, untrusted server drop packet events will not be logged for the specified <duration>.

Client packet destined to untrusted port <port-number> dropped.

Indicates that the destination of a DHCP client unicast packet is on an untrusted port. This event is recognized when a client unicast packet is dropped because the destination address is out a port configured as untrusted.

Ceasing untrusted port destination logs for %s.

More that one client unicast packet with an untrusted port destination was dropped. To avoid filling the log file with repeated attempts, untrusted port destination attempts will not be logged for the specified <duration>.

Unauthorized server <ip-address> detected on port <port-number>.

Indicates that an unauthorized DHCP server is attempting to send packets. This event is recognized when a server packet is dropped because there are configured authorized servers and a server packet is received from a server that is not configured as an authorized server.

Ceasing unauthorized server logs for <duration>.

More than one unauthorized server packet was dropped. To avoid filling the log file with repeated attempts, unauthorized server transmit attempts will not be logged for the specified <duration>.

Received untrusted relay information from client <mac-address> on port <port-number>.

Indicates the reception on an untrusted port of a client packet containing a relay information option field. This event is recognized when a client packet containing a relay information option field is dropped because it was received on a port configured as untrusted.

Ceasing untrusted relay information logs for <duration>.

More than one DHCP client packet received on an untrusted port with a relay information field was dropped. To avoid filling the log file with repeated attempts, untrusted relay information packets will not be logged for the specified <duration>.

Client address <mac-address> not equal to source MAC <mac-address> detected on port <port-number>.

Indicates that a client packet source MAC address does not match the "chaddr" field. This event is recognized when the dhcp-snooping agent is enabled to filter DHCP client packets that do not have a matching "chaddr" field and source MAC address.

Ceasing MAC mismatch logs for <duration>.

More than one DHCP client packet with a mismatched source MAC and chaddr field was dropped. To avoid filling the log file with repeated attempts, client address mismatch events will not be logged for the specified <duration>.

Attempt to release address <ip-address> leased to port <port-number> detected on port <port-number> dropped.

Indicates an attempt by a client to release an address when a DHCPRELEASE or DHCPDECLINE packet is received on a port different from the port the address was leased to.

Ceasing bad release logs for %s.

More than one bad DHCP client release packet was dropped. To avoid filling the log file with repeated bad release dropped packets, bad releases will not be logged for <duration>.

Lease table is full, DHCP lease was not added.

The lease table is full and this lease will not be added to it.

Write database to remote file failed errno (error-num).

An error occurred while writing the temporary file and sending it using tftp to the remote server.

DHCP packets being rate-limited.

Too many DHCP packets are flowing through the switch and some are being dropped.

Snooping table is full.

The DHCP binding table is full and subsequent bindings are being dropped.

IPv6 Network Defense

DSNOOPv6 and DIPLDv6

DSNOOPv6 enables network defenses for IPv6 on switches. It provides protection against network disruption by blocking unintended/rogue DHCPv6 Servers.

DSNOOPv6, when used with Dynamic IP Lockdown (DIPLD), provides network defense against source address spoofing. For example, a wireless access point with a DHCP Server running by default hands out IP addresses to wired clients that fall under a different subnet. DHCPv6 Snooping (DSNOOP) helps protect a network from unintended/rogue DHCP Servers handing out IP address leases to hosts on the network.

In an IPv6 network, addresses are predominately assigned via Router Advertisements. However, RA is limited in its ability to provide all of the network configurations to hosts. By managing their networks with DHCP(v4/v6) Servers, administrators can increase their network range and security. Since customer networks have both IPv4 and IPv6 configurations, enabling the DHCPv6-Snooping feature provides an additional network defense level of protection.

DSNOOPv6 operates similarly to DSNOOPv4. To decide which switch ports DHCPv6 packets are accepted from and forwarded to, packets are intercepted, examined and validated on DHCPv6 protocol fields. The Client IP address binding information is maintained by the switch in a binding table.

> **NOTE**
> The DIPLDv6 limits will be different on different switch platforms due to hardware limitations.

Configuring DHCPv6 snooping

> **NOTE**
> DHCPv6 snooping is currently configurable through SNMP using MIBs. For more information, see the *MIB and Trap Matrix*.

Enabling DHCPv6 snooping

To globally enable DHCPv6 snooping, enter:

```
HP Switch(config)# dhcpv6-snooping
```

Use the `no` form of the command to disable DHCPv6 snooping.

Enabling DHCPv6 snooping on VLANs

After you globally enable DHCPv6, use this command to enable DHCPv6 snooping on a VLAN or range of VLANs.

Syntax

```
[no] dhcpv6-snooping <vlan-id-range>
```

Use the `no` form of the command to disable DHCPv6 snooping on a VLAN.

vlan-id-range

Specifies the VLAN or range of VLANs on which to enable DHCPv6 snooping.

Configuring an authorized DHCPv6 server for snooping

Use this command to configure an authorized DHCPv6 server.

Syntax

```
dhcpv6-snooping authorized-server <IPv6-address>
```

IPv6-address

Specifies the IP address of a trusted DHCP server.

If no authorized servers are configured, all DHCP server addresses are considered valid. Maximum: 20 authorized servers.

Configuring a lease entry file for DHCPv6 snooping

Use this command to configure lease database transfer options for DHCPv6 snooping

Syntax

```
[no] dhcpv6-snooping database [file <ASCII string>] [delay <15-86400>] [timeout <0-86400>]
```

file <ASCII string>

Specifies the database URL in the form: "tftp://<IP-ADDR>/<FILENAME>" with a maximum length of 255 characters, IP-ADDR can be an IPv4 or an IPv6 address. IPv6 addresses must be enclosed in square brackets.

delay <15-86400>

Specifies the seconds to delay before writing to the lease database file. Valid values are 15 to -86400. Default is 300 seconds.

timeout <0-86400>

Specifies the seconds to wait for the lease file transfer to finish before a failure message is displayed. Valid values are 0 to 86400. Default is 300 seconds. If 0 is specified, the file transfer is retried indefinitely.

Configuring DHCPv6 snooping max binding

Use this command to configure the maximum number of binding addresses allowed per port. . If you configure the max-bindings value before enabling DHCPv6-snooping, the limit you enter is immediately applied, and the bindings are not allowed to exceed the max-bindings value. If you set the max-bindings value after enabling DHCPv6-snooping, the following occurs:

- If current bindings are greater than the max-binding value, the configuration is applied when clients release their IPv6 addresses.

- If current bindings are lesser than that of the max-binding value, the configuration is immediately applied.

Syntax

```
[no] dhcpv6-snooping max-bindings <port-list> <1-8192>
```

port-list

Specifies the ports on which to apply max-bindings.

1-8192

Specifies the maximum number of binding addresses.

Configuring traps for DHCPv6 snooping

Use this command to configure traps for DHCPv6 snooping.

Syntax

```
[no] snmp-server enable traps dhcpv6-snooping [[out-of-resources] | [errant-reply]]
```

out-of-resources

Sends a trap message when the number of bindings exceeds the maximum limit of 8192 bindings.

errant-reply

Sends a trap message when a DHCPv6 reply packet is received on an untrusted port or from an unauthorized server.

Clearing DHCPv6 snooping statistics

Use this command in switch config mode to clear DHCPv6 snooping statistics.

Syntax

```
clear dhcpv6-snooping statistics
```

Enabling debug logging for DHCPv6 snooping

To enable debug logging for DHCPv6 snooping, use this command.

Syntax

```
[no] debug security dhcpv6-snooping [config|event|packet]
```

config

Displays DHCPv6 snooping configuration messages.

event

Displays DHCPv6 snooping event messages.

packet

Displays DHCPv6 snooping packet messages.

DHCPv6 show commands

Use this command to show DHCPv6 snooping information.

Syntax

```
show dhcpv6-snooping [stats] [bindings]
```

stats

Shows DHCPv6 snooping statistics.

bindings

Shows DHCPv6 binding state entries in a tabular format.

Examples

The following example shows all available DHCPv6 snooping information.

```
HP Switch(config)# show dhcpv6 snooping
DHCP Snooping Information
DHCP Snooping        : Yes
Enabled VLANs        : 1 13 16
Remote-ID            : MAC
Store Lease Database : Yes
URL                  : tftp://120.93.49.9/avi
Read at boot         : no
Write Delay          : 300
Write Timeout        : 300
File Status          : up-to-date
Write Attempts       : 0
Write Failures       : 0
Last Successful File Update

                Max         Current  Bindings
Port    Trust   Bindings    Static   Dynamic

_____   _____   _____    _____  _____
  1     Yes        -           -         -
  2     No         20          20        3
  4     No         3*          3         6
  4     No         543         231       10
 13     No         -           3         6
 48     Yes        -           -         -

Ports 3,5-12,14-47 are untrusted.
Note that show commands list only those ports that have bindings on them.
Ports 3, 5, 6,8 are untrusted as they are not listed in table and they do
not have associated bindings.
```

The following example shows DHCPv6 snooping statistics.

```
HP Switch(config)# show dhcpv6 snooping stats

Packet Type    Action    Reason                           Count

_____   _____   _____                          _____
server         forward   from trusted port                0
client         forward   to trusted port                  0
server         drop      received on validating port      0
server         drop      unauthorized server              0
client         drop      destination on validating port   0
client         drop      relay reply on validating port   0
client         drop      bad DHCPv6 release request        0
```

```
client        drop        failed verify MAC check        0
client        drop        failed on max-binding limit    0
```

Dynamic ARP protection

On the VLAN interfaces of a routing switch, dynamic ARP protection ensures that only valid ARP requests and responses are relayed or used to update the local ARP cache. ARP packets with invalid IP-to-MAC address bindings advertised in the source protocol address and source physical address fields are discarded. For more information about the ARP cache, see "ARP Cache Table" in the multicast and routing guide.

ARP requests are ordinarily broadcast and received by all devices in a broadcast domain. Most ARP devices update their IP-to-MAC address entries each time they receive an ARP packet even if they did not request the information. This behavior makes an ARP cache vulnerable to attacks.

Because ARP allows a node to update its cache entries on other systems by broadcasting or unicasting a gratuitous ARP reply, an attacker can send his own IP-to-MAC address binding in the reply that causes all traffic destined for a VLAN node to be sent to the attacker's MAC address. As a result, the attacker can intercept traffic for other hosts in a classic "man-in-the-middle" attack. The attacker gains access to any traffic sent to the poisoned address and can capture passwords, e-mail, and VoIP calls or even modify traffic before resending it.

Another way in which the ARP cache of known IP addresses and associated MAC addresses can be poisoned is through unsolicited ARP responses. For example, an attacker can associate the IP address of the network gateway with the MAC address of a network node. In this way, all outgoing traffic is prevented from leaving the network because the node does not have access to outside networks. As a result, the node is overwhelmed by outgoing traffic destined to another network.

Dynamic ARP protection is designed to protect your network against ARP poisoning attacks in the following ways:

- Allows you to differentiate between trusted and untrusted ports.

- Intercepts all ARP requests and responses on untrusted ports before forwarding them.

- Verifies IP-to-MAC address bindings on untrusted ports with the information stored in the lease database maintained by DHCP snooping and user-configured static bindings (in non-DHCP environments):

  ◦ If a binding is valid, the switch updates its local ARP cache and forwards the packet.

  ◦ If a binding is invalid, the switch drops the packet, preventing other network devices from receiving the invalid IP-to-MAC information.

DHCP snooping intercepts and examines DHCP packets received on switch ports before forwarding the packets. DHCP packets are checked against a database of DHCP binding information. Each binding consists of a client MAC address, port number, VLAN identifier, leased IP address, and lease time. The DHCP binding database is used to validate packets by other security features on the switch.

If you have already enabled DHCP snooping on a switch, you may also want to add static IP-to-MAC address bindings to the DHCP snooping database so that ARP packets from devices that have been assigned static IP addresses are also verified.

Supports additional checks to verify source MAC address, destination MAC address, and IP address.ARP packets that contain invalid IP addresses or MAC addresses in their body that do not match the addresses in the Ethernet header are dropped.

When dynamic ARP protection is enabled, only ARP request and reply packets with valid IP-to-MAC address bindings in their packet header are relayed and used to update the ARP cache.

Dynamic ARP protection is implemented in the following ways on a switch:

- You can configure dynamic ARP protection only from the CLI; you cannot configure this feature from the WebAgent or menu interfaces.

- Line rate—Dynamic ARP protection copies ARP packets to the switch CPU, evaluates the packets, and then re-forwards them through the switch software. During this process, if ARP packets are received at too high a line rate, some ARP packets may be dropped and will need to be retransmitted.

- The SNMP MIB, HP-ICF-ARP-PROTECT-MIB, is created to configure dynamic ARP protection and to report ARP packet-forwarding status and counters.

Enabling dynamic ARP protection

To enable dynamic ARP protection for VLAN traffic on a routing switch, enter the `arp-protect vlan` command at the global configuration level.

Syntax:

`[no] arp-protect vlan [`*`vlan-range`*`]`

`vlan-range`: Specifies a VLAN ID or a range of VLAN IDs from one to 4094; for example, 1–200.

An example of the `arp-protect vlan` command is shown here:

`HP Switch(config)# arp-protect vlan 1-101`

Configuring trusted ports

In a similar way to DHCP snooping, dynamic ARP protection allows you to configure VLAN interfaces in two categories: trusted and untrusted ports. ARP packets received on trusted ports are forwarded without validation.

By default, all ports on a switch are untrusted. If a VLAN interface is untrusted:

- The switch intercepts all ARP requests and responses on the port.

- Each intercepted packet is checked to see if its IP-to-MAC binding is valid. If a binding is invalid, the switch drops the packet.

Configure trusted ports carefully. For example, in the topology in **Figure 110: Trusted ports for dynamic ARP protection** on page 481, Switch B may not see the leased IP address that Host 1 receives from the DHCP server. If the port on Switch B that is connected to Switch A is untrusted and if Switch B has dynamic ARP protection enabled, it will see ARP packets from Host 1 as invalid, resulting in a loss of connectivity.

In contrast, if Switch A does not support dynamic ARP protection and you configure the port on Switch B connected to Switch A as trusted, Switch B opens itself to possible ARP poisoning from hosts attached to Switch A.

Figure 110: *Trusted ports for dynamic ARP protection*

Consider the following configuration guidelines when you use dynamic ARP protection in your network:

• Configure ports connected to other switches in the network as trusted ports. In this way, all network switches can exchange ARP packets and update their ARP caches with valid information.

• Switches that do not support dynamic ARP protection must be separated by a router in their own Layer 2 domain. Because ARP packets do not cross Layer 3 domains, the unprotected switches cannot unknowingly accept ARP packets from an attacker and forward them to protected switches through trusted ports.

To configure one or more Ethernet interfaces that handle VLAN traffic as trusted ports, enter the `arp-protect trust` command at the global configuration level. The switch does not check ARP requests and responses received on a trusted port.

Syntax:

```
[no] arp-protect trust <port-list>
```

`port-list`: Specifies a port number or a range of port numbers. Separate individual port numbers or ranges of port numbers with a comma; for example: 13-15, 17.

An example of the `arp-protect trust` command is shown here:

```
HP Switch(config)# arp-protect trust 5-8, 17
```

Adding an IP-to-MAC binding to the DHCP database

A routing switch maintains a DHCP binding database, which is used for DHCP and ARP packet validation. Both the DHCP snooping and DHCP Option 82 insertion features maintain the lease database by learning the IP-to-MAC bindings on untrusted ports. Each binding consists of the client MAC address, port number, VLAN identifier, leased IP address, and lease time.

If your network does not use DHCP or if some network devices have fixed, user-configured IP addresses, you can enter static IP-to-MAC address bindings in the DHCP binding database. The switch uses manually configured static bindings for DHCP snooping and dynamic ARP protection.

Clearing the DHCP snooping binding table

To remove the IP-to-MAC binding from the database, use the `no` form of the `ip source-binding` command.

Adding a static binding

To add the static configuration of an IP-to-MAC binding for a port to the database, enter the `ip source-binding` command at the global configuration level. Use the `no`form of the command to remove the IP-to-MAC binding from the database.

Syntax:

```
[no] ip source-binding <mac-address> vlan <vlan-id> <ip-address> interface <port-number>
```

<mac-address>

> Specifies a MAC address to bind with a VLAN and IP address on the specified port in the DHCP binding database.

vlan <vlan-id>

> Specifies a VLAN ID number to bind with the specified MAC and IP addresses on the specified port in the DHCP binding database.

<ip-address>

> Specifies an IP address to bind with a VLAN and MAC address on the specified port in the DHCP binding database.

interface <port-number >

> Specifies the port number on which the IP-to- MAC address and VLAN binding is configured in the DHCP binding database.

An example of the `ip source-binding` command is shown here:

```
HP Switch(config)# ip source-binding 0030c1-7f49c0
interface vlan 100 10.10.20.1 interface A4
```

> **NOTE**
>
> The `ip source-binding` command is the same command used by the Dynamic IP Lockdown feature to configure static bindings. The Dynamic ARP Protection and Dynamic IP Lockdown features share a common list of source IP-to-MAC bindings.

Configuring additional validation checks on ARP packets

Dynamic ARP protection can be configured to perform additional validation checks on ARP packets. By default, no additional checks are performed. To configure additional validation checks, enter the `arp-protect validate` command at the global configuration level.

Syntax:

```
[no] arp-protect validate <[src-mac]|[dest-mac]|[ip]>
```

src-mac

> (Optional) Drops any ARP request or response packet in which the source MAC address in the Ethernet header does not match the sender MAC address in the body of the ARP packet.

dest-mac

> (Optional) Drops any unicast ARP response packet in which the destination MAC address in the Ethernet header does not match the target MAC address in the body of the ARP packet.

ip

(Optional) Drops any ARP packet in which the sender IP address is invalid. Drops any ARP response packet in which the target IP address is invalid. Invalid IP addresses include: 0.0.0.0, 255.255.255.255, all IP multicast addresses, and all Class E IP addresses.

You can configure one or more of the validation checks. The following example of the `arp-protect validate` command shows how to configure the validation checks for source MAC address and destination AMC address:

```
HP Switch(config)# arp-protect validate src-mac dest-mac
```

Verifying the configuration of dynamic ARP protection

To display the current configuration of dynamic ARP protection, including the additional validation checks and the trusted ports that are configured, enter the `show arp-protect` command:

Output for the show arp-protect command

```
HP Switch(config)# show arp-protect

 ARP Protection Information

 Enabled Vlans : 1-4094
 Validate : dest-mac, src-mac

 Port  Trust
 ----- -----
 5     Yes
 6     Yes
 7     No
 8     No
 9     No
```

Displaying ARP packet statistics

To display statistics about forwarded ARP packets, dropped ARP packets, MAC validation failure, and IP validation failures, enter the `show arp-protect statistics <vid-range>` command:

Output for the show arp-protect statistics command

```
HP Switch(config)# show arp-protect statistics 1-2

Status and Counters - ARP Protection Counters for VLAN 1

Forwarded pkts : 10       Bad source mac     : 2
Bad bindings   : 1        Bad destination mac: 1
Malformed pkts : 0        Bad IP address     : 0

Status and Counters - ARP Protection Counters for VLAN 2

Forwarded pkts : 1        Bad source mac     : 1
Bad bindings   : 1        Bad destination mac: 1
Malformed pkts : 1        Bad IP address     : 1
```

Monitoring dynamic ARP protection

When dynamic ARP protection is enabled, you can monitor and troubleshoot the validation of ARP packets with the `debug arp-protect` command. Use this command when you want to debug the following conditions:

- The switch is dropping valid ARP packets that should be allowed.

- The switch is allowing invalid ARP packets that should be dropped.

Output for the debug arp-protect command

```
HP Switch(config)# debug arp-protect

1. ARP request is valid
"DARPP: Allow ARP request 000000-000001,10.0.0.1 for 10.0.0.2 port 1,
vlan "

2. ARP request detected with an invalid binding
"DARPP: Deny ARP request 000000-000003,10.0.0.1 port 1, vlan 1"

3. ARP response with a valid binding
"DARPP: Allow ARP reply 000000-000002,10.0.0.2 port 2, vlan 1"

4. ARP response detected with an invalid binding
"DARPP: Deny ARP reply 000000-000003,10.0.0.2 port 2, vlan 1"
```

Dynamic IP lockdown

The Dynamic IP lockdown feature is used to prevent IP source address spoofing on a per-port and per-VLAN basis. When dynamic IP lockdown is enabled, IP packets in VLAN traffic received on a port are forwarded only if they contain a known source IP address and MAC address binding for the port. The IP-to-MAC address binding can either be statically configured or learned by the DHCP Snooping feature.

Protection against IP source address spoofing

Many network attacks occur when an attacker injects packets with forged IP source addresses into the network. Also, some network services use the IP source address as a component in their authentication schemes. For example, the BSD "r" protocols (rlogin, rcp, rsh) rely on the IP source address for packet authentication. SNMPv1 and SNMPv2c also frequently use authorized IP address lists to limit management access. An attacker that is able to send traffic that appears to originate from an authorized IP source address may gain access to network services for which he is not authorized.

Dynamic IP lockdown provides protection against IP source address spoofing by means of IP-level port security. IP packets received on a port enabled for dynamic IP lockdown are only forwarded if they contain a known IP source address and MAC address binding for the port.

Dynamic IP lockdown uses information collected in the DHCP Snooping lease database and through statically configured IP source bindings to create internal, per-port lists. The internal lists are dynamically created from known IP-to-MAC address bindings to filter VLAN traffic on both the source IP address and source MAC address.

Prerequisite: DHCP snooping

Dynamic IP lockdown requires that you enable DHCP snooping as a prerequisite for its operation on ports and VLAN traffic:

- Dynamic IP lockdown only enables traffic for clients whose leased IP addresses are already stored in the lease database created by DHCP snooping or added through a static configuration of an IP-to-MAC binding.Therefore, if you enable DHCP snooping after dynamic IP lockdown is enabled, clients with an existing DHCP-assigned address must either request a new leased IP address or renew their existing DHCP-assigned

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

address. Otherwise, a client's leased IP address is not contained in the DHCP binding database. As a result, dynamic IP lockdown will not allow inbound traffic from the client.

- It is recommended that you enable DHCP snooping a week before you enable dynamic IP lockdown to allow the DHCP binding database to learn clients' leased IP addresses. You must also ensure that the lease time for the information in the DHCP binding database lasts more than a week. Alternatively, you can configure a DHCP server to re-allocate IP addresses to DHCP clients. In this way, you repopulate the lease database with current IP-to-MAC bindings.

- The DHCP binding database allows VLANs enabled for DHCP snooping to be known on ports configured for dynamic IP lockdown. As new IP-to-MAC address and VLAN bindings are learned, a corresponding permit rule is dynamically created and applied to the port (preceding the final deny any vlan <VLAN_IDs> rule. These VLAN_IDs correspond to the subset of configured and enabled VLANs for which DHCP snooping has been configured.

- For dynamic IP lockdown to work, a port must be a member of at least one VLAN that has DHCP snooping enabled.

- Disabling DHCP snooping on a VLAN causes Dynamic IP bindings on Dynamic IP Lockdown-enabled ports in this VLAN to be removed. The port reverts back to switching traffic as usual.

Filtering IP and MAC addresses per-port and per-VLAN

This section contains an example that shows the following aspects of the Dynamic IP Lockdown feature:

- Internal Dynamic IP lockdown bindings dynamically applied on a per-port basis from information in the DHCP Snooping lease database and statically configured IP-to-MAC address bindings

- Packet filtering using source IP address, source MAC address, and source VLAN as criteria.

In this example, the following DHCP leases have been learned by DHCP snooping on port 5. VLANs 2 and 5 are enabled for DHCP snooping.

Table 36: *Sample DHCP snooping entries*

IP Address	MAC Address	VLAN ID
10.0.8.5	001122–334455	2
10.0.8.7	001122–334477	2
10.0.10.3	001122–334433	5

The following example shows an IP-to-MAC address and VLAN binding that have been statically configured in the lease database on port 5.

IP Address	MAC Address	VLAN ID
10.0.10.1	001122–110011	5

Assuming that DHCP snooping is enabled and that port 5 is untrusted, dynamic IP lockdown applies the following dynamic VLAN filtering on port 5:

Internal statements used by dynamic IP lockdown

```
permit 10.0.8.5 001122-334455 vlan 2

permit 10.0.8.7 001122-334477 vlan 2

permit 10.0.10.3 001122-334433 vlan 5

permit 10.0.10.1 001122-110011 vlan 5

deny any vlan 1-10

permit any
```

> **NOTE**
> The `deny any` statement is applied only to VLANs for which DHCP snooping is enabled. The `permit any` statement is applied only to all other VLANs.

Enabling Dynamic IP Lockdown

> **NOTE**
> Dynamic IPv6 Lockdown (DIPLDv6) is currently configurable through SNMP using MIBs. For more information, see the MIB and Trap Matrix.

IPv4

To enable dynamic IP lockdown on all ports or specified ports, enter this command at the global configuration level.

Syntax

```
[no] ip source-lockdown <port-list>
```

port-list

Specifies one or more ports on which to enable IP source lockdown.

Use the `no` form of the command to disable dynamic IP lockdown.

IPv6
Enabling dynamic IPv6 source lockdown

To enable dynamic IPv6 lockdown on all ports or specified ports, enter this command at the global configuration level.

Syntax

```
[no] ipv6 source-lockdown <port-list>
```

port-list

Specifies one or more ports on which to enable IP source lockdown.

Use the `no` form of the command to disable dynamic IP lockdown.

Enabling traps for dynamic IPv6 source lockdown

Use this command to configure traps for IPv6 source lockdown.

Syntax

```
[no] snmp-server enable traps dyn-ipv6-lockdown [[out-of-resources] | [violations]]
```

out-of-resources

Sends a trap message when resources are unavailable for configuring dynamic IPv6 source lockdown.

violations

Sends a trap message when a source lockdown violation occurs.

Enabling debug logging for dynamic IPv6 source lockdown

Syntax

```
[no] debug dynamic-ipv6-lockdown [config|event|packet]
```

config

Displays dynamic lockdown configuration messages.

event

Displays dynamic lockdown event messages.

packet

Displays dynamic lockdown packet messages.

Operational notes

- Dynamic IP lockdown is enabled at the port configuration level and applies to all bridged or routed IP packets entering the switch. The only IP packets that are exempt from dynamic IP lockdown are broadcast DHCP request packets, which are handled by DHCP snooping.

- DHCP snooping is a prerequisite for Dynamic IP Lockdown operation. The following restrictions apply:

  ◦ DHCP snooping is required for dynamic IP lockdown to operate. To enable DHCP snooping, enter the `dhcp-snooping` command at the global configuration level.

  ◦ Dynamic IP lockdown only filters packets in VLANs that are enabled for DHCP snooping. In order for Dynamic IP lockdown to work on a port, the port must be configured for at least one VLAN that is enabled for DHCP snooping.To enable DHCP snooping on a VLAN, enter the `dhcp-snooping vlan [vlan-id-range]` command at the global configuration level or the `dhcp-snooping` command at the VLAN configuration level.

  ◦ Dynamic IP lockdown is not supported on a trusted port. (However, note that the DHCP server must be connected to a trusted port when DHCP snooping is enabled.)By default, all ports are untrusted. To remove the trusted configuration from a port, enter the `no dhcp-snooping trust <port-list>` or `no dhcp6-snooping trust <port-list>` command at the global configuration level.

- After you enter the `ip source-lockdown` command (enabled globally with the desired ports entered in <port-list> the dynamic IP lockdown feature remains disabled on a port if any of the following conditions exist:

  ◦ If DHCP snooping has not been globally enabled on the switch.

  ◦ If the port is not a member of at least one VLAN that is enabled for DHCP snooping.

  ◦ If the port is configured as a trusted port for DHCP snooping.

  Dynamic IP lockdown is activated on the port only after you make the following configuration changes:

◦ Enable DHCP snooping on the switch.

◦ Configure the port as a member of a VLAN that has DHCP snooping enabled.

◦ Remove the trusted-port configuration.

- You can configure dynamic IP lockdown only from the CLI; this feature cannot be configured from the WebAgent or menu interface.

- If you enable dynamic IP lockdown on a port, you cannot add the port to a trunk.

- Dynamic IP lockdown must be removed from a trunk before the trunk is removed.

Adding an IP-to-MAC binding to the DHCP binding database

A switch maintains a DHCP binding database, which is used for dynamic IP lockdown as well as for DHCP and ARP packet validation. The DHCP snooping feature maintains the lease database by learning the IP-to-MAC bindings of VLAN traffic on untrusted ports. Each binding consists of the client MAC address, port number, VLAN identifier, leased IP address, and lease time.

Dynamic IP lockdown supports a total of 4K static and dynamic bindings with up to 64 bindings per port. When DHCP snooping is enabled globally on a VLAN, dynamic bindings are learned when a client on the VLAN obtains an IP address from a DHCP server. Static bindings are created manually with the CLI or from a downloaded configuration file.

When dynamic IP lockdown is enabled globally or on ports the bindings associated with the ports are written to hardware. This occurs during these events:

- Switch initialization

- Hot swap

- A dynamic IP lockdown-enabled port is moved to a DHCP snooping enabled VLAN

- DHCP snooping or dynamic IP lockdown characteristics are changed such that dynamic IP lockdown is enabled on the ports.

Potential issues with bindings

- When dynamic IP lockdown enabled, and a port or switch has the maximum number of bindings configured, the client DHCP request will be dropped and the client will not receive an IP address through DHCP.

- When dynamic IP lockdown is enabled and a port is configured with the maximum number of bindings, adding a static binding to the port will fail.

- When dynamic IP lockdown is enabled globally, the bindings for each port are written to hardware. If global dynamic IP lockdown is enabled and disabled several times, it is possible to run out of buffer space for additional bindings. The software will delay adding the bindings to hardware until resources are available.

Adding a static binding

To add the static configuration of an IP-to-MAC binding for a port to the database, enter the `ip source-binding` command at the global configuration level. Use the `no` form of the command to remove the IP-to-MAC binding from the database.

For IPv4

Syntax

```
[no] ip source-binding <mac-address> vlan <vlan-id> <ip-address>interface <port-number>
```

mac-address

Specifies a MAC address to bind with a VLAN and IP address on the specified port in the DHCP binding database.

vlan-id

Specifies a VLAN ID number to bind with the specified MAC and IP addresses on the specified port in the DHCP binding database.

ip-address

Specifies an IP address to bind with a VLAN and MAC address on the specified port in the DHCP binding database.

<port-number>

Specifies the port number on which the IP-to- MAC address and VLAN binding is configured in the DHCP binding database.

An example of the `ip source-binding` command is shown here:

```
HP Switch(config)# ip source-binding 0030c1-7f49c0
interface vlan 100 10.10.20.1 interface A4
```

> **NOTE**
>
> The `ip source-binding` command is the same command used by the Dynamic IP Lockdown feature to configure static bindings. The Dynamic ARP Protection and Dynamic IP Lockdown features share a common list of source IP-to-MAC bindings.

For IPv6

Syntax

```
[no] ipv6 source-binding <mac-address> vlan <vlan-id> <ip-address> interface <port-number>
```

mac-address

Specifies a MAC address to bind with a VLAN and IP address on the specified port in the DHCP binding database.

vlan-id

Specifies a VLAN ID number to bind with the specified MAC and IP addresses on the specified port in the DHCP binding database.

ip-address

Specifies an IPv6 address to bind with a VLAN and MAC address on the specified port in the DHCP binding database.

<port-number>

Specifies the port number on which the IP-to- MAC address and VLAN binding is configured in the DHCP binding database.

Verifying the dynamic IP lockdown configuration

To display the ports on which dynamic IP lockdown is configured, enter the `show ip source-lockdown status` command at the global configuration level.

For IPv4

Syntax

```
show ip source-lockdown status
```

Output for the `show ip source-lockdown status` command is shown in the following example.

Output for the show ip source-lockdown status command

```
HP Switch(config)# show ip source-lockdown status
 Dynamic IP Lockdown Status Information

Global State: Enabled

    Port     Operational State
    -------- -----------------
    1        Active
    2        Not in DHCP Snooping vlan
    3        Disabled
    4        Disabled
    5        Trusted port, Not in DHCP Snooping vlan
    . . . .  . . . . . . . . . . . .
```

For IPv6

Syntax

```
show ipv6 source-lockdown status
```

Displaying the static configuration of IP-to-MAC bindings

To display the static configurations of IP-to-MAC bindings stored in the DHCP lease database, enter the `show ip source-lockdown bindings` or `show ipv6 source-lockdown bindings` command.

For IPv4

Syntax

```
show ip source-lockdown bindings [port-number]
```

port-number

> (Optional) Specifies the port number on which source IP-to-MAC address and VLAN bindings are configured in the DHCP lease database.

The following example shows output from the `show ip source-lockdown bindings` command.

Output for the `show ip source-lockdown bindings` command

```
HP Switch (config)# show ip source-lockdown bindings

Dynamic IP Lockdown (DIPLD) Bindings

Mac Address     IP Address    VLAN   Port   Not in HW
-----------     ----------    ----   ----   ----------
001122-334455   10.10.10.1    1111   x11
```

```
005544-332211   10.10.10.2   2222   Trk11   YES
. . . . . . . . . . . . . . . . . . . . . . . . .
```

In the `show ip source-lockdown bindings` command output, the "Not in HW" column specifies whether or not (YES or NO) a statically configured IP-to- MAC and VLAN binding on a specified port has been combined in the lease database maintained by the DHCP Snooping feature.

For IPv6

Syntax

```
show ipv6 source-lockdown bindings [port-number]
```

Debugging dynamic IP lockdown

To enable the debugging of packets dropped by dynamic IP lockdown, enter the `debug dynamic-ip-lockdown` command.

Syntax:

```
debug dynamic-ip-lockdown
```

To send command output to the active CLI session, enter the `debug destination session` command.

Counters for denied packets are displayed in the `debug dynamic-ip-lockdown` command output. Packet counts are updated every five minutes. An example of the command output is shown in **Output for the debug dynamic-ip-lockdown command** on page 491.

When dynamic IP lockdown drops IP packets in VLAN traffic that do not contain a known source IP-to-MAC address binding for the port on which the packets are received, a message is entered in the event log.

Output for the debug dynamic-ip-lockdown command

```
HP Switch(config)# debug dynamic-ip-lockdown

DIPLD 01/01/90 00:01:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 1 packets
DIPLD 01/01/90 00:06:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 294 packets
DIPLD 01/01/90 00:11:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:16:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:21:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 299 packets
DIPLD 01/01/90 00:26:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:31:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
DIPLD 01/01/90 00:36:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 299 packets
DIPLD 01/01/90 00:41:25 : denied ip 192.168.2.100 (0)
(PORT 4) -> 192.168.2.1 (0), 300 packets
```

Using the instrumentation monitor

The instrumentation monitor can be used to detect anomalies caused by security attacks or other irregular operations on the switch. The following table shows the operating parameters that can be monitored at pre-determined intervals, and the possible security attacks that may trigger an alert:

Table 37: *Parameters for monitoring*

Parameter name	Description
pkts-to-closed-ports	The count of packets per minute sent to closed TCP/UDP ports. An excessive amount of packets could indicate a port scan, in which an attacker is attempting to expose a vulnerability in the switch.
arp-requests	The count of ARP requests processed per minute. A large amount of ARP request packets could indicate an host infected with a virus that is trying to spread itself.
ip-address-count	The number of destination IP addresses learned in the IP forwarding table. Some attacks fill the IP forwarding table causing legitimate traffic to be dropped.
system-resource-usage	The percentage of system resources in use. Some Denial-of- Service (DoS) attacks will cause excessive system resource usage, resulting in insufficient resources for legitimate traffic.`<1-2147483647>` – Set the threshold value`low` – Low threshold`med` – Medium threshold`high` – High threshold
login-failures/min	The count of failed CLI login attempts or SNMP management authentication failures. This indicates an attempt has been made to manage the switch with an invalid login or password. Also, it might indicate a network management station has not been configured with the correct SNMP authentication parameters for the switch.
port-auth-failures/min	The count of times a client has been unsuccessful logging into the network.
system-delay	The response time, in seconds, of the CPU to new network events such as BPDU packets or packets for other network protocols. Some DoS attacks can cause the CPU to take too long to respond to new network events, which can lead to a breakdown of Spanning Tree or other features. A delay of several seconds indicates a problem.
mac-address-count	The number of MAC addresses learned in the forwarding table. Some attacks fill the forwarding table so that new conversations are flooded to all parts of the network.
mac-moves/min	The average number of MAC address moves from one port to another per minute. This usually indicates a network loop, but can also be caused by DoS attacks.
learn-discards/min	Number of MAC address learn events per minute discarded to help free CPU resources when busy.

Operating notes

- To generate alerts for monitored events, you must enable the instrumentation monitoring log and/or SNMP trap. The threshold for each monitored parameter can be adjusted to minimize false alarms (see **Configuring instrumentation monitor** on page 493.

- When a parameter exceeds its threshold, an alert (event log message and/or SNMP trap) is generated to inform network administrators of this condition. The following example shows an event log message that occurs when the number of MAC addresses learned in the forwarding table exceeds the configured threshold:

Figure 111: *Event log message generated by instrumentation monitor*

Alerts are automatically rate limited to prevent filling the log file with redundant information. The following is an example of alerts that occur when the device is continually subject to the same attack (too many MAC addresses in this instance):

Figure 112: *Rate limiting when multiple messages are generated*

```
W 01/01/90 00:05:00 inst-mon: Limit for MAC addr count (300) is exceeded (321)
W 01/01/90 00:10:00 inst-mon: Limit for MAC addr count (300) is exceeded (323)
W 01/01/90 00:15:00 inst-mon: Limit for MAC addr count (300) is exceeded (322)
W 01/01/90 00:20:00 inst-mon: Limit for MAC addr count (300) is exceeded (324)
W 01/01/90 00:20:00 inst-mon: Ceasing logs for MAC addr count for 15 minutes
```

In the preceding example, if a condition is reported 4 times (persists for more than 15 minutes) then alerts cease for 15 minutes. If after 15 minutes the condition still exists, the alerts cease for 30 minutes, then for 1 hour, 2 hours, 4 hours, 8 hours, and after that the persisting condition is reported once a day. As with other event log entries, these alerts can be sent to a server.

Known Limitations: The instrumentation monitor runs once every five minutes. The current implementation does not track information such as the port, MAC, and IP address from which an attack is received.

Configuring instrumentation monitor

The following commands and parameters are used to configure the operational thresholds that are monitored on the switch. By default, the instrumentation monitor is disabled.

Syntax:

```
[no] instrumentation monitor [parameterName|all] [<low|med|high|limitValue>]
```

[log]

Enables/disables instrumentation monitoring log so that event log messages are generated every time there is an event which exceeds a configured threshold. (Default threshold setting when instrumentation monitoring is enabled: enabled)

[all]

Enables/disables all counter types on the switch but does not enable/disable instrumentation monitor logging. (Default threshold setting when enabled: see parameter listings below)

[arp-requests]

The number of arp requests that are processed each minute. (Default threshold setting when enabled: 1000 med)

[ip-address-count]

The number of destination IP addresses learned in the IP forwarding table. (Default threshold setting when enabled: 1000 med)

[learn-discards]

The number of MAC address learn events per minute discarded to help free CPU resources when busy. (Default threshold setting when enabled: 100 med)

[login-failures]

The count of failed CLI login attempts or SNMP management authentication failures per hour. (Default threshold setting when enabled: 10 med)

[mac-address-count]

The number of MAC addresses learned in the forwarding table. You must enter a specific value in order to enable this feature. (Default threshold setting when enabled: 1000 med)

[mac-moves]

The average number of MAC address moves per minute from one port to another. (Default threshold setting when enabled: 100 med)

[pkts-to-closed-ports]

The count of packets per minute sent to closed TCP/UDP ports. (Default threshold setting when enabled: 10 med)

[port-auth-failures]

The count of times per minute that a client has been unsuccessful logging into the network. (Default threshold setting when enabled: 10 med)

[system-resource-usage]

The percentage of system resources in use. (Default threshold setting when enabled: 50 med)

`<1-2147483647>` – Set the threshold value

`low` – Low threshold

`med` – Medium threshold

`high` – High threshold

[system-delay]

The response time, in seconds, of the CPU to new network events such as BPDU packets or packets for other network protocols. (Default threshold setting when enabled: 3 seconds med)

[trap]

> Enables or disables SNMP trap generation. (Default setting when instrumentation monitoring is enabled: disabled)

To enable instrumentation monitor using the default parameters and thresholds, enter the general `instrumentation monitor` command. To adjust specific settings, enter the name of the parameter that you wish to modify, and revise the threshold limits as needed.

Examples:

To turn on monitoring and event log messaging with the default medium values:

```
HP Switch(config)# instrumentation monitor
```

To turn off monitoring of the system delay parameter:

```
HP Switch(config)# no instrumentation monitor system-
delay
```

To adjust the alert threshold for the MAC address count to the low value:

```
HP Switch(config)# instrumentation monitor mac-address-
count low
```

To adjust the alert threshold for the MAC address count to a specific value:

```
HP Switch(config)# instrumentation monitor mac-address-
count 767
```

To enable monitoring of learn discards with the default medium threshold value:

```
HP Switch(config)# instrumentation monitor learn-
discards
```

To disable monitoring of learn discards:

```
HP Switch(config)# no instrumentation monitor
learn-discards
```

To enable or disable SNMP trap generation:

```
HP Switch(config)# [no] instrumentation monitor
trap
```

Viewing the current instrumentation monitor configuration

The `show instrumentation monitor configuration` command displays the configured thresholds for monitored parameters.

Viewing the instrumentation monitor configuration

```
HP Switch# show instrumentation monitor configuration

 PARAMETER                LIMIT
 ------------------------ ---------------
 mac-address-count        1000 (med)
 ip-address-count         1000 (med)
 system-resource-usage    50 (med)
 system-delay             5 (high)
 mac-moves/min            100 (med)
 learn-discards/min       100 (med)
 ip-port-scans/min        10 (med)
 arp-requests/min         100 (low)
 login-failures/min       10 (med)
 port-auth-failures/min   10 (med)
```

```
SNMP trap generation for alerts: enabled
Instrumentation monitoring log : enabled
```

An alternate method of determining the current Instrumentation Monitor configuration is to use the `show run` command. However, the `show run` command output does not display the threshold values for each limit set.

Overview

You can enhance in-band security and improve control over access to network resources by configuring static filters to forward (the default action) or drop unwanted traffic. That is, you can configure a traffic filter to either forward or drop all network traffic moving to outbound (destination) ports and trunks (if any) on the switch.

Filter limits

The switch accepts up to 101 static filters. These limitations apply:

- Source-port filters: up to 78

- Multicast filters: up to 16 with 1024 or fewer VLANs configured. Up to 8 with more than 1024 VLANs configured.

- Protocol filters: up to 7

Using port trunks with filter

The switch manages a port trunk as a single source or destination for sourceport filtering. If you configure a port for filtering before adding it to a port trunk, the port retains the filter configuration, but suspends the filtering action while a member of the trunk. If you want a trunk to perform filtering, first configure the trunk, then configure the trunk for filtering. See **Configuring a filter on a port trunk** on page 509.

Filter types and operation

The following table represents the types of static filters and their selection criteria:

Table 38: *Filter types and criteria*

Static filter type	Selection criteria
Source-port	Inbound traffic from a designated, physical source-port will be forwarded or dropped on a per-port (destination) basis.
Multicast	Inbound traffic having a specified multicast MAC address will be forwarded to outbound ports (the default) or dropped on a per-port (destination) basis.
Protocol	Inbound traffic having the selected frame (protocol) type will be forwarded or dropped on a per-port (destination) basis.

Source-port filters

This filter type enables the switch to forward or drop traffic from all end nodes on the indicated source-port to specific destination ports.

Figure 113: *Source-port filer application*

Configuring a source-port filter to drop traffic received on port 1 with an outbound destination of port 2 means that End Nodes A, B, and C cannot send traffic to the server. To block traffic in the opposite direction, you would also configure a source-port filter to drop traffic received on port 2 with an outbound destination of port 1.

Operating rules for source-port filters

- You can configure one source-port filter for each physical port and port trunk on the switch. (See **Defining and configuring named source-port filters** on page 500.)

- You can include all destination ports and trunks in the switch on a single source-port filter.

- Each source-port filter includes:

  ○ One source port or port trunk (trk1, trk2, ...trk**n**)

  ○ A set of destination ports and/or port trunks that includes all untrunked LAN ports and port trunks on the switch

  ○ An action (forward or drop) for each destination port or port trunk

  When you create a source-port filter, the switch automatically sets the filter to forward traffic from the designated source to all destinations for which you do not specifically configure a "drop" action. Thus, it is not necessary to configure a source-port filter for traffic you want the switch to forward unless the filter was previously configured to drop the desired traffic.

- When you create a source port filter, all ports and port trunks (if any) on the switch appear as destinations on the list for that filter, even if routing is disabled and separate VLANs and/or subnets exist. Where traffic would normally be allowed between ports and/or trunks, the switch automatically forwards traffic to the outbound ports and/or trunks you do not specifically configure to drop traffic. (Destination ports that comprise a trunk are listed collectively by the trunk name—such as Trk1— instead of by individual port name.)

- Packets allowed for forwarding by a source-port filter are subject to the same operation as inbound packets on a port that is not configured for source-port filtering.

With multiple IP addresses configured on a VLAN, and routing enabled on the switch, a single port or trunk can be both the source and destination of packets moving between subnets in that same VLAN. In this case, you can

prevent the traffic of one subnet from being routed to another subnet of the same port by configuring the port or trunk as both the source and destination for traffic to drop.

Example:

If you wanted to prevent server "A" from receiving traffic sent by workstation "X", but do not want to prevent any other servers or end nodes from receiving traffic from workstation "X", you would configure a filter to drop traffic from port 5 to port 7. The resulting filter would drop traffic from port 5 to port 7, but would forward all other traffic from any source port to any destination port. (See **Figure 114: Filter blocking traffic only from Port 5 to Server A** on page 499 and **Figure 115: Filter for the actions shown in Filter blocking traffic only from Port 5 to Server A** on page 499.

Figure 114: *Filter blocking traffic only from Port 5 to Server A*

Figure 115: *Filter for the actions shown in Filter blocking traffic only from Port 5 to Server A*

Name source-port filters

You can specify named source-port filters that may be used on multiple ports and port trunks. A port or port trunk can only have one source-port filter, but by using this capability you can define a source-port filter once and apply it to multiple ports and port trunks. This can make it easier to configure and manage source-port filters on your switch. The commands to define, configure, apply, and display the status of named source-port filters are described below.

Operating rules for named source-port filters

- A port or port trunk may only have one source-port filter, named or not named.

- A named source-port filter can be applied to multiple ports or port trunks.

- Once a named source-port filter is defined, subsequent changes only modify its action, they don't replace it.

- To change the named source-port filter used on a port or port trunk, the current filter must first be removed, using the `no filter source-port named-filter <filter-name>` command.

- A named source-port filter can only be deleted when it is not applied to any ports.

Defining and configuring named source-port filters

The `named source-port filter` command operates from the global configuration level.

Syntax:

`[no] filter source-port named-filter <filter-name>`

Defines or deletes a named source-port filter. The `<filter-name >` may contain a maximum of 20 alpha-numeric characters (longer names may be specified, but they are not displayed.) A filter-name cannot be a valid port or port trunk name.

The maximum number of named source-port filters that can be used is equal to the number of ports on a switch.

A named source-port filter can only be removed if it is not in use (use the `show filter source-port` command to check the status). Named source-port filters are not automatically deleted when they are no longer used.

Use the `no` option to delete an unused named source-port filter.

Syntax:

`filter source-port named-filter <filter-name> drop <destination-port-list>`

Configures the named source-port filter to drop traffic having a destination on the ports and/or port trunks in the `<destination-port-list >`. Can be followed by the `forward`option if you have other destination ports or port trunks previously set to drop that you want to change to `forward`.

For example: `filter source-port named-filter <filter-name> drop <destination-port-list> forward <destination-port-list> `.

The `destination-port-list` may contain ports, port trunks, and ranges (for example 3-7 or trk4-trk9) separated by commas.

Syntax:

`filter source-port named-filter <filter-name> forward <destination-port-list>`

Configures the named source-port filter to forward traffic having a destination on the ports and/or port trunks in the `<destination-port-list>`. Since "forward" is the default state for destinations in a filter, this command is useful when destinations in an existing filter are configured for "drop" and you want to change them to "forward". Can be followed by the `drop`option if you have other destination ports set to `forward` that you want to change to `drop`.

For example: `filter source-port named-filter <filter-name> forward <destination-port-list> drop <destination-port-list>`

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

A named source-port filter must first be defined and configured before it can be applied. In the following example two named source-port filters are defined, `web-only` and `accounting`.

```
HP Switch(config)# filter source-port named-filter web-only
```

```
HP Switch(config)# filter source-port named-filter accounting
```

By default, these two named source-port filters forward traffic to all ports and port trunks.

To configure a named source-port filter to prevent inbound traffic from being forwarded to specific destination switch ports or port trunks, the `drop`option is used. For example, on a 26-port switch, to configure the named source-port filter `web-only` to drop any traffic except that for destination ports 1 and 2, the following command would be used:

```
HP Switch(config)# filter source-port named-filter
web-only drop 3-26
```

A named source-port filter can be defined and configured in a single command by adding the `drop` option, followed by the required destination-port-list.

Viewing a named source-port filter

You can list all source-port filters configured in the switch, both named and unnamed, and their action using the `show` command below.

Syntax:

```
show filter source-port
```

Displays a listing of configured source-port filters, where each filter entry includes a Filter Name, Port List, and Action:

Filter Name

The filter-name used when a named source-port filter is defined. Non-named source-port filters are automatically assigned the port or port trunk number of the source port.

Port List

Lists the port and port trunk destinations using the filter. Named source-port filters that are not in use display NOT USED.

Action

Lists the ports and port trunks dropped by the filter. If a named source-port filter has been defined but not configured, this field is blank.

[index]

For the supplied index (IDX) displays the action taken (Drop or Forward) for each destination port on the switch.

Using named source-port filters

A company wants to manage traffic to the Internet and its accounting server on a 26-port switch. Their network is pictured in **Figure 116: Network configuration for named source-port filters** on page 502. Switch port 1

connects to a router that provides connectivity to a WAN and the Internet. Switch port 7 connects to the accounting server. Two workstations in accounting are connected to switch ports 10 and 11.

Figure 116: *Network configuration for named source-port filters*

Example:

While named source-port filters may be defined and configured in two steps, this is not necessary. Here we define and configure each of the named source-port filters for our example network in a single step.

Figure 117: *Applying example named source-port filters*

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Once the named source-port filters have been defined and configured we now apply them to the switch ports.

Figure 118: *Source port filters applied to switch ports*

```
HP Switch(config)# filter source-port 2-6,8,9,12-26 named-filter web-only
HP Switch(config)# filter source-port 7,10,11 named-filter accounting
HP Switch(config)# filter source-port 1 named-filter no-incoming-web
HP Switch(config)#
```

The `show filter` command shows what ports have filters applied.

Figure 119: *Example of the show filter command*

```
HP Switch(config)# show filter
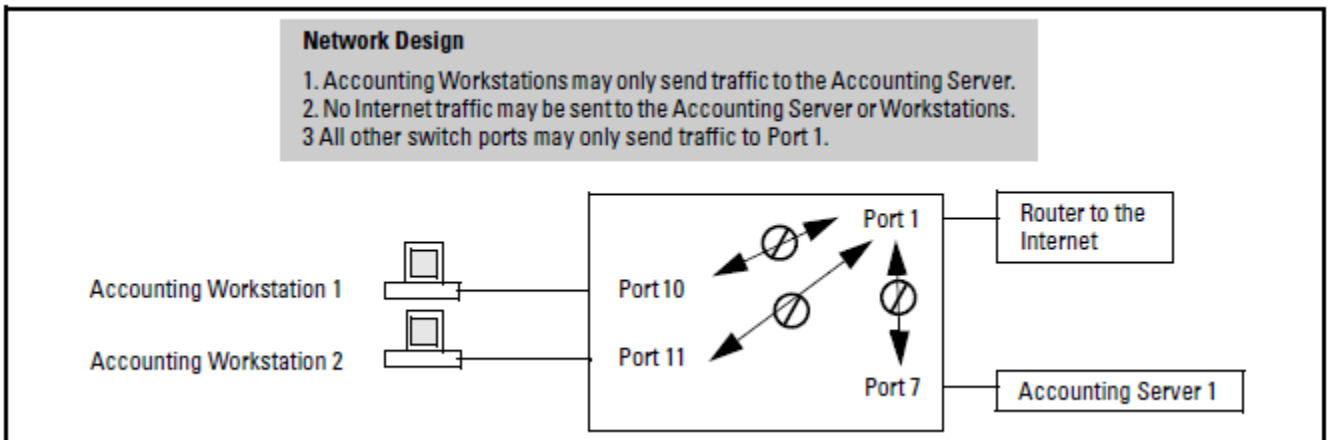
Traffic/Security Filters

  IDX Filter Type   | Value
  --------------- + -------------------
   1    Source Port  | 2
   2    Source Port  | 3
   3    Source Port  | 4
   4    Source Port  | 5
   5    Source Port  | 6
   6    Source Port  | 8
   7    Source Port  | 9
   8    Source Port  | 12

  20    Source Port  | 24
  21    Source Port  | 25
  22    Source Port  | 26
  23    Source Port  | 7
  24    Source Port  | 10
  25    Source Port  | 11
  26    Source Port  | 1
```

Indicates the port number or port-trunk name of the source port or trunk assigned to the filter.

An automatically assigned index number used to identify the filter for a detailed information listing. A filter retains its assigned IDX number for as long as the filter exists in the switch. The switch assigns the lowest available IDX number to a new filter. This can result in a newer filter having a lower IDX number than an older filter if a previous (source-port or named source-port) filter deletion created a gap in the filter listing.

Using the IDX value in the `show filter` command, we can see how traffic is filtered on a specific port (Value). The two outputs below show a non-accounting and an accounting switch port.

Figure 120: *Showing traffic filtered on specific ports*

```
HP Switch(config)# show filter 4          HP Switch(config)# show filter 24

 Traffic/Security Filters                  Traffic/Security Filters

  Filter Type : Source Port                 Filter Type : Source Port
  Source Port : 5                           Source Port : 10

  Dest   Port Type      |  Action           Dest   Port Type      | Action
  -------- --------- + -------            -------- --------- + -----------
  1        10/100TX     | Forward           1        10/100TX     | Drop
  2        10/100TX     | Drop              2        10/100TX     | Drop
  3        10/100TX     | Drop              3        10/100TX     | Drop
  4        10/100TX     | Drop              4        10/100TX     | Drop
  5        10/100TX     | Drop              5        10/100TX     | Drop
  6        10/100TX     | Drop              6        10/100TX     | Drop
  7        10/100TX     | Drop              7        10/100TX     | Forward
  8        10/100TX     | Drop              8        10/100TX     | Drop
  9        10/100TX     | Drop              9        10/100TX     | Drop
  10       10/100TX     | Drop              10       10/100TX     | Drop
  11       10/100TX     | Drop              11       10/100TX     | Drop
  12       10/100TX     | Drop              12       10/100TX     | Drop

     .            .            .                .            .            .
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

The same command, using IDX 26, shows how traffic from the Internet is handled.

Figure 121: *Source port filtering with internet traffic*

```
HP Switch(config)# show filter 26

 Traffic/Security Filters

  Filter Type : Source Port
  Source Port : 1

  Dest      Port Type       | Action
  --------- --------- + -----------------------
  1           10/100TX        | Forward
  2           10/100TX        | Forward
  3           10/100TX        | Forward
  4           10/100TX        | Forward
  5           10/100TX        | Forward
  6           10/100TX        | Forward
  7           10/100TX        | Drop
  8           10/100TX        | Forward
  9           10/100TX        | Forward
  10          10/100TX        | Drop
  11          10/100TX        | Drop
  12          10/100TX        | Forward
       .           .              .
```

As the company grows, more resources are required in accounting. Two additional accounting workstations are added and attached to ports 12 and 13. A second server is added attached to port8.

Figure 122: *Expanded network configuration for named source-port filters*

Network Design
1. Accounting Workstations may only send traffic to the Accounting Server.
2. No Internet traffic may be sent to the Accounting Server or Workstations.
3 All other switch ports may only send traffic to Port 1.

The following revisions to the named source-port filter definitions maintain the desired network traffic management, as shown in the Action column of the `show` command.

Figure 123: *Network traffic management with source port filters*

```
HP Switch(config)# filter source-port named-filter accounting forward 8,12,13
HP Switch(config)# filter source-port named-filter no-incoming-web drop 8,12,13
HP Switch(config)#
HP Switch(config)# show filter source-port
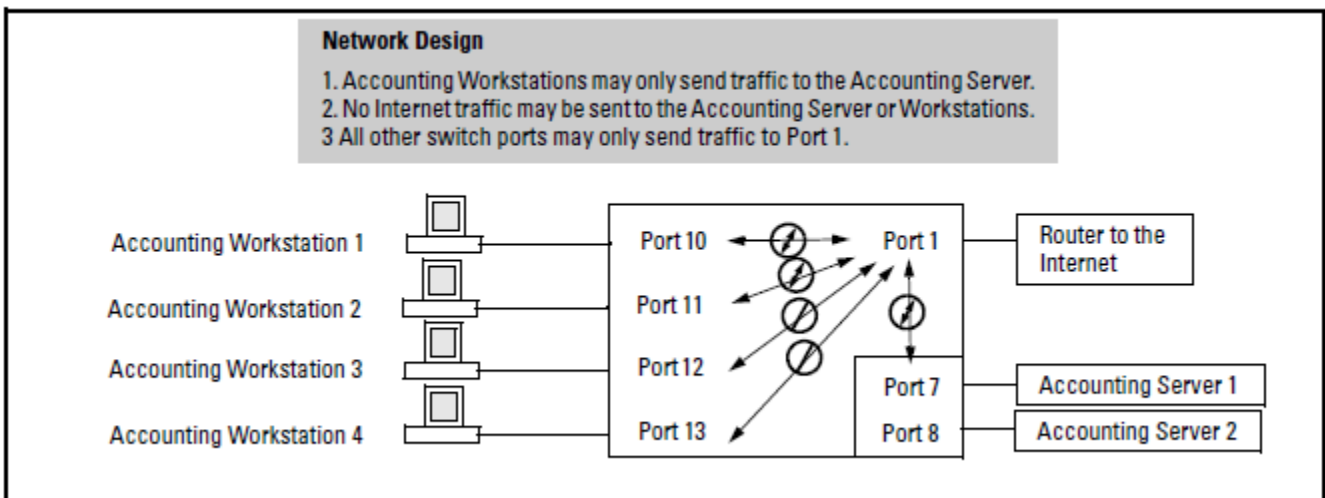
 Traffic/Security Filters

 Filter Name         | Port List        | Action
 ------------------ + ---------------- + ----------------
 web-only            | 2-6,8-9,12-26    | drop 2-26
 accounting          | 7,10-11          | drop 1-6,9,14-26
 no-incoming-web     | 1                | drop 7-8,10-13

HP Switch(config)#
```

We next apply the updated named source-port filters to the appropriate switch ports. As a port can only have one source-port filter (named or not named), before applying the new named source-port filters we first remove the existing source-port filters on the port.

Figure 124: *No filter source-port*

```
HP Switch(config)# no filter source-port 8,12,13
HP Switch(config)# filter source-port 8,12,13 named-filter accounting
HP Switch(config)#
```

The named source-port filters now manage traffic on the switch ports as shown below, using the `show filter source-port` command.

Figure 125: *Named source-port filters managing traffic*

```
HP Switch(config)# show filter source-port

 Traffic/Security Filters

 Filter Name         | Port List        | Action
 ------------------ + ---------------- + ----------------
 web-only            | 2-6,8-9,14-26    | drop 2-26
 accounting          | 7-8,10-13        | drop 1-6,9,14-26
 no-incoming-web     | 1                | drop 7-8,10-13

HP Switch(config)#
```

Static multicast filters

This filter type enables the switch to forward or drop multicast traffic to a specific set of destination ports. This helps to preserve bandwidth by reducing multicast traffic on ports where it is unnecessary, and to isolate multicast traffic to enhance security.

You can configure up to 16 static multicast filters (defined by the `filter` command). However, if an IGMP-controlled filter for a joined multicast group has the same multicast address as a static multicast filter configured on a given port, the IGMP-controlled filter overrides the static multicast filter configured on that port. Note that in the default configuration, IGMP is disabled on VLANs configured in the switch. To enable IGMP on a specific VLAN, use the `vlan <vid> ip igmp` command. (For more on this command, see "Multimedia Traffic Control with IP Multicast (IGMP)" in the multicast and routing guide for your switch.)

The total of static multicast filters and IGMP multicast filters together can range from 389 to 420, depending on the current `max-vlans` setting in the switch. If multiple VLANs are configured, then each filter is counted once per VLAN in which it is used.

Table 39: *Multicast filer limits*

Max-VLANs setting	Max # multicast filters (static and IGMP combined)
1 (minimum)	420
8 (default)	413
32 or higher	389

Per-Port IP Multicast Filters

The static multicast filters described in this section filter traffic having a multicast address you specify. To filter all multicast traffic on a per-VLAN basis, see "Configuring and Displaying IGMP" in the multicast and routing guide for your switch.

IP Multicast Filters

Multicast filters are configured using the Ethernet format for the multicast address. IP multicast addresses occur in the range of 224.0.0.0 through 239.255.255.255 (which corresponds to the Ethernet multicast address range of 01005e-000000 through 01005e-7fffff). Any static Traffic/ Security filters configured with a multicast filter type and a multicast address in this range will continue to be in effect unless IGMP learns of a multicast group destination in this range. In this case, IGMP takes over the filtering function for the multicast destination address(es) for as long as the IGMP group is active. If the IGMP group subsequently deactivates, the static filter resumes control over traffic to the multicast address.

> **CAUTION**
> If Spanning Tree is enabled, then the MSTP multicast MAC address (0180c2- 000000) should not be filtered. (STP will not operate properly if the MSTP multicast MAC address is filtered.)

Protocol filters

This filter type enables the switch to forward or drop, on the basis of protocol type, traffic to a specific set of destination ports on the switch. Filtered protocol types include:

- Appletalk
- ARP
- IPX
- NetBEUI
- SNA

Only one filter for a particular protocol type can be configured at any one time. For example, a separate protocol filter can be configured for each of the protocol types listed above, but only one of those can be an IP filter. Also, the destination ports for a protocol filter can be on different VLANs.

You can configure up to seven protocol filters.

Configuring traffic/security filters

Use this procedure to specify the type of filters to use on the switch and whether to forward or drop filtered packets for each filter you specify.

Procedure

1. Select the static filter type(s).

2. For inbound traffic matching the filter type, determine the filter action you want for each outbound (destination) port on the switch (forward or drop). The default action for a new filter is to forward traffic of the specified type to all outbound ports.

3. Configure the filter.

4. Use `show filter` to check the filter listing to verify that you have configured correct action for the desired outbound ports.

Configuring a source-port traffic filter

Syntax:

```
[no] filter [source-port <port-number|trunk-name>]
```

Specifies one inbound port or trunk. Traffic received inbound on this interface from other devices will be filtered. The `no` form of the command deletes the source-port filter for <port-number> and returns the destination ports for that filter to the Forward action. (Default: Forward on all ports.)

> **NOTE**
>
> If multiple VLANs are configured, the source-port and the destination port(s) must be in the same VLAN unless routing is enabled. Similarly, if a VLAN containing both the source and destination is multinetted, the source and destination ports and/or trunks must be in the same subnet unless routing is enabled.

Syntax:

```
[drop] <destination-port-list> [forward <port-list>]
```

Configures the filter to drop traffic for the ports and/or trunks in the designated `<destination-port-list>` . Can be followed by `forward <destination-port-list>` if you have other destination ports set to `drop` that you want to change to `forward`. If no drop or forward action is specified, the switch automatically creates a filter with a `forward` action from the designated source port (or trunk) to all destination ports (or trunks) on the switch.

Syntax:

```
[forward] <port-list>
```

Configures the filter to forward traffic for the ports and/ or trunks in the designated `<destination-port-list>` . Because `forward` is the default state for destinations in a filter, this command is useful when destinations in an existing filter are configured for `drop` and you want to change them to `forward`. Can be followed by `drop <destination-port-list>` if you have other destination ports set to `forward` that you want to change to `drop`. If no drop or forward action is specified, the switch automatically creates a filter with a forward action from the designated source port (or trunk) to all destination ports (or trunks) on the switch.

Example:

For example, assume that you want to create a source-port filter that drops all traffic received on port 5 with a destination of port trunk 1 (Trk1) and any port in the range of port 10 to port 15. To create this filter you would execute this command:

```
HP Switch(config)# filter source-port 5 drop trk1,10-15
```

Later, suppose you wanted to shift the destination port range for this filter up by two ports; that is, to have the filter drop all traffic received on port 5 with a destination of any port in the range of port 12 to port 17. (The Trk1

destination is already configured in the filter and can remain as-is.)With one command you can restore forwarding to ports 10 and 11 while adding ports 16 and 17 to the "drop" list:

```
HP Switch(config)# filter source-port 5 forward 10-11 drop
16-17
```

Configuring a filter on a port trunk

This operation uses the same command as is used for configuring a filter on an individual port. However, the configuration process requires two steps:

1. Configure the port trunk.

2. Configure a filter on the port trunk by using the trunk name (trk1, trk2, ...trk6) instead of a port name.

For example, to create a filter on port trunk 1 to drop traffic received inbound for trunk 2 and ports 10-15:

```
HP Switch(config)# filter source-port trk1 drop trk2,
10-15
```

Note that if you first configure a filter on a port and then later add the port to a trunk, the port remains configured for filtering but the filtering action will be suspended while the port is a member of the trunk. That is, the trunk does not adopt filtering from the port configuration. You must still explicitly configure the filter on the port trunk. If you use the `show filter <index>` command for a filter created before the related source port was added to a trunk, the port number appears between asterisks (*), indicating that the filter action has been suspended for that

filter. For example, if you create a filter on port 5, then create a trunk with ports 5 and 6, and display the results, you would see the following:

Figure 126: *Switch response to adding a filtered source port to a trunk*

```
HP Switch(config)# filter source-port 5 drop 2
HP Switch(config)# trunk 5-6 trk1
HP Switch(config)# show filter

 Traffic/Security Filters

  IDX Filter Type  | Value
  --- ----------- + -----------
  1    Source Port | *5*

HP Switch(config)# show filter 1

 Traffic/Security Filters

  Filter Type : Source Port
  Source Port : *5*

  Dest Port Type         | Action
  --------- --------- + -------
  1         100/1000T | Forward
  2         100/1000T | Forward
  3         100/1000T | Forward
  4         100/1000T | Forward
  .              .          .
  .              .          .
  .              .          .
```

The *5* shows that port 5 is configured for filtering, but the filtering action has been suspended while the port is a member of a trunk.

If you want the trunk to which port 5 belongs to filter traffic, then you must explicitly configure filtering on the trunk.

Note: If you configure an existing trunk for filtering and later add another port to the trunk, the switch will apply the filter to all traffic moving on any link in the trunk. If you remove a port from the trunk it returns to the configuration it had before it was added to the trunk

Editing a source-port filter

The switch includes in one filter the action(s) for all destination ports and/or trunks configured for a given source port or trunk. Thus, if a source-port filter already exists and you want to change the currently configured action for some destination ports or trunks, use the `filter source-port` command to update the existing filter. For example, suppose you configure a filter to drop traffic received on port 8 and destined for ports 1 and 2. The resulting filter is shown on the left in **Assigning additional destination ports to an existing filter** on page 510. Later, you update the filter to drop traffic received on port 8 and destined for ports 3 through 5. Since only one filter exists for a given source port, the filter on traffic from port 8 appears as shown on the right in **Assigning additional destination ports to an existing filter** on page 510:

Assigning additional destination ports to an existing filter

```
HP Switch(config)# show filter 1      HP Switch(config)# show filter 1
 Traffic/Security Filters              Traffic/Security Filters
  Filter Type : Source Port             Filter Type : Source Port
  Source Port : 8                       Source Port : 8

  Dest Port Type      | Action          Dest Port Type      | Action
  --------- --------- + -------         --------- --------- + -------
  1         100/1000T | Drop            1         100/1000T | Drop
  2         100/1000T | Drop            2         100/1000T | Drop
```

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

```
3          100/1000T | Forward      3          100/1000T | Drop
4          100/1000T | Forward      4          100/1000T | Drop
5          100/1000T | Forward      5          100/1000T | Drop
6          100/1000T | Forward      6          100/1000T | Forward
7          100/1000T | Forward      7          100/1000T | Forward
```

Configuring a multicast or protocol traffic filter

Syntax:

```
[no] filter
```

[multicast *<mac-address>*]

Specifies a multicast address. Inbound traffic received (on any port) with this multicast address will be filtered. (Default: Forward on all ports.)

The `no` form of the command deletes the multicast filter for the <mac-address> multicast address and returns the destination ports for that filter to the `Forward` action.

[<forward|drop> *<port-list>*]

Specifies whether the designated destination port(s) should forward or drop the filtered traffic.

[protocol <ip|ipx|arp|appletalk|sna|netbeui>]

Specifies a protocol type. Traffic received (on any port) with this protocol type will be filtered. (Default: Forward on all ports.)

The `no` form of the command deletes the protocol filter for the specified protocol and returns the destination ports for that filter to the `Forward`action.

[<forward|drop> *<port-list>*]

Specifies whether the designated destination port(s) should forward or drop the filtered traffic.

Example:

Suppose you wanted to configure the filters in the following table on a switch. (For more on source-port filters, see **Configuring a source-port traffic filter** on page 508.

Table 40: *Filter example*

Filter type	Filter value	Action	Destination ports
source-port	Inbound ports: 1, 2	Drop	1-4
multicast	010000-123456	Drop	5-8, 9-12
multicast	010000-224466	Drop	13-15
protocol	Appletalk	Drop	16-19, 1
protocol	ARP	Drop	22, 23-24

The following commands configure the filters listed above:

Configuring various traffic/security filters

```
HP Switch(config)# filter source-port 1 drop e 1-4
HP Switch(config)# filter source-port 2 drop 1-4
```

```
HP Switch(config)# filter multicast 010000-123456 drop e 5-8,9-12
HP Switch(config)# filter multicast 010000-224466 drop e 13-15
HP Switch(config)# filter protocol appletalk drop e 16-19,1
HP Switch(config)# filter protocol arp drop e 22,23-24
```

Filtering index

The switch automatically assigns each new filter to the lowest-available index (IDX) number. The index numbers are included in the `show filter` command described in the next section and are used with the `show filter <index>` command to display detailed information about a specific filter.

If there are no filters currently configured, and you create three filters in succession, they will have index numbers 1-3. However, if you then delete the filter using index number "2" and then configure two new filters, the first new filter will receive the index number "2" and the second new filter will receive the index number "4". This is because the index number "2" was made vacant by the earlier deletion, and was therefore the lowest index number available for the next new filter.

Displaying traffic/security filters

This command displays a listing of all filters by index number and also enables you to use the index number to display the details of individual filters.

Syntax:

```
show filter
```

Lists the filters configured in the switch, with corresponding filter index (IDX) numbers. IDX: An automatically assigned index number used to identify the filter for a detailed information listing. A filter retains its assigned IDX number for as long as the filter exists in the switch. The switch assigns the lowest available IDX number to a new filter. This can result in a newer filter having a lower IDX number than an older filter if a previous filter deletion created a gap in the filter listing.

Filter Type

Indicates the type of filter assigned to the IDX number (source-port, multicast, or protocol).

Value

Indicates the port number or port-trunk name of the source port or trunk assigned to the filter.

[*index*]

Lists the filter type and other data for the filter corresponding to the index number in the show filter output. Also lists, for each outbound destination port in the switch, the port number, port type, and filter action (forward or drop). The switch assigns the lowest available index number to a new filter. If you delete a filter, the index number for that filter becomes available for the next filter you create.

Example:

To display the filters created in **Configuring a multicast or protocol traffic filter** and then list the details of the multicast filter for multicast address 010000-224466:

Figure 127: *Displaying filter data*

```
                    HP Switch(config)# show filter          Lists all filters configured
                                                            in the switch.
                    Traffic/Security Filters

                    IDX  Filter Type   | Value
 Filter Index Numbers  ---  -----------  + -----------       Criteria for Individual
 (Automatically Assigned)  1   | Source Port   | 1            Filters
                    2   | Source Port   | 2
                    3   | Multicast     | 010000-123456 |
                    4   | Multicast     | 010000-224466 |
                    5   | Protocol      | Appletalk     |
                    6   | Protocol      | ARP           |

                    HP Switch(config)# show filter          Uses the index number
                                                            (IDX) for a specific filter
                    Traffic/Security Filters                to list the details for that
                                                            filter only.
                    Filter Type : Multicast
                    Multi-cast Addres : 010000-224466

                    Dest Port  Type       | Action
                    ---------  ---------  + -------
                    1          1000LX     | Forward
                    2                     | Forward
                    3                     | Forward
                    4          1000SX     | Forward
                    5          100/1000T  | Drop
                    6          100/1000T  | Drop
                    7          100/1000T  | Drop
                    8          100/1000T  | Drop
                    9          10/100TX   | Forward
                    10         10/100TX   | Forward
                    11         10/100TX   | Forward
                     .           .            .
```

Advanced Threat Detection

Enables an Aruba OS switch to accept different filters for different syslog servers. Filtering of syslog messages is a requirement for CPPM, Niara, or any other threat detection systems in an effort to curb security threats.

logging

Syntax

```
logging <IP-ADDR> [UDP 1024-49151 | TCP 1024-49151]
 filter <FILTER-NAME>

no logging <IP-ADDR> [UDP 1024-49151 | TCP 1024-49151]
 filter <FILTER-NAME>
```

Description

Enables the configuration of a named filter to a specific server. The attributes associated with this defined filter will be applied to the chosen server.

The `no` form of this command removes the associated filter from the designated server.

Command context

`Config`

Parameters

<IP-ADDR>

Specifies the IP address for logging into the syslog server.

<UDP 1024-49151 | TCP 1024-49151>

Specifies the configuration of the UDP or TCP transport protocol for the transmission of logging messages to a syslog server.

<FILTER-NAME>

Specifies the file to be filtered in syslog server.

Examples

```
Logging to syslog server using IP address

switch# logging <IP-ADDR> <UDP 1024-49151 | TCP 1024-49151>
 filter <FILTER-NAME> logging 10.10.10.1 tcp 2500 filter cppm_filter

switch# logging <IP-ADDR> filter <FILTER-NAME>
logging 10.10.10.1 filter cppm_filter
```

logging filter

Syntax

```
logging filter <NAME> [per-ip] [severity <SEVERITY>
| event-list <NUM-RANGE> [permit|deny] | system-module
<MODULE-NAME>]

no logging filter <NAME> [per-ip] [severity <SEVERITY>
| event-list <NUM-RANGE> [permit|deny]| system-module
<MODULE-NAME>]
```

Description

Creates a syslog filter in the logging context. The `per-ip` filter is enabled by default.

Using PER-IP filter, you can create the global filter to match on few event logs to be sent to the syslog servers. However, this filter will automatically get applied on all the syslog servers configured on switch. It is not possible to apply the filter on few servers. Hence, PER-IP filter is created to match on few event logs based on event-id, severity and system modules. For more information on PER-IP filters, see the *IPv6 Guide* for your switch.

The `no` form of this command removes the `per-ip` filter.

Command context

`Config`

Parameters

<NAME>

Specifies the name that identifies the filter.

<SEVERITY>

Specifies the severity of an event-major, warning, error, info or debug.

<NUM-RANGE>

Specifies an event number list to match. User can create either permit list or deny list.

<MODULE-NAME>

Specifies the system module name.

Examples

```
Creating Per IP severity:
switch# logging filter <NAME> [per-ip] [severity <SEVERITY>]
logging filter cppm_filter per-ip severity major

Creating Per IP Event list:
switch# logging filter <NAME> [per-ip]  [event-list <EVENT-NUM-RANGE>] [permit|
deny]
logging filter cppm_filter per-ip event-list 2-4 permit

Creating Per IP System Module:
switch# logging filter <NAME> [per-ip] [system-module <module-name>]
logging filter cppm_filter per-ip system-module igmp

Creating a Default Per IP filter:
switch# logging filter <NAME> [per-ip] [default]
logging filter cppm_filter per-ip default

Creating Per IP filter on server:
switch# logging filter <IP> filter [per-ip]
logging 192.123.4.5 filter cppm_filter
```

logging filter enable | disable

Syntax

```
logging filter <NAME> enable | disable

no logging filter <NAME> enable | disable
```

Description

Enables a log filter. In case of global filter, only one filter can be enabled at a time. An enabled filter automatically disables a previously enabled filter.

The `no` form of this command will remove the `enable | disable` filter.

Command context

```
Config
```

Parameters

<NAME>

Specifies the name that identifies the filter.

Examples

```
switch# logging filter <NAME> enable | disable
logging filter cppm_filter enable
```

show logging filter

Syntax

```
Show logging filter [<NAME> | per-ip]
```

Description

This command displays the filter information of syslog server.

Command context

```
config
```

Examples

```
switch (config)# show logging filter

Name           Enabled      IP Address
------------   ----------   ---------------
aaa            Yes          111.222.222.111…
bbb            Yes          1.2.3.4
global         No

switch (config)# show logging filter per-ip
Status and Counters - Per IP Log Filters Information

Name               : Policy1
Severity           : MAJOR
System Module      : snmp dldp
Event Action       : Permit
Event List         : permit
Regular Exp Action : Permit
Regular Exp        : Null

Name               : Policy2
Severity           : MAJOR
System Module      : snmp dldp
Event Action       : Permit
Event List         : 1-45
Regular Exp Action : Permit
Regular Exp        : Null

switch (config)# show logging filter aaa
Status and Counters - Per IP Log Filters Information

Name        : aaa
IP Support  : Yes
URL Support : No

IP Address
--------------
1.2.3.4
4.3.2.1
10.1.2.3
20.1.2.3
40.1.2.3
```

```
URL
-----
https://ARUBA-CENTRAL-URL1:8080/

Type            Value           Action
--------------  --------------  --------------
Severity        major           Default
```

show syslog configuration

Syntax

```
show syslog configuration
```

Description

This command display the syslog server configuration. If the server is associated with the Per IP filter, then the show command will display the Per IP as "Yes".

Command context

```
config
```

Examples

```
switch (config)# show syslog configuration

Syslog Facility            : user
Syslog Severity            : debug
Syslog System Module       : all-pass
Syslog Priority Description :

Syslog Server Details:

Syslog Server Address     L4     Port    Syslog Control Descr    PerIp
----------------------    ----   ------  ----------------------  -----
1.2.3.4                   UDP    514                             No
2.3.4.5                   UDP    2500                            Yes
```

Overview

Why use port or user-based access control?

Local Area Networks are often deployed in a way that allows unauthorized clients to attach to network devices, or allows unauthorized users to get access to unattended clients on a network. Also, the use of DHCP services and zero configuration make access to networking services easily available. This exposes the network to unauthorized use and malicious attacks. While access to the network should be made easy, uncontrolled and unauthorized access is usually not desirable. 802.1X simplifies security management by providing access control along with the ability to control user profiles from up to three RADIUS servers while allowing a given user to use the same entering valid user credentials for access from multiple points within the network.

General features

802.1X on the switches covered in this guide includes the following:

- Switch operation as both an authenticator (for supplicants having a point-to-point connection to the switch) and as a supplicant for point-to-point connections to other 802.1X-aware switches.

    ◦ Authentication of 802.1X access using a RADIUS server and either the EAP or CHAP protocol.

    ◦ Provision for enabling clients that do not have 802.1 supplicant software to use the switch as a path for downloading the software and initiating the authentication process (802.1X Open VLAN mode).

    ◦ User-Based access control option with support for up to 32 authenticated clients per-port.

    ◦ Port-Based access control option allowing authentication by a single client to open the port. This option does not force a client limit and, on a port opened by an authenticated client, allows unlimited client access without requiring further authentication.

    ◦ Supplicant implementation using CHAP authentication and independent user credentials on each port.

- The local operator password configured with the `password` command for management access to the switch is no longer accepted as an 802.1X authenticator credential. The `password port-access` command configures the local operator username and password used as 802.1X authentication credentials for access to the switch. The values configured can be stored in a configuration file using the `include-credentials` command. For information about the `password port-access` command, see **General setup procedure for 802.1X access control** on page 524.

- On-demand change of a port's configured VLAN membership status to support the current client session.

- Session accounting with a RADIUS server, including the accounting update interval.

- Use of `show` commands to display session counters.

- Support for concurrent use of 802.1X and either Web authentication or MAC authentication on the same port.

- For unauthenticated clients that do not have the necessary 802.1X supplicant software (or for other reasons related to unauthenticated clients), there is the option to configure an Unauthorized-Client VLAN. This mode

allows you to assign unauthenticated clients to an isolated VLAN through which you can provide the necessary supplicant software and/or other services you want to extend to these clients.

User authentication methods

The switch offers two methods for using 802.1X access control. Generally, the "Port Based" method supports one 802.1X-authenticated client on a port, which opens the port to an unlimited number of clients. The "User-Based" method supports up to 32 802.1X-authenticated clients on a port. In both cases, there are operating details to be aware of that can influence your choice of methods.

802.1X user-based access control

802.1X operation with access control on a per-user basis provides client-level security that allows LAN access to individual 802.1X clients (up to 32 per port), where each client gains access to the LAN by entering valid user credentials. This operation improves security by opening a given port only to individually authenticated clients, while simultaneously blocking access to the same port for clients that cannot be authenticated. All sessions must use the same untagged VLAN. Also, an authenticated client can use any tagged VLAN memberships statically configured on the port, provided the client is configured to use the tagged VLAN memberships available on the port. Note: The session total includes any sessions begun by the Web Authentication or MAC Authentication features covered in **Option for authenticator ports: configure port-security to allow only 802.1X-authenticated devices** on page 548.

802.1X port-based access control

802.1X port-based access control provides port-level security that allows LAN access only on ports where a single 802.1X-capable client (supplicant) has entered authorized RADIUS user credentials. For reasons outlined below, this option is recommended for applications where only one client at a time can connect to the port. Using this option, the port processes all IP traffic as if it comes from the same client. Thus, in a topology where multiple clients can connect to the same port at the same time:

- If the first client authenticates and opens the port, and then another client authenticates, the port responds as if the original client has initiated a reauthentication. With multiple clients authenticating on the port, the RADIUS configuration response to the latest client authentication replaces any other configuration from an earlier client authentication. If all clients use the same configuration this should not be a problem. But if the RADIUS server responds with different configurations for different clients, then the last client authenticated will effectively lock out any previously authenticated client. When any client to authenticate closes its session, the port will also close and remain so until another client successfully authenticates.

- The most recent client authentication determines the untagged VLAN membership for the port. Also, any client able to use the port can access any tagged VLAN memberships statically configured on the port, provided the client is configured to use the available, tagged VLAN memberships.

- If the first client authenticates and opens the port, and then one or more other clients connect without trying to authenticate, then the port configuration as determined by the original RADIUS response remains unchanged and all such clients will have the same access as the authenticated client. When the authenticated client closes the session, the port will also be closed to any other, unauthenticated clients that may have also been using the port.

This operation unblocks the port while an authenticated client session is in progress. In topologies where simultaneous, multiple client access is possible this can allow unauthorized and unauthenticated access by another client while an authenticated client is using the port. If you want to allow only authenticated clients on the port, then user-based access control should be used instead of port-based access control. Using the user-based method enables you to specify up to 32 authenticated clients. See **802.1X user-based access control** on page 519.

Port-Based 802.1X can operate concurrently with Web-Authentication or MAC-Authentication on the same port. However, this is not a commonly used application and is not generally recommended. For more information, see **Operating rules and notes**.

Alternative to using a RADIUS server

Note that you can also configure 802.1X for authentication through the switch local username and password instead of a RADIUS server, but doing so increases the administrative burden, decentralizes user credential administration, and reduces security by limiting authentication to one operator password set for all users.

Accounting

The switches covered in this guide also provide RADIUS Network accounting for 802.1X access. See **RADIUS Authentication, Authorization, and Accounting**.

General 802.1X authenticator operation

This operation provides security on a point-to-point link between a client and the switch, where both devices are 802.1X-aware. If you expect desirable clients that do not have the necessary 802.1X supplicant software, you can provide a path for downloading such software by using the 802.1X Open VLAN mode, see **802.1X Open VLAN mode** on page 535.

Example of the authentication process

Suppose that you have configured a port on the switch for 802.1X authentication operation, which blocks access to the LAN through that port. If you then connect an 802.1X-aware client (supplicant) to the port and attempt to log on:

1. The switch responds with an identity request.

2. The client responds with a user name that uniquely defines this request for the client.

3. The switch responds in one of the following ways:

   • If 802.1X on the switch is configured for RADIUS authentication, the switch then forwards the request to a RADIUS server.

     a. The server responds with an access challenge which the switch forwards to the client.

     b. The client then provides identifying credentials (such as a user certificate), which the switch forwards to the RADIUS server.

     c. The RADIUS server then checks the credentials provided by the client.

     d. If the client is successfully authenticated and authorized to connect to the network, then the switch allows access to the client. Otherwise, access is denied and the port remains blocked for that client.

   • If 802.1X on the switch is configured for local authentication, then:

     a. The switch compares the client credentials to the username and password configured in the switch (operator level).

     b. If the client is successfully authenticated and authorized to connect to the network, then the server notifies the switch to allow access to the client. Otherwise, access is denied and the port remains blocked.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

> **NOTE**
>
> HPE switches use either 802.1X port-based authentication or 802.1X user-based authentication. For more information, see **User authentication methods**.

VLAN membership priority

Following client authentication, an 802.1X port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. The port also becomes an untagged member of one VLAN according to the following order of options:

Procedure

1. 1st Priority: The port joins a VLAN to which it has been assigned by a RADIUS server during client authentication.

2. 2nd Priority: If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the VLAN entered in the port's 802.1X configuration as an Authorized-Client VLAN, if configured.

3. 3rd Priority: If the port does not have an Authorized-Client VLAN configured, but does have a static, untagged VLAN membership in its configuration, then the switch assigns the port to this VLAN.

A port assigned to a VLAN by an Authorized-Client VLAN configuration (or a RADIUS server) will be an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN.

NOTE

On switches, using the same port for both RADIUS-assigned clients and clients using a configured, Authorized-Client VLAN is not recommended. Doing so can result in authenticated clients with mutually exclusive VLAN priorities, meaning some authenticated clients can be denied access to the port.

Figure 128: *Priority of VLAN assignment for an authenticated client*

General operating rules and notes

- In the user-based mode, when there is an authenticated client on a port, the following traffic movement is allowed:

  ◦ Multicast and broadcast traffic

  ◦ Unicast traffic to authenticated clients

  ◦ All traffic from authenticated clients.

- When a port on the switch is configured as either an authenticator or supplicant and is connected to another device, rebooting the switch causes a reauthentication of the link.

- Using user-based 802.1X authentication, when a port on the switch is configured as an authenticator the port allows only authenticated clients up to the currently configured client limit. For clients without proper 802.1X supplicant software, the optional 802.1X Open VLAN mode can be used to open a path for downloading 802.1X supplicant software to a client or to provide other services for unauthenticated clients. See **802.1X Open VLAN mode** on page 535.

- Using port-based 802.1X authentication when a port on the switch is configured as an authenticator, one authenticated client opens the port. Other clients not running an 802.1X supplicant application can have access to the switch and network through the opened port. If another client uses an 802.1X supplicant application to access the opened port, reauthentication occurs using the RADIUS configuration response for the latest client to authenticate. To control access by all clients, use the user-based method.

- Where a switch port is configured with user-based authentication to accept multiple 802.1X (and/or web- or MAC authentication) client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session. Thus, on a port where one or more authenticated client sessions are already running, all such clients will be on the same untagged VLAN. If a RADIUS server subsequently authenticates a new client, but attempts to reassign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail. For more on this topic, see **802.1X Open VLAN mode** on page 535.

  NOTE: If the port is statically configured with any tagged VLAN memberships, any authenticated client configured to use these tagged VLANs will have access to them.

- You can configure a port as both an 802.1X authenticator and an 802.1X supplicant.

- On a port configured for 802.1X with RADIUS authentication, if the RADIUS server specifies a VLAN for the supplicant and the port is a trunk member, the port will be blocked. If the port is later removed from the trunk, the port will allow authentication of the supplicant. Similarly, if the supplicant is authenticated and later the port becomes a trunk member, the port will be blocked. If the port is removed from the trunk, it will allow the supplicant to reauthenticate.

- If a client already has access to a switch port when you configure the port for 802.1X authenticator operation, the port will block the client from further network access until it can be authenticated.

- Meshing is not supported on ports configured for 802.1X port-access security.

- A port can be configured as an authenticator or an 802.1X supplicant, or both. Some configuration instances block traffic flow or allow traffic to flow without authentication. See **Configuring switch ports to operate as supplicants for 802.1X connections to other switches** on page 549.

- To help maintain security, 802.1X and LACP cannot both be enabled on the same port. If you try to configure 802.1X on a port already configured for LACP (or the reverse), you will see a message similar to the following:

```
Error configuring port X: LACP and 802.1X cannot be run together.
```

Applying web-based authentication or MAC authentication concurrently with port-based 802.1X authentication:

While 802.1X port-based access control can operate concurrently with Web Authentication or MAC Authentication, port-based access control is subordinate to web-based authentication and MAC authentication operation. If 802.1X operates in port-based mode and MAC or web-based authentication is enabled on the same port, any 802.1X authentication has no effect on the ability of a client to access the controlled port. That is, the client access will be denied until the client authenticates through web-based authentication or MAC authentication on the port. Note: A client authenticating with port-based 802.1X does not open the port in the same way that it does when web-based authentication or MAC authentication are not enabled. Any unauthenticated client attempting to access the port after another client authenticates with port-based 802.1X still has to authenticate through web or MAC.

General setup procedure for 802.1X access control

Do these steps before you configure 802.1X operation.

Procedure

1. Configure a local username and password on the switch for both the operator (login) and manager (enable) access levels. (While this may or may not be required for your 802.1X configuration, Hewlett Packard Enterprise recommends that you use a local username and password pair at least until your other security measures are in place.)

2. Enable `include-credentials`. The `port-access` option is available only if `include-credentials` is enabled. See **MAC authentication**.

   For switches covered in this guide, the local operator password configured with the password command is not accepted as an 802.1X authenticator credential. The port-access command is used to configure the operator username and password that are used as 802.1X credentials for network access to the switch. 802.1X network access is not allowed unless a password has been configured using the `password port-access` command.

   ```
   password port-access [user-name <name>]<password>
   ```

   Configures the operator username and password used to access the network through 802.1X authentication.

   user-name <name>

   > operator username (text string) used only for local authentication of 802.1X clients. This value is different from the local operator username configured with the `password` command for management access.

   <password>

   > operator password (text string) used only for local authentication of 802.1X clients. This value is different from the local operator password configured with the `password` command for management access.

   The password port-access command

   ```
   HP Switch(config)# password port-access user-name Jim secret3
   ```

   You can save the port-access password for 802.1X authentication in the configuration file by using the `include-credentials` command. For more information, see **Saving security credentials in a config file** on page 52.

3. Determine the switch ports that you want to configure as authenticators and/or supplicants, and disable LACP on these ports.

   To display the current configuration of 802.1X, Web-based, and MAC authentication on all switch ports, enter the `show port-access config` command.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Output for the show port-access config command

```
HP Switch (config)# show port-access config

Port-access authenticator activated [No] : No
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

     Supplicant Authenticator Web-Auth Mac-Auth LMA-Auth Ctrl  Mixed Speed
Port Enabled    Enabled      Enabled  Enabled  Enabled  Dir   Mode  VSA   MBV
---- --------   ------------ -------- -------- -------- ----- ----  ----  ---
C1   No         Yes          No       No       No       In    No    Yes   Yes
C2   No         Yes          No       No       No       Both  Yes   Yes   Yes
C3   No         Yes          No       No       No       Both  No    No    Yes
C4   No         Yes          No       No       Yes      Both  No    Yes   Yes
...
```

4. Determine whether to use user-based access control, see **802.1X user-based access control** on page 519 or port-based access control, see **802.1X port-based access control** on page 519.

5. Determine whether to use the optional 802.1X Open VLAN mode for clients that are not 802.1X-aware; that is, for clients that are not running 802.1X supplicant software. (This will require you to provide downloadable software that the client can use to enable an authentication session.) See **802.1X Open VLAN mode** on page 535.

6. For any port you want to operate as a supplicant, determine the user credentials. You can either use the same credentials for each port or use unique credentials for individual ports or subgroups of ports. (This can also be the same local username/password pair that you assign to the switch.)

7. Unless you are using only the switch's local username and password for 802.1X authentication, configure at least one RADIUS server to authenticate access requests coming through the ports on the switch from external supplicants (including switch ports operating as 802.1X supplicants). You can use up to three RADIUS servers for authentication; one primary and two backups. See the documentation provided with your RADIUS application.

Overview: configuring 802.1X authentication on the switch

This section outlines the steps for configuring 802.1X on the switch. For detailed information on each step, see the following:

- **802.1X user-based access control** on page 519

- **802.1X port-based access control** on page 519

- **Configuring switch ports to operate as supplicants for 802.1X connections to other switches** on page 549

1. Enable 802.1X user-based or port-based authentication on the individual ports you want to serve as authenticators. On the ports you will use as authenticators, either accept the default 802.1X settings or change them, as necessary. Note that, by default, the port-control parameter is set to `auto` for all ports on the switch. This requires a client to support 802.1X authentication and to provide valid credentials to get network access. See **Enable 802.1X authentication on selected ports** on page 526.

2. If you want to provide a path for clients without 802.1X supplicant software to download the software so that they can initiate an authentication session, enable the 802.1X Open VLAN mode on the ports you want to support this feature. See **802.1X Open VLAN mode** on page 535.

3. Configure the 802.1X authentication type. Options include:

- Local operator username and password (using the `password port-access` command).

- EAP RADIUS: This option requires your RADIUS server application to support EAP authentication for 802.1X

- CHAP (MD5) RADIUS: This option requires your RADIUS server application to support CHAP (MD5) authentication. See **Configuring the 802.1X authentication method** on page 530.

4. If you select either `eap-radius` or `chap-radius` for step 3, use the `radius host` command to configure up to three RADIUS server IP address(es) on the switch. See **Enter the RADIUS host IP address(es)** on page 531.

5. Enable 802.1X authentication on the switch. See **Enable 802.1X authentication on selected ports** on page 526.

6. Test both the authorized and unauthorized access to your system to ensure that the 802.1X authentication works properly on the ports you have configured for port-access.

> **NOTE**
> If you want to implement the optional port security feature (step 7) on the switch, you should first ensure that the ports you have configured as 802.1X authenticators operate as expected.

7. If you are using Port Security on the switch, configure the switch to allow only 802.1X access on ports configured for 802.1X operation, and (if desired) the action to take if an unauthorized device attempts access through an 802.1X port. See **Port-Security** on page 549.

8. If you want a port on the switch to operate as a supplicant on a port operating as an 802.1X authenticator on another device, then configure the supplicant operation. (See **Configuring switch ports to operate as supplicants for 802.1X connections to other switches** on page 549.

Configuring switch ports as 802.1X authenticators

Enable 802.1X authentication on selected ports

This task configures the individual ports you want to operate as 802.1X authenticators for point-to-point links to 802.1X-aware clients or switches, and consists of two steps:

1. Enable the selected ports as authenticators.

2. Specify either user-based or port-based 802.1X authentication.

(Actual 802.1X operation does not commence until you activate 802.1X authentication on the switch.)

> **NOTE**
> If you enable 802.1X authentication on a port, the switch automatically disables LACP on that port. However, if the port is already operating in an LACP trunk, you must remove the port from the trunk before you can configure it for 802.1X authentication.

Enable the selected ports as authenticators and enable the (default) port-based authentication

Syntax:

```
[no] aaa port-access authenticator <port-list>
```

Enables specified ports to operate as 802.1X authenticators and enables port-based authentication. (To enable user-based authentication, execute this command first, and then execute the client-limit <port-list> version of this

command described in the next section.) The `no` form of the command removes 802.1X authentication from <port-list>. To activate configured 802.1X operation, you must enable 802.1X authentication. See **Enable 802.1X authentication on the switch** on page 532.

Specify user-based authentication or return to port-based authentication

User-based 802.1X authentication:

Syntax:

```
aaa port-access authenticator <port-list> client-limit <1-32>
```

Used after executing `aaa port-access authenticator <port-list>` to convert authentication from port-based to user-based. Specifies user-based 802.1X authentication and the maximum number of 802.1X-authenticated client sessions allowed on each of the ports in `<port-list>`. If a port currently has no authenticated client sessions, the next authenticated client session the port accepts determines the untagged VLAN membership to which the port is assigned during the session. If another client session begins later on the same port while an earlier session is active, the later session will be on the same untagged VLAN membership as the earlier session.

> **NOTE**
>
> The client limit is 256 clients per-port for MAC-auth and Web-auth; the client limit for 802.1X is 32 clients per port. The MAC-auth and Web-auth limit of 256 clients only applies when there are fewer than 16,384 authentication clients on the entire switch. After the limit of 16,384 clients is reached, no additional authentication clients are allowed on any port for any method.

Port-based 802.1X authentication:

Syntax:

```
[no] aaa port-access authenticator <port-list> client-limit
```

Used to convert a port from user-based authentication to port-based authentication, which is the default setting for ports on which authentication is enabled. (Executing `aaa port-access authenticator <port-list>` enables 802.1X authentication on `<port-list>` and enables port-based authentication.) If a port currently has no authenticated client sessions, the next authenticated client session the port accepts determines the untagged VLAN membership to which the port is assigned during the session. If another authenticated client session begins later on the same port while an earlier session is active, the later session replaces the currently active session and will be on the untagged VLAN membership specified by the RADIUS server for the later session.

Configuring user-based 802.1X authentication on page 527 Enables ports 10-12 to operate as authenticators, and then configures the ports for user-based authentication.

Configuring user-based 802.1X authentication

```
HP Switch(config)# aaa port-access authenticator 10-12
HP Switch(config)# aaa port-access authenticator 10-12 client-limit 4
```

Configuring port-based 802.1X authentication

```
HP Switch(config)# aaa port-access authenticator 13-15
HP Switch(config)# no aaa port-access authenticator 13-15 client-limit
```

Reconfigure settings for port-access

The commands in this section are initially set by default and can be reconfigured as needed.

Syntax

```
aaa port-access authenticator <port-list> [<item>]
```

Parameters

<port-list>

Specifies the ports acted on by this command.

<item>

Specifies one of these items:

`auth-vid` *<vlan-id>*

Configures an existing, static VLAN to be the Authorized-Client VLAN.

`clear-statistics`

Clears authenticator statistics counters.

`client-limit` *<1-32>*

Set the maximum number of clients to allow on the port. With no client limit, authentication happens in port-based mode, otherwise it happens in client-based mode.

`control {authorized | auto | unauthorized}`

Controls authentication mode on the specified port.

`authorized`

Also termed "Force Authorized". Gives access to a device connected to the port. In this case, the device does not have to provide 802.1X credentials or support 802.1X authentication. (You can still configure console, Telnet, or SSH security on the port.)

`auto`

This is the default. The device connected to the port must support 802.1X authentication and provide valid credentials to get network access. (Optional: You can use the Open VLAN mode to provide a path for clients without 802.1X supplicant software to download this software and begin the authentication process.)

`initialize`

On the specified ports, blocks inbound and outbound traffic and restarts the 802.1X authentication process. This happens only on ports configured with `control auto` and actively operating as 802.1X authenticators.

NOTE

If a specified port is configured with `control authorized` and `port-security`, and the port has learned an authorized address, the port will remove this address and learn a new one from the first packet it receives.

`logoff-period` *<1-999999999>*

Configures the time the switch waits for client activity before removing an inactive client from the port. (Default: 300 seconds)

`max-requests` *<1-10>*

Sets the number of authentication attempts that must time out before authentication fails and the authentication session ends. If you are using the Local authentication option, or are using RADIUS authentication with only one host server, the switch will not start another session until a client tries a new access attempt. If you are using RADIUS authentication with two or three host servers, the switch will open a session with each server, in turn, until authentication occurs or there are no more servers to try. During the `quiet-period`, if any, you cannot reconfigure this parameter. (Default: 2)

quiet-period *<0-65535>*

Sets the period during which the port does not try to acquire a supplicant. The period begins after the last attempt authorized by the `max-requests` parameter fails. (Default: 60 seconds)

reauth-period *<0-9999999>*

Sets the time after which clients connected must be reauthenticated. When the timeout is set to 0, the reauthentication is disabled (Default: 0 second)

reauthenticate

Forces reauthentication (unless the authenticator is in 'HELD' state).

server-timeout *<1-300>*

Sets the time the switch waits for a server response to an authentication request. If there is no response within the configured time frame, the switch assumes that the authentication attempt has timed out. Depending on the current `max-requests` setting, the switch will either send a new request to the server or end the authentication session. (Default: 30 seconds)

supplicant-timeout *<1-300>*

Sets the time the switch waits for a supplicant response to an EAP request. If the supplicant does not respond within the configured time frame, the session times out. (Default: 30 seconds)

tx-period *<0-65535>*

Sets the time the port waits to retransmit the next EAPOL PDU during an authentication session. (Default: 30 seconds)

unauth-period *<0-255>*

Specifies a delay in seconds for placing a port on the Unauthorized-Client VLAN. This delay allows more time for a client with 802.1X supplicant capability to initiate an authentication session. If a connected client does not initiate a session before the timer expires, the port is assigned to the Unauthenticated-Client VLAN. (Default: 0 seconds)

unauth-vid *<vlan-id>*

Configures an existing static VLAN to be the Unauthorized-Client VLAN. This enables you to provide a path for clients without supplicant software to download the software and begin an authentication session.

About `tx-period` and identity request triggers

The actual period between EAPOL PDU retransmits is influenced by the state of authenticating or connecting clients. The trigger for EAPOL identity requests depends on the following:

- The `tx-period` configured.

- The number of clients connected to the switch and the state of the clients.

If there is one client connected and:

- The client is in the authenticated state, `tx-period` expiry will not trigger an identity request.

- The client is in the connecting state, `tx-period` expiry will trigger an identity request to the client MAC.

- The client MAC address is not known, then upon `tx-period` expiry, the switch will send the next identity request to the well-known client MAC (EAPOL group multicast address).

If there are two clients connected, and:

- One client is in the connecting state, `tx-period` expiry will trigger an identity request to the client MAC. In this case, it is assumed that there is no traffic from the second client and that the switch is not aware of the second client.

- Two clients are in the connecting state (and if the logoff period does not expire before `tx-period` expiry), then each client will maintain separate timers and identity requests will be sent at regular intervals.

- One client is in the authenticated state and the second client is in the connecting state, then the identity request will be triggered upon expiry of any client timer. In this case, if the first client timer expires, then the first client MAC will send an identity request to the second client MAC. Therefore, the identity request send interval may be different than what is set for `tx-period`.

- Two clients are in the authenticated state, upon `tx-period` expiry, the switch will not send an identity request.

- Both clients are not sending any traffic, the switch will send identity requests to the well-known client MAC (EAPOL group multicast address).

Configuring the 802.1X authentication method

This task specifies how the switch authenticates the credentials provided by a supplicant connected to a switch port configured as an 802.1X authenticator.

Syntax:

```
aaa authentication port-access <chap-radius|eap-radius|local>
```

Configures `local`, `chap-radius`, or `eap-radius` as the primary password authentication method for port-access. The default primary authentication is `local`. (See the documentation for your RADIUS server application.)

For switches covered in this guide, you must use the `password port-access` command to configure the operator username and password for 802.1X access.

```
[<none|authorized>]
```

Provides options for secondary authentication. The `none` option specifies that a backup authentication method is not used. The `authorized` option allows access without authentication. (default: `none`).

To enable the switch to perform 802.1X authentication using one or more EAP-capable RADIUS servers:

Figure 129: *802.1X (port-access) authentication*

```
HP Switch(config)# aaa authentication port-access eap-radius
HP Switch(config)# show authentication

 Status and Counters - Authentication Information

  Login Attempts : 3
  Respect Privilege : Disabled

              | Login          Login          Login
  Access Task | Primary        Server Group   Secondary
  ----------- + ----------    ------------    ----------
  Console     | Local                          None
  Telnet      | Local                          None          802.1X (Port-Access)
  Port-Access | EapRadius   ◄──────────────────────────────  configured for EAP-
  Webui       | Local                          None          RADIUS authentication.
  SSH         | Local                          None
  Web-Auth    | ChapRadius                     None
  MAC-Auth    | ChapRadius                     None

              | Enable          Enable         Enable
  Access Task | Primary        Server Group   Secondary
  ----------- + ----------    ------------    ----------
  Console     | Local                          None
  Telnet      | Local                          None
  Webui       | Local                          None
  SSH         | Local                          None
```

Enter the RADIUS host IP address(es)

If you select either `eap-radius` or `chap-radius` for the authentication method, configure the switch to use 1, 2, or 3 RADIUS servers for authentication. The following syntax shows the basic commands. For coverage of all commands related to RADIUS server configuration, see **RADIUS Authentication, Authorization, and Accounting** on page 194.

Syntax:

`radius host <ip-address>`

Adds a server to the RADIUS configuration.

`[key <server-specific key-string>]`

Optional. Specifies an encryption key for use during authentication (or accounting) sessions with the specified server. This key must match the key used on the RADIUS server. Use this option only if the specified server requires a different key than configured for the global encryption key. The tilde (~) character is allowed in the string. It is not backward compatible; the "~" character is lost if you use a software version that does not support the "~" character.

Syntax:

```
radius-server key <global key-string>
```

Specifies the global encryption key the switch uses for sessions with servers for which the switch does not have a server-specific key. This key is optional if all RADIUS server addresses configured in the switch include a server-specific encryption key. The tilde (~) character is allowed in the string, for example, `radius-server key hp~switch`. It is not backward compatible; the "~" character is lost if you use a software version that does not support the "~" character.

Default: Null

The `no` form of the command removes the global encryption key.

Enable 802.1X authentication on the switch

After configuring 802.1X authentication as described in the preceding four sections, activate it with this command:

Syntax:

```
aaa port-access authenticator active
```

Activates 802.1X port-access on ports you have configured as authenticators.

Reset authenticator operation (optional)

While 802.1X authentication is operating, you can use the following `aaa port-access authenticator` commands to reset 802.1X authentication and statistics on specified ports.

Syntax:

```
aaa port-access authenticator <port-list>
```

```
[initialize]
```

On the specified ports, blocks inbound and outbound traffic and restarts the 802.1X authentication process. This happens only on ports configured with `control auto` and actively operating as 802.1X authenticators.

```
[reauthenticate]
```

On the specified ports, forces reauthentication (unless the authenticator is in "HELD" state).

```
[clear-statistics]
```

On the specified ports, clears authenticator statistics counters.

Configure 802.1X controlled direction (optional)

After you enable 802.1X authentication on specified ports, you can use the `aaa port-access controlled-direction` command to configure how a port transmits traffic before it successfully authenticates a client and enters the authenticated state.

As documented in the IEEE 802.1X standard, an 802.1X-aware port that is unauthenticated can control traffic in either of the following ways:

- In both ingress and egress directions by disabling both the reception of incoming frames and transmission of outgoing frames

- Only in the ingress direction by disabling only the reception of incoming frames.

Prerequisite:

As documented in the IEEE 802.1X standard, the disabling of incoming traffic and transmission of outgoing traffic on an 802.1X-aware egress port in an unauthenticated state (using the `aaa port-access controlled-direction in` command) is supported only if:

- The port is configured as an edge port in the network using the `spanning-tree edge-port` command.

- The 802.1s Multiple Spanning Tree Protocol (MSTP) or 802.1w Rapid Spanning Tree Protocol (RSTP) is enabled on the switch. MSTP and RSTP improve resource utilization while maintaining a loop-free network.

For information on how to configure the prerequisites for using the `aaa port-access controlled-direction in` command, see "Multiple Instance Spanning-Tree Operation" in the advanced traffic management guide.

Syntax:

`aaa port-access <`*`port-list`*`> controlled-direction <both|in>`

`both` (default): Incoming and outgoing traffic is blocked on an 802.1X-aware port before authentication occurs.

`in`: Incoming traffic is blocked on an 802.1X-aware port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on unauthenticated 802.1X-aware ports.

Wake-on-LAN Traffic

The Wake-on-LAN feature is used by network administrators to remotely power on a sleeping workstation (for example, during early morning hours to perform routine maintenance operations, such as patch management and software updates).

The `aaa port-access controlled-direction in` command allows Wake-on-LAN traffic to be transmitted on an 802.1X-aware egress port that has not yet transitioned to the 802.1X authenticated state; the `controlled-direction both` setting prevents Wake-on-LAN traffic to be transmitted on an 802.1X-aware egress port until authentication occurs.

NOTE: Although the `controlled-direction in` setting allows Wake-on-LAN traffic to traverse the switch through unauthenticated 802.1X-aware egress ports, it does not guarantee that the Wake-on-LAN packets will arrive at their destination. For example, firewall rules on other network devices and VLAN rules may prevent these packets from traversing the network.

Operating notes

- Using the `aaa port-access controlled-direction in` command, you can enable the transmission of Wake-on-LAN traffic on unauthenticated egress ports that are configured for any of the following port-based security features:

  - 802.1X authentication

  - MAC authentication

  - Web authentication

Because a port can be configured for more than one type of authentication to protect the switch from unauthorized access, the last setting you configure with the `aaa port-access controlled-direction` command is applied to all authentication methods configured on the switch. See **Web and MAC Authentication** on page 99.

- To display the currently configured 802.1X Controlled Direction value, enter the `show port-access authenticator config` command.

- When an 802.1X-authenticated port is configured with the `controlled-direction in` setting, eavesdrop prevention is not supported on the port.

Configuring 802.1X controlled directions on page 534 shows how to enable the transmission of Wake-on-LAN traffic in the egress direction on an 802.1X-aware port before it transitions to the 802.1X authenticated state and successfully authenticates a client device.

Configuring 802.1X controlled directions

```
HP Switch(config)# aaa port-access authenticator a10
HP Switch(config)# aaa authentication port-access eap-radius
HP Switch(config)# aaa port-access authenticator active
HP Switch(config)# aaa port-access a10 controlled-direction in
```

Unauthenticated VLAN access (guest VLAN access)

When a PC is connected through an IP phone to a switch port that has been authorized using 802.1X or Web/MAC authentication, the IP phone is authenticated using client-based 802.1X or Web/MAC authentication and has access to secure, tagged VLANs on the port. If the PC is unauthenticated, it needs to have access to the insecure guest VLAN (unauthenticated VLAN) that has been configured for 802.1X or Web/MAC authentication. 802.1X and Web/MAC authentication normally do not allow authenticated clients (the phone) and unauthenticated clients (the PC) on the same port.

Mixed port access mode allows 802.1X and Web/MAC authenticated and unauthenticated clients on the same port when the guest VLAN is the same as the port's current untagged authenticated VLAN for authenticated clients, or when none of the authenticated clients are authorized on the untagged authenticated VLAN. Instead of having just one client per port, multiple clients can use the guest VLAN.

Authenticated clients always have precedence over guests (unauthenticated clients) if access to a client's untagged VLAN requires removal of a guest VLAN from the port. If an authenticated client becomes authorized on its untagged VLAN as the result of initial authentication or because of an untagged packet from the client, then all 802.1X or Web/MAC authenticated guests are removed from the port and the port becomes an untagged member of the client's untagged VLAN.

Characteristics of mixed port access mode

- The port keeps tagged VLAN assignments continuously.

- The port sends broadcast traffic from the VLANs even when there are only guests authorized on the port.

- Guests cannot be authorized on any tagged VLANs.

- Guests can use the same bandwidth, rate limits and QoS settings that may be assigned for authenticated clients on the port (via RADIUS attributes).

- When no authenticated clients are authorized on the untagged authenticated VLAN, the port becomes an untagged member of the guest VLAN for as long as no untagged packets are received from any authenticated clients on the port.

- New guest authorizations are not allowed on the port if at least one authenticated client is authorized on its untagged VLAN and the guest VLAN is not the same as the authenticated client's untagged VLAN.

> **NOTE**
>
> If you disable mixed port access mode, this does not automatically remove guests that have already been authorized on a port where an authenticated client exists. New guests are not allowed after the change, but the existing authorized guests will still be authorized on the port until they are removed by a new authentication, an untagged authorization, a port state change, and so on.

Configuring mixed port access mode

Syntax:

`[no] aaa port-access <port-list> mixed`

Enables or disables guests on ports with authenticated clients.

Default: Disabled; guests do not have access

Configuring mixed port access mode

```
HP Switch(config)# aaa port-access 6 mixed
```

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

http://www.hpe.com/support/selfrepair

802.1X Open VLAN mode

Introduction

This section describes how to use the 802.1X Open VLAN mode to provide a path for clients that need to acquire 802.1X supplicant software before proceeding with the authentication process. The Open VLAN mode involves options for configuring unauthorized-client and authorized-client VLANs on ports configured as 802.1X authenticators.

Configuring the 802.1X Open VLAN mode on a port changes how the port responds when it detects a new client. In earlier releases, a "friendly" client computer not running 802.1X supplicant software could not be authenticated on a port protected by 802.1X access security. As a result, the port would become blocked and the client could not access the network. This prevented the client from:

- Acquiring IP addressing from a DHCP server

- Downloading the 802.1X supplicant software necessary for an authentication session

The 802.1X Open VLAN mode solves this problem by temporarily suspending the port's static VLAN memberships and placing the port in a designated Unauthorized-Client VLAN (sometimes termed a guest VLAN). In this state the client can proceed with initialization services, such as acquiring IP addressing and 802.1X client software, and starting the authentication process.

NOTE	On ports configured to allow multiple sessions using 802.1X user-based access control, all clients must use the same untagged VLAN. On a given port where there are no currently active, authenticated clients, the first authenticated client determines the untagged VLAN in which the port will operate for all subsequent, overlapping client sessions.
	If the switch operates in an environment where some valid clients will not be running 802.1X supplicant software and need to download it from your network. Then, because such clients would need to use the Unauthorized-Client VLAN and authenticated clients would be using a different VLAN (for security reasons), allowing multiple clients on an 802.1X port can result in blocking some or all clients needing to use the Unauthorized-Client VLAN.
	On ports configured for port-based 802.1X access control, if multiple clients try to authenticate on the same port, the most recently authenticated client determines the untagged VLAN membership for that port. Clients that connect without trying to authenticate will have access to the untagged VLAN membership that is currently assigned to the port.

VLAN membership priorities

Following client authentication, an 802.1X port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. The port also becomes an untagged member of one VLAN according to the following order of options:

- **1st Priority:**

  The port joins a VLAN to which it has been assigned by a RADIUS server during client authentication.

- **2nd Priority:**

  If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the VLAN entered in the port's 802.1X configuration as an Authorized-Client VLAN, if configured.

- **3rd Priority:**

  If the port does not have an Authorized-Client VLAN configured, but does have a static, untagged VLAN membership in its configuration, then the switch assigns the port to this VLAN.

A port assigned to a VLAN by an Authorized-Client VLAN configuration (or a RADIUS server) will be an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN.

NOTE	After client authentication, the port resumes membership in any tagged VLANs for which it is configured. If the port is a tagged member of a VLAN used for 1 or 2 listed above, then it also operates as an untagged member of that VLAN while the client is connected. When the client disconnects, the port reverts to tagged membership in the VLAN.

Use models for 802.1X Open VLAN modes

You can apply the 802.1X Open VLAN mode in more than one way. Depending on your use, you will need to create one or two static VLANs on the switch for exclusive use by per-port 802.1X Open VLAN mode authentication:

- **Unauthorized-Client VLAN:**

Configure this VLAN when unauthenticated, friendly clients will need access to some services before being authenticated or instead of being authenticated.

- **Authorized-Client VLAN:**

Configure this VLAN for authenticated clients when the port is not statically configured as an untagged member of a VLAN you want clients to use, or when the port is statically configured as an untagged member of a VLAN you do not want clients to use. (A port can be configured as untagged on only one port-based VLAN. When an Authorized-Client VLAN is configured, it will always be untagged and will block the port from using a statically configured, untagged membership in another VLAN.) After client authentication, the port returns to membership in any tagged VLANs for which it is configured. See **VLAN membership priorities**.

Table 41: *802.1X per-port configuration*

802.1X per-port configuration	Port response
No Open VLAN mode:	The port automatically blocks a client that cannot initiate an authentication session.
Open VLAN mode with both of the following configured:	**Unauthorized-client VLAN:** - When the port detects a client without 802.1X supplicant capability, it automatically becomes an untagged member of this VLAN. If you previously configured the port as a static, tagged member of the VLAN, membership temporarily changes to untagged while the client remains unauthenticated. - If the port already has a statically configured, untagged membership in another VLAN, then the port temporarily closes access to this other VLAN while in the Unauthorized-Client VLAN. - To limit security risks, the network services and access available on the Unauthorized-Client VLAN must include only what a client requires to enable an authentication session. If the port is statically configured as a tagged member of any other VLANs, access to these VLANs is blocked while the port is a member of the Unauthorized-Client VLAN. **NOTE** **For a port configured to allow multiple client sessions:** If any previously authenticated clients are using a port assigned to a VLAN other than the Unauthorized-Client VLAN, then a later client that is not running 802.1X supplicant software is blocked on the port until all other, authenticated clients on the port have disconnected.

Table Continued

802.1X per-port configuration	Port response

	Authorized-client VLAN:
	
	• After client authentication, the port drops membership in the Unauthorized-Client VLAN and becomes an untagged member of this VLAN.
	
	    NOTE: If the client is running an 802.1X supplicant application when the authentication session begins, and is able to authenticate itself before the switch assigns the port to the Unauthorized-Client VLAN, then the port does not become a member of the Unauthorized-Client VLAN. On the switches covered in this guide, you can use the `unauth-period` command to delay moving the port into the Unauthorized-Client VLAN.
	
	    If RADIUS authentication assigns a VLAN and there are no other authenticated clients on the port, then the port becomes a member of the RADIUS-assigned VLAN—instead of the Authorized-Client VLAN—while the client is connected.
	
	• If the port is statically configured as a tagged member of a VLAN, and this VLAN is used as the Authorized-Client VLAN, then the port temporarily becomes an untagged member of this VLAN when the client becomes authenticated.
	
	• If the port is statically configured as a tagged member of a VLAN, the port returns to tagged membership in this VLAN upon successful authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. If the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an untagged member of that VLAN for the duration of the client connection. After the client disconnects, the port returns to tagged membership in that VLAN.

Table Continued

802.1X per-port configuration	Port response
Open VLAN mode with only an unauthorized-client VLAN configured:	• When the port detects a client, it automatically becomes an untagged member of this VLAN. To limit security risks, the network services and access available on this VLAN should include only what a client requires to enable an authentication session. If the port is statically configured as an untagged member of another VLAN, the switch temporarily removes the port from membership in this other VLAN while membership in the Unauthorized-Client VLAN exists. • After the client is authenticated, and if the port is statically configured as an untagged member of another VLAN, the port access to this other VLAN is restored. NOTE: If RADIUS authentication assigns the port to a VLAN, this assignment overrides any statically configured, untagged VLAN membership on the port (while the client is connected). If the port is statically configured as a tagged member of a VLAN, the port returns to tagged membership in this VLAN upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. If the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an untagged member of that VLAN for the duration of the client connection. NOTE: **For a port configured to allow multiple client sessions:** If any previously authenticated clients are using a port assigned to a VLAN other than the Unauthorized-Client VLAN (such as a RADIUS-assigned VLAN), then a later client that is not running 802.1X supplicant software is blocked on the port until all other, authenticated clients on the port have disconnected.
Open VLAN mode with only an authorized-client VLAN configured:	• Port automatically blocks a client that cannot initiate an authentication session. • If the client successfully completes an authentication session, the port becomes an untagged member of this VLAN. • If the port is statically configured as a tagged member of any other VLAN, the port returns to tagged membership in this VLAN upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN. If the port is already configured as a tagged member of a VLAN that RADIUS assigns as an authorized VLAN, then the port becomes an untagged member of that VLAN for the duration of the client connection. After the client disconnects, the port returns to tagged membership in that VLAN.

802.1X per-port configuration	Port response
	An authorized-client VLAN configuration can be overridden by a RADIUS authentication that assigns a VLAN.

Operating rules for authorized and unauthorized-client VLANs

Table 42: *Condition for authorized client and unauthorized client VLANs*

Condition	Rule
Static VLANs used as authorized-client or unauthorized-client VLANs	These must be configured on the switch before you configure an 802.1X authenticator port to use them. (Use the `vlan <vlan-id>` command or the VLAN Menu screen in the Menu interface.)
VLAN assignment received from a RADIUS server	If the RADIUS server specifies a VLAN for an authenticated supplicant connected to an 802.1X authenticator port, this VLAN assignment overrides any Authorized-Client VLAN assignment configured on the authenticator port. This is because membership in both VLANs is untagged, and the switch allows only one untagged, port-based VLAN membership per-port. For example, suppose you configured port A4 to place authenticated supplicants in VLAN 20. If a RADIUS server authenticates supplicant "A" and assigns this supplicant to VLAN 50, then the port can access VLAN 50 as an untagged member while the client session is running. When the client disconnects from the port, then the port drops these assignments and uses the untagged VLAN memberships for which it is statically configured. (After client authentication, the port resumes any tagged VLAN memberships for which it is already configured.

Table Continued

Condition	Rule
Temporary VLAN membership during a client session	• Port membership in a VLAN assigned to operate as the Unauthorized-Client VLAN is temporary, and ends when the client receives authentication or the client disconnects from the port, whichever is first. In the case of the multiple clients allowed on switches covered in this guide, the first client to authenticate determines the untagged VLAN membership for the port until all clients have disconnected. Any other clients that cannot operate in that VLAN are blocked at that point. • Port membership in a VLAN assigned to operate as the Authorized-Client VLAN ends when the client disconnects from the port. If a VLAN assignment from a RADIUS server is used instead, the same rule applies. In the case of the multiple clients allowed on switches, the port maintains the same VLAN as long as there is any authenticated client using the VLAN. When the last client disconnects, then the port reverts to only the VLAN(s) for which it is statically configured as a member.
Effect of unauthorized-client VLAN session on untagged port VLAN membership	• When an unauthenticated client connects to a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Unauthorized-Client VLAN (also untagged). (While the Unauthorized-Client VLAN is in use, the port does not access any other VLANs.) • If the client disconnects, the port leaves the Unauthorized-Client VLAN and re-acquires membership in all the statically configured VLANs to which it belongs. • If the client becomes authenticated, the port leaves the Unauthenticated-Client VLAN and joins the appropriate VLAN. See **VLAN membership priorities** on page 536. • In the case of the multiple clients allowed on switches, if an authenticated client is already using the port for a different VLAN, then any other unauthenticated clients needing to use the Unauthorized-Client VLAN are blocked.

Table Continued

Condition	Rule
Effect of authorized-client VLAN session on untagged port VLAN membership.	• When a client becomes authenticated on a port that is already configured with a static, untagged VLAN, the switch temporarily moves the port to the Authorized-Client VLAN (also untagged). While the Authorized-Client VLAN is in use, the port does not have access to the statically configured, untagged VLAN. • When the authenticated client disconnects, the switch removes the port from the Authorized-Client VLAN and moves it back to the untagged membership in the statically configured VLAN. (After client authentication, the port resumes any tagged VLAN memberships for which it is already configured.) **NOTE** This rule assumes: • No alternate VLAN has been assigned by a RADIUS server. • No other authenticated clients are already using the port.
Multiple authenticator ports using the same unauthorized-client and authorized-client VLANs	You can use the same static VLAN as the Unauthorized-Client VLAN for all 802.1X authenticator ports configured on the switch. Similarly, you can use the same static VLAN as the Authorized-Client VLAN for all 802.1X authenticator ports configured on the switch. **CAUTION** Do not use the same static VLAN for both the unauthorized-client VLAN and the authorized-client VLAN. Using one VLAN for both creates a security risk by defeating the isolation of unauthenticated clients.
Effect of filed client authentication attemptThis rule assumes no other authenticated clients are already using the port on a different VLAN.	When there is an Unauthorized-Client VLAN configured on an 802.1X authenticator port, an unauthorized client connected to the port has access only to the network resources belonging to the Unauthorized- Client VLAN. This access continues until the client disconnects from the port. (If there is no Unauthorized-Client VLAN configured on the authenticator port, the port simply blocks access for any unauthorized client.)

Table Continued

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Condition	Rule
Effect of RADIUS-assigned VLANThis rule assumes no other authenticated clients are already using the port on a different VLAN.	The port joins the RADIUS-assigned VLAN as an untagged member.
IP Addressing for a client connected to a port configured for 802.1X Open VLAN mode	A client can either acquire an IP address from a DHCP server or use a manually configured IP address before connecting to the switch.
802.1X supplicant software for a client connected to a port configured for 802.1X Open VLAN mode	A friendly client, without 802.1X supplicant software, connecting to an authenticator port must be able to download this software from the Unauthorized-Client VLAN before authentication can begin.

Table Continued

Condition	Rule

Switch with a port configured to allow multiple authorized-client sessions	When a new client is authenticated on a given port:
	
	• If no other clients are authenticated on that port, then the port joins one VLAN in the following order of precedence:
	
	  1. A RADIUS-assigned VLAN, if configured.
	
	  2. An Authenticated-Client VLAN, if configured.
	
	  3. A static, port-based VLAN to which the port belongs as an untagged member.
	
	  4. Any VLAN(s) to which the port is configured as a tagged member (provided that the client can operate in that VLAN).
	
	• If another client is already authenticated on the port, then the port is already assigned to a VLAN for the previously-existing client session, and the new client must operate in this same VLAN, regardless of other factors. (This means that a client without 802.1X client authentication software cannot access a configured, Unauthenticated-Client VLAN if another, authenticated client is already using the port.)

Limitation on using an unauthorized-client VLAN on an 802.1X port configured to allow multiple-client access	You can optionally enable switches to allow up to 32 clients per-port. The Unauthorized-Client VLAN feature can operate on an 802.1X-configured port regardless of how many clients the port is configured to support. However, all clients on the same port must operate through the same untagged VLAN membership.(See **MAC-based VLANs** on page 221). This means that any client accessing a given port must be able to authenticate and operate on the same VLAN as any other previously authenticated clients that are currently using the port. Thus, an Unauthorized-Client VLAN configured on a switch port that allows multiple 802.1X clients cannot be used if there is already an authenticated client using the port on another VLAN. Also, a client using the Unauthenticated-Client VLAN will be blocked when another client becomes authenticated on the port. For this reason, the best utilization of the Unauthorized-Client VLAN feature is in instances where only one client is allowed per-port. Otherwise, unauthenticated clients are subject to being blocked at any time by authenticated clients using a different VLAN. (Using the same VLAN for authenticated and unauthenticated clients can create a security risk and is not recommended.)

NOTE

If you use the same VLAN as the Unauthorized-Client VLAN for all authenticator ports, unauthenticated clients on different ports can communicate with each other.

Setting up and configuring 802.1X Open VLAN mode

Preparation:

This section assumes use of both the unauthorized-client and authorized-client VLANs.

Before you configure the 802.1X Open VLAN mode on a port:

- Statically configure an Unauthorized-Client VLAN in the switch. The only ports that should belong to this VLAN are ports offering services and access you want available to unauthenticated clients. (802.1X authenticator ports do not have to be members of this VLAN.)

CAUTION

Do not allow any port memberships or network services on this VLAN that would pose a security risk if exposed to an unauthorized client.

- Statically configure an Authorized-Client VLAN in the switch. The only ports that should belong to this VLAN are ports offering services and access you want available to authenticated clients. 802.1X authenticator ports do not have to be members of this VLAN.Note that if an 802.1X authenticator port is an untagged member of another VLAN, the port's access to that other VLAN will be temporarily removed while an authenticated client is connected to the port.For example, if:

  1. Port 5 is an untagged member of VLAN 1 (the default VLAN).

  2. You configure port 5 as an 802.1X authenticator port.

  3. You configure port 5 to use an Authorized-Client VLAN.

  Then, if a client connects to port 5 and is authenticated, port 5 becomes an untagged member of the Authorized-Client VLAN and is temporarily suspended from membership in the default VLAN.

- If you expect friendly clients to connect without having 802.1X supplicant software running, provide a server on the Unauthorized-Client VLAN for downloading 802.1X supplicant software to the client, and a procedure by which the client initiates the download.

- A client must either have a valid IP address configured before connecting to the switch, or download one through the Unauthorized-Client VLAN from a DHCP server. In the latter case, you will need to provide DHCP services on the Unauthorized-Client VLAN.

- Ensure that the switch is connected to a RADIUS server configured to support authentication requests from clients using ports configured as 802.1X authenticators. (The RADIUS server should not be on the Unauthorized-Client VLAN.)

NOTE

As an alternative, you can configure the switch to use local password authentication instead of RADIUS authentication. However, this is less desirable because it means that all clients use the same passwords and have the same access privileges. Also, you must use 802.1X supplicant software that supports the use of local switch passwords.

CAUTION

Ensure that you do not introduce a security risk by allowing Unauthorized- Client VLAN access to network services or resources that could be compromised by an unauthorized client.

Configuring general 802.1X operation

These steps enable 802.1X authentication, and must be done before configuring 802.1X VLAN operation.

Procedure

1. Enable 802.1X authentication on the individual ports you want to serve as authenticators. (The switch automatically disables LACP on the ports on which you enable 802.1X.) On the ports you will use as authenticators with VLAN operation, ensure that the port-control parameter is set to `auto` (the default). (See **Enable 802.1X authentication on selected ports** on page 526.) This setting requires a client to support 802.1X authentication (with 802.1X supplicant operation) and to provide valid credentials to get network access.

   ```
   aaa port-access authenticator <port-list> control auto
   ```

   Activates 802.1X port-access on ports you have configured as authenticators.

2. Configure the 802.1X authentication type. Options include:

   ```
   aaa authentication port-access <local|eap-radius|chap-radius>
   ```

   Determines the type of RADIUS authentication to use.

3. If you selected either `eap-radius` or `chap-radius` for step 2, use the `radius host` command to configure up to three RADIUS server IP address(es) on the switch.

   ```
   radius host <ip-address>
   ```

   Adds a server to the RADIUS configuration.

   ```
   [key <server-specific key-string>]
   ```

   Optional. Specifies an encryption key for use with the specified server. This key must match the key used on the RADIUS server. Use this option only if the specified server requires a different key than configured for the global encryption key The tilde (~) character is allowed in the string. It is not backward compatible; the "~" character is lost if you use a software version that does not support the "~" character.

   ```
   radius-server key <global key-string>
   ```

   Specifies the global encryption key the switch uses for sessions with servers for which the switch does not have a server-specific key. This key is optional if all RADIUS server addresses configured in the switch include a server- specific encryption key. The tilde (~) character is allowed in the string, for example, `radius-server key hp~switch`. It is not backward compatible; the "~" character is lost if you use a software version that does not support the "~" character.

   Default: Null

   The `no` form of the command removes the global encryption key.

4. Activate authentication on the switch.

   ```
   aaa port-access authenticator active
   ```

   Activates 802.1X port-access on ports you have configured as authenticators.

5. Test both the authorized and unauthorized access to your system to ensure that the 802.1X authentication works properly on the ports you have configured for port-access.

	If you want to implement the optional port-security feature on the switch, you should first ensure that the ports you have configured as 802.1X authenticators operate as expected. Then see **Option for authenticator ports: configure port-security to allow only 802.1X-authenticated devices** on page 548 .
NOTE	

After you complete steps 1 and 2, the configured ports are enabled for 802.1X authentication (without VLAN operation), and you are ready to configure VLAN Operation.

Configuring 802.1X Open VLAN mode

Use these commands to actually configure Open VLAN mode. For a listing of the steps needed to prepare the switch for using Open VLAN mode, see **Setting up and configuring 802.1X Open VLAN mode** on page 545.

Syntax:

```
aaa port-access authenticator <port-list>
```

```
[auth-vid <vlan-id>]
```

Configures an existing, static VLAN to be the Authorized-Client VLAN.

```
[<unauth-vid <vlan-id>]
```

Configures an existing, static VLAN to be the Unauthorized-Client VLAN.

For example, suppose you want to configure 802.1X port-access with Open VLAN mode on ports 10-20 and

- These two static VLANs already exist on the switch:

    ◦ Unauthorized, VID = 80

    ◦ Authorized, VID = 81

- Your RADIUS server has an IP address of 10.28.127.101. The server uses rad4all as a server-specific key string. The server is connected to a port on the Default VLAN.

- The switch's default VLAN is already configured with an IP address of 10.28.127.100 and a network mask of 255.255.255.0

```
HP Switch(config)# aaa authentication port-access eap-radius
```

Configures the switch for 802.1X authentication using an EAP-RADIUS server.

```
HP Switch(config)# aaa port-access authenticator 10-20
```

Configures ports 10 - 20 as 802.1 authenticator ports.

```
HP Switch(config)# radius host 10.28.127.101 key rad4all
```

Configures the switch to look for a RADIUS server with an IP address of 10.28.127.101 and an encryption key of rad4all.

```
HP Switch(config)# aaa port-access authenticator e 10-20 unauth-vid 80
```

Configures ports 10 - 20 to use VLAN 80 as the Unauthorized-Client VLAN.

```
HP Switch(config)# aaa port-access authenticator e 10-20 auth-vid 81
```

Configures ports 10 - 20 to use VLAN 81 as the Authorized-Client VLAN.

```
HP Switch(config)# aaa port-access authenticator active
```

Activates 802.1X port-access on ports you have configured as authenticators.

Inspecting 802.1X Open VLAN mode operation

For information and an example on viewing current Open VLAN mode operation, see **Viewing 802.1X Open VLAN mode status** on page 558.

802.1X Open VLAN operating notes

- Although you can configure Open VLAN mode to use the same VLAN for both the Unauthorized-Client VLAN and the Authorized-Client VLAN, this is not recommended. Using the same VLAN for both purposes allows unauthenticated clients access to a VLAN intended only for authenticated clients, which poses a security breach.

- While an Unauthorized-Client VLAN is in use on a port, the switch temporarily removes the port from any other statically configured VLAN for which that port is configured as a member. Note that the Menu interface will still display the port's statically configured VLAN(s).

- A VLAN used as the Unauthorized-Client VLAN should not allow access to resources that must be protected from unauthenticated clients.

- If a port is configured as a tagged member of VLAN "X", then the port returns to tagged membership in VLAN "X" upon successful client authentication. This happens even if the RADIUS server assigns the port to another, authorized VLAN "Y". Note that if RADIUS assigns VLAN "X" as an authorized VLAN, then the port becomes an untagged member of VLAN "X" for the duration of the client connection. (If there is no Authorized-Client or RADIUS-assigned VLAN, then an authenticated client without tagged VLAN capability can access only a statically configured, untagged VLAN on that port.)

- When a client's authentication attempt on an Unauthorized-Client VLAN fails, the port remains a member of the Unauthorized-Client VLAN until the client disconnects from the port.

- During an authentication session on a port in 802.1X Open VLAN mode, if RADIUS specifies membership in an untagged VLAN, this assignment overrides port membership in the Authorized-Client VLAN. If there is no Authorized-Client VLAN configured, then the RADIUS assignment overrides any untagged VLAN for which the port is statically configured.

- If the only authenticated client on a port loses authentication during a session in 802.1X Open VLAN mode, the port VLAN membership reverts back to the Unauthorized-Client VLAN. If there is no Unauthorized-Client VLAN configured, then the client loses access to the port until it can reauthenticate itself. If there are multiple clients authenticated on the port, if one client loses access and attempts to re-authenticate, that client will be handled as a new client on the port.

- The first client to authenticate on a port configured to support multiple clients will determine the port's VLAN membership for any subsequent clients that authenticate while an active session is already in effect.

Option for authenticator ports: configure port-security to allow only 802.1X-authenticated devices

If 802.1X authentication is disabled on a port or set to `authorized` (Force Authorize), the port can allow access to an unauthenticated client. Port-Security operates with 802.1X authentication only if the selected ports are configured as 802.1X with the `control` mode in the port-access authenticator command set to auto (the default setting). For example, if port 10 was at a nondefault 802.1X setting and you wanted to configure it to support the port-security option, you would use the following `aaa port-access` command:

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Port-access support for port-security operation

```
HP Switch(config)# aaa port-access authenticator 10 control auto
HP Switch(config)# show port-access authenticator 10 config

Port Access Authenticator Configuration

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

     | Re-auth Access  Max  Quiet  TX       Supplicant Server  Cntrl
Port | Period  Control Reqs Period Timeout  Timeout    Timeout Dir
---- + ------- ------- ---- ------ ------- ---------- ------- -----
10   | No      Auto    2    60     30       30         30      both
```

Port-Security

NOTE

If 802.1X port-access is configured on a given port, then port-security `learn-mode` for that port must be set to either `continuous` (the default) or `port-access`.

In addition to the above, to use port-security on an authenticator port, use the per-port `client-limit` option to control how many MAC addresses of 802.1X-authenticated devices the port is allowed to learn. (Using `client-limit` sets 802.1X to user-based operation on the specified ports.) When this limit is reached, no further devices can be authenticated until a currently authenticated device disconnects and the current delay period or logoff period has expired.

Configure the port access type

Syntax:

```
aaa port-access authenticator <port-list> client-limit <1-32>
```

```
aaa port-access authenticator <port-list> client-limit <1-2>
```

Configures user-based 802.1X authentication on the specified ports and sets the number of authenticated devices the port is allowed to learn. For more on this command, see **Configuring switch ports as 802.1X authenticators** on page 526.

–or–

Syntax:

```
no aaa port-access authenticator <port-list> client-limit
```

Configures port-based 802.1X authentication on the specified ports, which opens the port. (See **User authentication methods** on page 519.)

Configuring switch ports to operate as supplicants for 802.1X connections to other switches

A switch port can operate as a supplicant in a connection to a port on another 802.1X-aware switch to provide security on links between 802.1X-aware switches. (A port can operate as both an authenticator and a supplicant.)

Example:

Suppose that you want to connect two switches, where:

- Switch "A" has port 1 configured for 802.1X supplicant operation.

- You want to connect port 1 on switch "A" to port 5 on switch "B".

Figure 130: *Supplicant operation*

Procedure

1. When port 1 on switch "A" is first connected to a port on switch "B", or if the ports are already connected and either switch reboots, port 1 begins sending start packets to port 5 on switch "B".

   - If, after the supplicant port sends the configured number of start packets, it does not receive a response, it assumes that switch "B" is not 802.1X-aware, and transitions to the authenticated state. If switch "B" is operating properly and is not 802.1X-aware, then the link should begin functioning normally, but without 802.1X security. and password.

   - If, after sending one or more start request packets, port 1 receives a request packet from port 5, then switch "B" is operating as an 802.1X authenticator. The supplicant port then sends a response/ID packet. If switch "B" is configured for RADIUS authentication, it forwards this request to a RADIUS server. If switch "B" is configured for Local 802.1X authentication, the authenticator compares the switch "A" response to its local username

2. The RADIUS server then responds with an MD5 access challenge that switch "B" forwards to port 1 on switch "A".

3. Port 1 replies with an MD5 hash response based on its username and password or other unique credentials. Switch "B" forwards this response to the RADIUS server.

4. The RADIUS server then analyzes the response and sends either a "success" or "failure" packet back through switch "B" to port 1.

   - A "success" response unblocks port 5 to normal traffic from port 1.

   - A "failure" response continues the block on port 5 and causes port 1 to wait for the "held-time" period before trying again to achieve authentication through port 5

Supplicant port configuration

Enabling a switch port as a supplicant

You can configure a switch port as a supplicant for a point-to-point link to an 802.1X-aware port on another switch. Configure the port as a supplicant before configuring any supplicant-related parameters.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Syntax:

```
[no] aaa port-access supplicant [ethernet] <port-list>
```

Configures a port as a supplicant with either the default supplicant settings or any previously configured supplicant settings, whichever is most recent. The `no` form of the command disables supplicant operation on the specified ports.

Configuring a supplicant switch port

Enable supplicant operation on a port before changing the supplicant configuration. Execute the supplicant command once without any other parameters, then execute it again with a supplicant parameter that you want to configure. If the intended authenticator port uses RADIUS authentication, then use the `identity` and `secret` options to configure the RADIUS-expected credentials on the supplicant port. If the intended authenticator port uses Local 802.1X authentication, then use the `identity` and `secret` options to configure the authenticator switch local username and password on the supplicant port.

Syntax:

```
aaa port-access supplicant [ethernet] <port-list>
```

To enable supplicant operation on the designated ports, execute this command without any other parameters. To configure supplicant operation, use the same command again with the following parameters. Use one instance of the command for each parameter you want to configure. The `no` form disables supplicant operation on the designated ports.

```
[identity <username>]
```

Sets the username and password to send in response to an authentication request. If the intended authenticator port is configured for RADIUS authentication, then <*username*> and <*password*> must be the username and password expected by the RADIUS server. If the intended authenticator port is configured for Local authentication, then <*username*> and <*password*> must be the username and password configured on the Authenticator switch. (Default: Null.)

```
[secret]
```

Enter secret: <*password*>

Repeat secret: <*password*>

Sets the secret password to be used by the port supplicant when an MD5 authentication request is received from an authenticator. The switch prompts you to enter the secret password after the command is invoked.

```
[auth-timeout <1-300>]
```

Sets the delay period that the port waits to receive a challenge from the authenticator. If the request times out, the port sends another request, up to the number of attempts specified by the `max-start` parameter. (Default: 30 seconds).

```
[max-start <1-10>]
```

Defines the maximum number of times the supplicant port requests authentication. See step 1 in **Configuring switch ports to operate as supplicants for 802.1X connections to other switches** on page 549 for a description of how the port reacts to the authenticator response. (Default: 3).

```
[held-period <0-65535>]
```

Sets the time period that the supplicant port waits after an active 802.1X session fails before trying to reacquire the authenticator port. (Default: 60 seconds)

```
[start-period <1-300>]
```

Sets the delay between Start packet retransmissions. That is, after a supplicant sends a start packet, it waits during the `start-period` for a response. If no response comes during the `start-period`, the supplicant sends a new start packet. The `max-start` setting specifies how many start attempts are allowed in the session. (Default: 30 seconds)

```
[initialize]
```

On the specified ports, blocks inbound and outbound traffic and restarts the 802.1X authentication process. Affects only ports configured as 802.1X supplicants.

```
[clear-statistics]
```

Clears and restarts the 802.1X supplicant statistics counters.

Displaying 802.1X configuration, statistics, and counters

Show commands for port-access authenticator

show port-access authenticator

The port-access authenticator command

Syntax:

```
show port-access authenticator [port-list] [config|statistics|session-counters|vlan|clients]|[detailed]
```

If you enter the `show port-access authenticator` command without an optional value, the following configuration information is displayed for all switch ports, or specified ports, that are enabled for 802.1X port-access authentication:

- Port-access authenticator activated: Are any switch ports configured to operate as 802.1X authenticators using the `aaa port-access authenticator` command? `Yes` or `No`

- Allow RADIUS-assigned dynamic (GVRP) VLANs: Are RADIUS-assigned dynamic (GVRP-learned) VLANs supported for authenticated and unauthenticated client sessions on the switch? `Yes` or `No`

- Auth Clients: Number of authorized clients

- Unauth Clients: Number of unauthorized clients

- Untagged VLAN: VLAN ID number of the untagged VLAN used in client sessions. If the switch supports MAC-based (untagged) VLANs, MAC-based is displayed to show that multiple untagged VLANs are configured for authentication sessions.

- Tagged VLANs: Are tagged VLANs (statically configured or RADIUS-assigned) used for authenticated clients? `Yes` or `No`

- Port CoS:

- Yes – Client-specific CoS (Class of Service) values are applied to more than one authenticated client on the port.

- No – No client-specific CoS values are applied to any authenticated client on the port.

- <CoS value> – Numerical value of the CoS (802.1p priority) applied to inbound traffic from one authenticated client. For client-specific per-port CoS values, enter the `show port-access web-based clients detailed` command.

• % In Limit: Inbound rate limit applied.

• RADIUS ACL: Are RADIUS-assigned ACLs used for authenticated clients? Yes or No

• Cntrl Dir: Direction in which flow of incoming and outgoing traffic is blocked on 802.1X-aware port that has not yet entered the authenticated state:

- Both: Incoming and outgoing traffic is blocked on port until authentication occurs.

- In: Only incoming traffic is blocked on port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on the unauthenticated 802.1X-aware port.

The show port-access authenticator command

```
HP Switch(config)# show port-access authenticator

Port Access Authenticator Status

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : Yes

     Auth    Unauth   Untagged Tagged              % In    RADIUS Cntrl
Port Clients Clients  VLAN     VLANs  Port COS     Limit   ACL    Dir
---- ------- -------- -------- ------ --------     ------  ------ -----
1    1       1        4006     Yes    77777777     No      Yes    both
2    2       0        MACbased No     No           No      Yes    both
3    4       0        1        Yes    No           No      No     both
```

show port-access authenticator config

The `show port-access authenticator config` command

Syntax:

```
show port-access authenticator config [port-list]
```

Displays 802.1X port-access authenticator configuration settings, including:

Whether port-access authentication is enabled.

Whether RADIUS-assigned dynamic VLANs are supported.

802.1X configuration of ports that are enabled as 802.1X authenticators. Use the `show running` command to view the current `client-limit` configuration available for switches.

You can display 802.1X port-access authenticator configuration for all switch ports or specified ports. 802.1X configuration information for ports that are not enabled as 802.1X authenticators is not displayed.

```
Switch(config)# show port-access authenticator config

Port Access Authenticator Configuration

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

     | Re-auth Access    Max    Quiet   TX         Supplicant Server   Cntrl
Port | Period Control    Reqs   Period  Timeout    Timeout    Timeout  Dir
---- + ------- --------- -----  ------- --------   ---------- -------- -----
1    | No      Auto      2      60      30         30         30       both
2    | No      Auto      2      60      30         30         30       in
```

Table 43: *Field table*

Field	Description
Port-access authenticator activated	802.1X authentication enabled/disabled status for specified ports.
Port	Port number on switch.
Re-auth Period	Time period (in seconds) after which clients connected to the port must be reauthenticated.
Access Control	Authentication mode of port: **Auto:** Network access is allowed to any connected device that supports 802.1X authentication and provides valid 802.1X credentials. **Authorized:** Network access is allowed to any device connected to the port, regardless of whether it meets 802.1X criteria. **Unauthorized:** Network access is blocked to any device connected to the port, regardless of whether the device meets 802.1X criteria.
Max reqs	Number of authentication attempts that must time out before authentication fails and the authentication session ends.
Quiet Period	Time period (in seconds) during which the port does not try to acquire a supplicant.
TX Timeout	Time period (in seconds) that the port waits to retransmit the next EAPOL PDU during an authentication session.

Table Continued

Field	Description
Supplicant Timeout	Time period (in seconds) that the switch waits for a supplicant response to an EAP request.
Server Timeout	Time period (in seconds) that the switch waits for a server response to an authentication request.
Cntrl Dir	Direction in which flow of incoming and outgoing traffic is blocked on 802.1X-aware port that has not yet entered the authenticated state: **Both:** Incoming and outgoing traffic is blocked on port until authentication occurs. **In:** Only incoming traffic is blocked on port before authentication occurs. Outgoing traffic with unknown destination addresses is flooded on the unauthenticated 802.1X-aware port.

show port-access authenticator statistics

The show port-access authenticator statistics command

Syntax:

show port-access authenticator statistics [*port-list*]

Displays statistical information for all switch ports or specified ports that are enabled as 802.1X authenticators, including:

- Whether port-access authentication is enabled

- Whether RADIUS-assigned dynamic VLANs are supported

- 802.1X supplicant's MAC address as determined by the content of the last EAPOL frame received on the port

- 802.1X traffic statistics from received and transmitted packets

802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.

```
HP Switch(config)# show port-access authenticator statistics

Port Access Authenticator Statistics

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

     Source        TX    TX   RX    RX     RX     RX   RX
Port MAC address   ReqId Req  Star  Logoff RespId Resp Errors
---- ------------- ----- ---- ----- ------ ------ ---- ------
2    001560-b3ea48 1     0    0     0      0      0    0
```

show port-access authenticator session-counters

The show port-access authenticator session-counters command

Syntax:

show port-access authenticator session-counters [*port-list*]

Displays information for active 802.1X authentication sessions on all switch ports or specified ports that are enabled as 802.1X authenticators, including:

- 802.1X frames received and transmitted on each port

- Duration and status of active 802.1X authentication sessions (`in-progress` or `terminated`)

- User name of 802.1X supplicant included in 802.1X response packets, configured with the `aaa port-access supplicant identity <username>` command

802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.

```
Port Access Authenticator Session Counters

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

                             Session    Session
Port Frames In  Frames Out  Time(sec.) Status      User
---- ---------- ----------- ---------- ----------- -----
2    45623      45623       20         in-progress bert
```

show port-access authenticator vlan

Syntax:

```
show port-access authenticator vlan [port-list]
```

Displays the following information on the VLANs configured for use in 802.1X port-access authentication on all switch ports, or specified ports, that are enabled as 802.1X authenticator:

- Authentication mode used on each port, configured with the `aaa port-access authenticator control` command

- VLAN ID (if any) to be used for traffic from 802.1X-authenticated clients

- VLAN ID (if any) to be used for traffic from unauthenticated clients

802.1X configuration information for ports that are not enabled as an 802.1X authenticators is not displayed.

```
HP Switch(config)# show port-access authenticator vlan

Port Access Authenticator VLAN Configuration

Port-access authenticator activated [No] : Yes
Allow RADIUS-assigned dynamic (GVRP) VLANs [No] : No

     Access    Unauth    Auth
Port Control   VLAN ID   VLAN ID
---- -------- -------- --------
2    Auto     0         0
```

Syntax:

```
show port-access authenticator clients [port-list]
```

Displays the session status, name, and address for each 802.1X port-access-authenticated client on the switch. Multiple authenticated clients may be displayed for the same port. The IP address displayed is taken from the DHCP binding table (learned through the DHCP Snooping feature).

- If DHCP snooping is not enabled on the switch, n/a (not available) is displayed for a client IP address.

- If an 802.1X-authenticated client uses an IPv6 address, n/a - IPv6 is displayed.

- If DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table, `n/a - no info` is displayed.

show port-access authenticator clients

The show port-access authenticator clients command

Syntax:

```
show port-access authenticator clients [port-list]
```

Displays the session status, name, and address for each 802.1X port-access-authenticated client on the switch. Multiple authenticated clients may be displayed for the same port. The IP address displayed is taken from the DHCP binding table (learned through the DHCP Snooping feature).

- If DHCP snooping is not enabled on the switch, n/a (not available) is displayed for a client IP address.

- If an 802.1X-authenticated client uses an IPv6 address, n/a - IPv6 is displayed.

- If DHCP snooping is enabled but no MAC-to-IP address binding for a client is found in the DHCP binding table, n/a - no info is displayed.

The show port-access authenticator clients command output

```
Switch (config)# show port-access authenticator clients

Port Access Authenticator Client Status

Port   Client Name MAC Address   IP Address       Session Status
-----  ----------- ------------- ---------------  --------------
1      webuser1     001321-eb8063 192.192.192.192 Authenticated
1      webuser2     001560-b3ea48 n/a - no info   Authenticating
1      webuser3     000000-111111 n/a - IPv6      Authenticating
3      webuser4     000000-111112 n/a             Authenticating
```

Syntax:

```
show port-access authenticator clients <port-list> detailed
```

Displays detailed information on the status of 802.1X-authenticated client sessions on specified ports, including the matches the switch detects for individual ACEs configured with the cnt (counter) option in an ACL assigned to the port by a RADIUS server.

Output for the show port-access authenticator clients detailed command

```
Switch (config)# show port-access authenticator clients 5 detailed

Port Access Authenticator Client Status Detailed

Client Base Details :
Port           : 5
Session Status : Open          Session Time(sec) : 999999999
Frames In      : 999999999     Frames Out        : 99999999
```

```
Username        : webuser1        MAC Address      : 001321-eb8063
IP              : 2001:fecd:ba23:cd1f:dcb1:1010:9234:4088

Access Policy Details :
COS Map         : 70000000             In Limit %  : 87
Untagged VLAN  : 3096                 Out Limit % : 100
Tagged VLANs    : 1, 3, 5, 6, 334, 2066
RADIUS-ACL List :
  deny in udp from any to 10.2.8.233 CNT
     Hit Count: 10
  permit in udp from any to 10.2.8.233 CNT
     Hit Count: 17
  deny in tcp from any to 10.2.8.233 CNT
     Hit Count: 1
  permit in tcp from any to 10.2.8.233 CNT
     Hit Count: 11
  permit in ip from any to any cnt
     Hit Count: 42
```

Viewing 802.1X Open VLAN mode status

You can examine the switch's current VLAN status by using the `show port-access authenticator vlan` and `show port-access authenticator <port-list>` commands as shown in this section. The table describes the data that these two commands display. **Figure 132: Example showing a VLAN with ports**

configured for Open VLAN mode on page 561 shows related VLAN data that can help you to see how the switch is using statically configured VLANs to support 802.1X operation.

Figure 131: *Example showing ports configured for Open VLAN mode*

```
HP Switch# show port-access authenticator vlan
 Port Access Authenticator VLAN Configuration

  Port-access authenticator activated [No] : Yes


         Access   Unauth   Auth
  Port Control   VLAN ID  VLAN ID
  ---- --------  -------- --------
   1     Auto     100       101
   2     Auto     100       101
   3     Auto     100       0
   4     Auto     100       101

HP Switch# show port-access authenticator 1-4
 Port Access Authenticator Status

  Port-access authenticator activated [No] : No


              Authenticator  Authenticator  Current   Current     % Curr. Rate
  Port Status State          Backend State  VLAN ID   Port COS    Limit Inbound
  ---- ------ -------------- -------------- --------  ----------  ---------------

   1    Closed Connecting     Idle           100       No-override No-override
   2    Open   Authorized     Idle           101       No-override No-override
   3    Closed Connecting     Idle           100       No-override No-override
   4    Closed Connecting     Idle           No PVID   No-override No-override
```

In these two **show** outputs, an Unauth VLAN ID appearing in the Current VLAN ID column for the same port indicates an unauthenticated client is connected to this port. (Assumes that the port is not a statically configured member of VLAN 100.)

Note: Series 5400zl switches do not include the **Authenticator State** and **Authenticator Backend State** fields shown in this figure.

Items 1 through 3 indicate that an authenticated client is connected to port 2:
1. **Open** in the Status column
2. **Authorized** in the Authenticator State column
3. The Auth VLAN ID (**101**) is also in the Current VLAN ID column. (This assumes that the port is not a statically configured member of VLAN 101.)
4. A "0" in the row for port 3 indicates there is no Authorized VLAN configured for port 3.
5. No PVID" means there is currently no untagged VLAN membership on port 4.

Notes on the above image:

- When the **Auth VLAN ID** is configured and matches the **Current VLAN ID**, an authenticated client is connected to the port. (This assumes the port is not a statically configured member of the VLAN you are using for Auth VLAN.)

- When the **Unauth VLAN ID** is configured and matches the **Current VLAN ID**, an unauthenticated client is connected to the port. (This assumes the port is not a statically configured member of the VLAN you are using for Unauth VLAN.)

NOTE

Because a temporary Open VLAN port assignment to either an authorized or unauthorized VLAN is an untagged VLAN membership, these assignments temporarily replace any other untagged VLAN membership that is statically configured on the port. For example, if port 12 is statically configured as an untagged member of VLAN 1, but is configured to use VLAN 25 as an authorized VLAN, then the port's membership in VLAN 1 will be temporarily suspended whenever an authenticated 802.1X client is attached to the port.

Table 44: *Output for determining Open VLAN mode status*

Status Indicator	Meaning

Access Control	This state is controlled by the following port-access Command syntax:
	
	 `HP Switch(config)# aaa port-access authenticator <port-list> control`
	 `<authorized|auto|unauthorized>`
	
	 Auto: Configures the port to allow network access to any connected device that supports 802.1X authentication and provides valid 802.1X credentials. (This is the default authenticator setting.)**Authorized:** Configures the port for "Force Authorized", which allows access to any device connected to the port, regardless of whether it meets 802.1X criteria. (You can still configure console, Telnet, or SSH security on the port.)**Unauthorized:** Configures the port for "Force Unauthorized", which blocks access to any device connected to the port, regardless of whether the device meets 802.1X criteria.

Unauthorized VLAN ID	**<vlan-id>:** Lists the VID of the static VLAN configured as the unauthorized VLAN for the indicated port.**0:** No unauthorized VLAN has been configured for the indicated port.
Authorized VLAN ID	**<vlan-id>:** Lists the VID of the static VLAN configured as the authorized VLAN for the indicated port.**0:** No authorized VLAN has been configured for the indicated port.
Status	**Closed:** Either no client is connected or the connected client has not received authorization through 802.1X authentication.**Open:** An authorized 802.1X supplicant is connected to the port.
Current VLAN ID	**<vlan-id>:** Lists the VID of the static, untagged VLAN to which the port currently belongs.**No PVID:** The port is not an untagged member of any VLAN.
Current Port CoS	See **RADIUS Authentication, Authorization, and Accounting** on page 194.
% Curr. Rate Limit Inbound	

Syntax:

```
show vlan <vlan-id>
```

Displays the port status for the selected VLAN, including an indication of which port memberships have been temporarily overridden by Open VLAN mode.

Figure 132: *Example showing a VLAN with ports configured for Open VLAN mode*

```
                    HP Switch(config)# show vlan 1

                     Status and Counters - VLAN Information - VLAN 1

                      VLAN ID : 1
                      Name : DEFAULT_VLAN
                      Status : Static
                      Voice : No
                      Jumbo : No

                      Port Information Mode      Unknown VLAN Status
                      ---------------- -------- ------------ ----------
                      A1               Untagged Learn        Up
                      A2               Untagged Learn        Up
                      A3               Untagged Learn        Up
                      A4               Untagged Learn        Up
                      B2               Untagged Learn        Up
                      B4               Untagged Learn        Up
                      .
                      .
                      .
                      B23              Untagged Learn        Up
                      B24              Untagged Learn        Up

                      Overriden Port VLAN configuration

                      Port Mode
                      ---- ----------
                      B1   Untagged
                      B3   Untagged
```

Note that ports B1 and B3 are not in the upper listing, but are included under "Overridden Port VLAN configuration". This shows that static, untagged VLAN memberships on ports B1 and B3 have been overridden by temporary assignment to the authorized or unauthorized VLAN. Using the **show port-access authenticator <** *port-list* **>** command shown in figure 13-18 provides details.

Show commands for port-access supplicant

Syntax:

```
show port-access supplicant [<port-list>] [statistics]
```

```
show port-access supplicant [<port-list>]
```

Shows the port-access supplicant configuration (excluding the `secret` parameter) for all ports or `<port-list >` ports configured on the switch as supplicants. The Supplicant State can include the following:

Connecting

Starting authentication.

Authenticated

Authentication completed (regardless of whether the attempt was successful).

Acquired

The port received a request for identification from an authenticator.

Authenticating

Authentication is in progress.

Held

Authenticator sent notice of failure. The supplicant port is waiting for the authenticator's held-period.

For descriptions of the supplicant parameters, see **Configuring a supplicant switch port**.


```
show port-access supplicant [<port-list>] statistics
```

Shows the port-access statistics and source MAC address(es) for all ports or `<port-list >` ports configured on the switch as supplicants. See **Note on supplicant statistics**.

Note on supplicant statistics

For each port configured as a supplicant, `show port-access supplicant statistics <port-list>]` displays the source MAC address and statistics for transactions with the authenticator device most recently detected on the port. If the link between the supplicant port and the authenticator device fails, the supplicant port continues to show data received from the connection to the most recent authenticator device until one of the following occurs:

- The supplicant port detects a different authenticator device.

- You use the `aaa port-access supplicant <port-list> clear-statistics` command to clear the statistics for the supplicant port.

- The switch reboots.

Thus, if the supplicant's link to the authenticator fails, the supplicant retains the transaction statistics it most recently received until one of the above events occurs. Also, if you move a link with an authenticator from one supplicant port to another without clearing the statistics data from the first port, the authenticator's MAC address will appear in the supplicant statistics for both ports.

How RADIUS/802.1X authentication affects VLAN operation

Static VLAN requirement:

RADIUS authentication for an 802.1X client on a given port can include a (static) VLAN requirement. (Refer to the documentation provided with your RADIUS application.) The static VLAN to which a RADIUS server assigns a client must exist on the switch. If it does not exist or is a dynamic VLAN (created by GVRP), authentication fails. Also, for the session to proceed, the port must be an untagged member of the required VLAN. If it is not, the switch temporarily reassigns the port.

If the port used by the client is not configured as an untagged member of the required static VLAN:

When a client is authenticated on port "N", if port "N" is not already configured as an untagged member of the static VLAN specified by the RADIUS server, then the switch temporarily assigns port "N" as an untagged member of the required VLAN (for the duration of the 802.1X session). At the same time, if port "N" is already configured as an untagged member of another VLAN, port "N" loses access to that other VLAN for the duration of the session. (This loss of access occurs because a port can be an untagged member of only one VLAN at a time.)

Using a RADIUS server to authenticate clients, you can provide port-level security protection from unauthorized network access for the following authentication methods:

- 802.1X: Port-based or client-based access control to open a port for client access after authenticating valid user credentials.

- MAC address: Authenticates a device MAC address to grant access to the network

- WebAgent: Authenticates clients for network access using a webpage for user login.

> **NOTE**
> You can use 802.1X (port-based or client-based) authentication and either Web or MAC authentication at the same time on a port, with a maximum of 32 clients allowed on the port. (The default is one client.) Web authentication and MAC authentication are mutually exclusive on the same port. Also, you must disable LACP on ports configured for any of these authentication methods.

VLAN assignment on a port

Following client authentication, VLAN configurations on a port are managed as follows when you use 802.1X, MAC, or Web authentication:

- The port resumes membership in any tagged VLANs for which it is already assigned in the switch configuration. Tagged VLAN membership allows a port to be a member of multiple VLANs simultaneously.

- The port is temporarily assigned as a member of an untagged (static or dynamic) VLAN for use during the client session according to the following order of options.

  1. The port joins the VLAN to which it has been assigned by a RADIUS server during client authentication.

  2. If RADIUS authentication does not include assigning the port to a VLAN, then the switch assigns the port to the authorized-client VLAN configured for the authentication method.

  3. If the port does not have an authorized-client VLAN configured, but is configured for membership in an untagged VLAN, the switch assigns the port to this untagged VLAN.

Operating notes

During client authentication, a port assigned to a VLAN by a RADIUS server or an authorized-client VLAN configuration is an untagged member of the VLAN for the duration of the authenticated session. This applies even if the port is also configured in the switch as a tagged member of the same VLAN. The following restrictions apply:

If the port is assigned as a member of an untagged static VLAN, the VLAN must already be configured on the switch. If the static VLAN configuration does not exist, the authentication fails. If the port is assigned as a member of an untagged dynamic VLAN that was learned through GVRP, the dynamic VLAN configuration must exist on the switch at the time of authentication and GVRP-learned dynamic VLANs for port-access authentication must be enabled. If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

- To enable the use of a GVRP-learned (dynamic) VLAN as the untagged VLAN used in an authentication session, enter the `aaa port-access gvrp-vlans` command.

- Enabling the use of dynamic VLANs in an authentication session offers the following benefits:

  ◦ You avoid the need of having static VLANs pre-configured on the switch.

  ◦ You can centralize the administration of user accounts (including user VLAN IDs) on a RADIUS server.

For information on how to enable the switch to dynamically create 802.1Q-compliant VLANs on links to other devices using the GARP VLAN Registration Protocol (GVRP), see "GVRP" in the advanced traffic management guide.

- For an authentication session to proceed, a port must be an untagged member of the (static or dynamic) VLAN assigned by the RADIUS server (or an authorized-client VLAN configuration). The port temporarily drops any current untagged VLAN membership.If the port is not already a member of the RADIUS-assigned (static or dynamic) untagged VLAN, the switch temporarily reassigns the port as an untagged member of the required VLAN (for the duration of the session). At the same time, if the port is already configured as an untagged member of a different VLAN, the port loses access to the other VLAN for the duration of the session. (A port can be an untagged member of only one VLAN at a time.)When the authentication session ends, the switch removes the temporary untagged VLAN assignment and re-activates the temporarily disabled, untagged VLAN assignment.

- If GVRP is already enabled on the switch, the temporary untagged (static or dynamic) VLAN created on the port for the authentication session is advertised as an existing VLAN. If this temporary VLAN assignment causes the switch to disable a different untagged static or dynamic VLAN configured on the port (as described in the preceding bullet) the disabled VLAN assignment is not advertised. When the authentication session ends, the switch:

  ◦ Removes the temporary untagged VLAN assignment and stops advertising it.

  ◦ Re-activates and resumes advertising the temporarily disabled, untagged VLAN assignment.

- If you modify a VLAN ID configuration on a port during an 802.1X, MAC, or Web authentication session, the changes do not take effect until the session ends.

- When a switch port is configured with RADIUS-based authentication to accept multiple 802.1X and/or MAC or Web authentication client sessions, all authenticated clients must use the same port-based, untagged VLAN membership assigned for the earliest, currently active client session.Therefore, on a port where one or more authenticated client sessions are already running, all such clients are on the same untagged VLAN. If a RADIUS server subsequently authenticates a new client, but attempts to re-assign the port to a different, untagged VLAN than the one already in use for the previously existing, authenticated client sessions, the connection for the new client will fail.

Example of untagged VLAN assignment in a RADIUS-based authentication session

The following example shows how an untagged static VLAN is temporarily assigned to a port for use during an 802.1X authentication session. In the example, an 802.1X-aware client on port A2 has been authenticated by a

RADIUS server for access to VLAN 22. However, port A2 is not configured as a member of VLAN 22 but as a member of untagged VLAN 33 as shown here:

Figure 133: *Active VLAN configuration*

```
===========================- CONSOLE - MANAGER MODE -===========================
               Switch Configuration - VLAN - VLAN Port Assignment

   Port    default_vlan      vlan_22          vlan_33          vlan_44
   ---- + ------------   ------------    ------------    ------------
   A1   | Untagged        Tagged          No               No
   A2   | No              No              Untagged         No
   A3   | Untagged        Forbid          Forbid           Forbid
   A4   | Untagged        Tagged          Tagged           Tagged
    :           :              :                :                :
    :           :              :                :                :

   Actions->    Cancel      Edit       Save       Help

   Cancel changes and return to previous screen.
   Use arrow keys to change action selection and <Enter> to execut
```

Scenario: An authorized 802.1X client requires access to VLAN 22 from port A2. However, access to VLAN 22 is blocked (not untagged or tagged) on port A2 and

In this example, if RADIUS authorizes an 802.1X client on port A2 with the requirement that the client use VLAN 22, then:

• VLAN 22 becomes available as Untagged on port A2 for the duration of the session.

• VLAN 33 becomes unavailable to port A2 for the duration of the session (because there can be only one untagged VLAN on any port).

To view the temporary VLAN assignment as a change in the active configuration, use the `show vlan <vlan-id>` command where `<vlan-id>` is the (static or dynamic) VLAN used in the authenticated client session.

Figure 134: *The active configuration for VLAN 22 temporarily changes for the 802.1X session*

```
HP Switch(config)# show vlan 22

 Status and Counters - VLAN Information - VLAN 22

  VLAN ID : 22
  Name : vlan 22
  Status : Static
  Voice : No
  Jumbo : No

  Port Information Mode        Unknown VLAN Status
  ---------------- --------    ------------ ----------
  A1               Tagged      Learn        Up
  A2               802.1X      Learn        Up
  A4               Tagged      Learn        Up
  .
  .
  .

  Overriden Port VLAN configuration

  Port Mode
  ---- ----------
  A2   No
```

This entry shows that port A2 is temporarily untagged on VLAN 22 for an 802.1X session. This is to accommodate an 802.1X client's access, authenticated by a RADIUS server, where the server included an instruction to put the client's access on VLAN 22.

Note: With the current VLAN configuration (figure 13-20), the only time port A2 appears in this **show vlan 22** listing is during an 802.1X session with an attached client. Otherwise, port A2 is not listed.

However, as shown in **Figure 133: Active VLAN configuration** on page 565, because VLAN 33 is configured as untagged on port A2 and because a port can be untagged on only one VLAN, port A2 loses access to VLAN 33 for the duration of the 802.1X session on VLAN 22.

You can verify the temporary loss of access to VLAN 33 by entering the `show vlan 33` command.

Figure 135: *The active configuration for VLAN 33 temporarily drops port 22 for the 802.1X session*

```
                    HP Switch(config)# show vlan 33

                     Status and Counters - VLAN Information - VLAN 33

                      VLAN ID : 33
                      Name : VLAN_33
                      Status : Static
                      Voice : No
                      Jumbo : No

                      Port Information Mode      Unknown VLAN Status
                      ---------------- --------  ------------ ---------
                      A4                 Tagged   Learn         Up

                      Overriden Port VLAN configuration

                      Port Mode
                      ---- ----------
                      A2   Untagged
```

Even though port A2 is configured as Untagged on (static) VLAN 33 (see figure 13-20), it does not appear in the VLAN 33 listing while the 802.1X session is using VLAN 22 in the Untagged status. However, after the 802.1X session with VLAN 22 ends, the active configuration returns port A2 to VLAN 33.

When the 802.1X client's session on port A2 ends, the port removes the temporary untagged VLAN membership. The static VLAN (VLAN 33) that is "permanently" configured as untagged on the port becomes available again. Therefore, when the RADIUS-authenticated 802.1X session on port A2 ends, VLAN 22 access on port A2 also ends, and the untagged VLAN 33 access on port A2 is restored.

Figure 136: *The active configuration for VLAN 33 restores port A2 after the 802.1X session ends*

```
                    HP Switch(config)# show vlan 33

                     Status and Counters - VLAN Information - VLAN 33

                      VLAN ID : 33
                      Name : VLAN_33
                      Status : Static
                      Voice : No
                      Jumbo : No
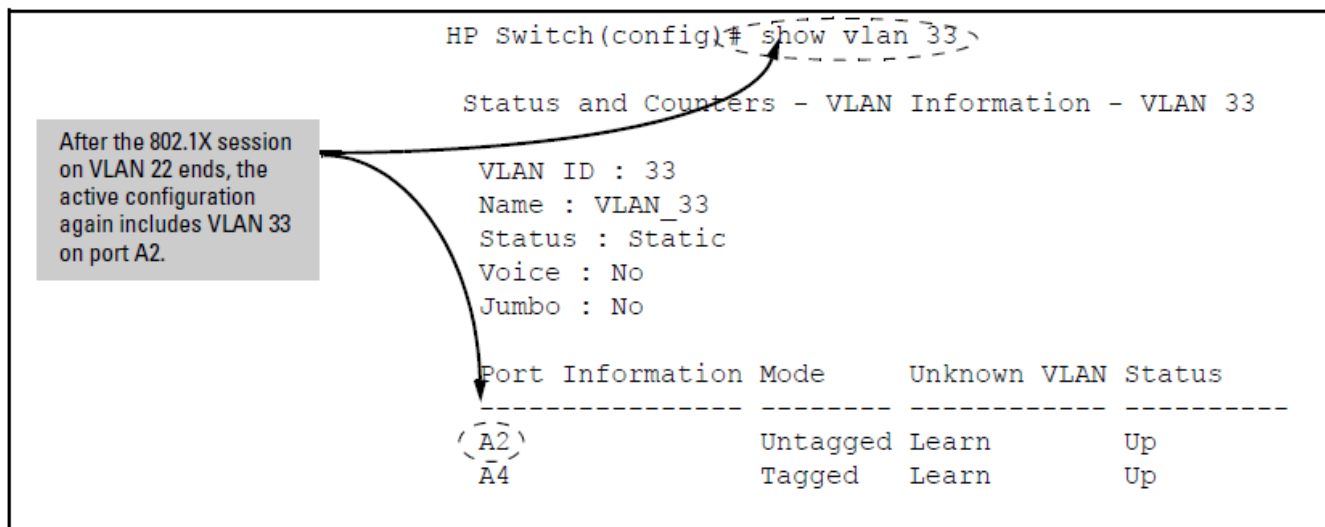
                      Port Information Mode      Unknown VLAN Status
                      ---------------- --------  ------------ ---------
                      A2                 Untagged Learn         Up
                      A4                 Tagged   Learn         Up
```

After the 802.1X session on VLAN 22 ends, the active configuration again includes VLAN 33 on port A2.

Enabling the use of GVRP-learned dynamic VLANs in authentication sessions

Syntax:

`aaa port-access gvrp-vlans`

Enables the use of dynamic VLANs (learned through GVRP) in the temporary untagged VLAN assigned by a RADIUS server on an authenticated port in an 802.1X, MAC, or Web authentication session.

Enter the `no` form of this command to disable the use of GVRP-learned VLANs in an authentication session.

For information on how to enable a switch to dynamically create 802.1Q-compliant VLANs, see "GVRP" in the advanced traffic management guide.

NOTE

1. If a port is assigned as a member of an untagged dynamic VLAN, the dynamic VLAN configuration must exist at the time of authentication and GVRP for port-access authentication must be enabled on the switch.If the dynamic VLAN does not exist or if you have not enabled the use of a dynamic VLAN for authentication sessions on the switch, the authentication fails.

2. After you enable dynamic VLAN assignment in an authentication session, it is recommended that you use the `interface unknown-vlans` command on a per-port basis to prevent denial-of-service attacks. The `interface unknown-vlans` command allows you to:

   • Disable the port from sending advertisements of existing GVRP-created VLANs on the switch.

   • Drop all GVRP advertisements received on the port. See "GVRP" in the advanced traffic management guide.

3. If you disable the use of dynamic VLANs in an authentication session using the `no aaa port-access gvrp-vlans` command, client sessions that were authenticated with a dynamic VLAN continue and are not deauthenticated.

   (This behavior differs from how static VLAN assignment is handled in an authentication session. If you remove the configuration of the static VLAN used to create a temporary client session, the 802.1X, MAC, or Web authenticated client is deauthenticated.)

   However, if a RADIUS-configured dynamic VLAN used for an authentication session is deleted from the switch through normal GVRP operation (for example, if no GVRP advertisements for the VLAN are received on any switch port), authenticated clients using this VLAN are deauthenticated.

NOTE

Any port VLAN-ID changes you make on 802.1X-aware ports during an 802.1X-authenticated session do not take effect until the session ends.

With GVRP enabled, a temporary, untagged static VLAN assignment created on a port by 802.1X authentication is advertised as an existing VLAN. If this temporary VLAN assignment causes the switch to disable a configured (untagged) static VLAN assignment on the port, then the disabled VLAN assignment is not advertised. When the 802.1X session ends, the switch:

• Eliminates and ceases to advertise the temporary VLAN assignment.

• Re-activates and resumes advertising the temporarily disabled VLAN assignment

---

     Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Messages related to 802.1X operation

Table 45: *802.1X operating messages*

Message	Meaning
`Port <port-list> is not an authenticator.`	The ports in the port list have not been enabled as 802.1X authenticators. Use this command to enable the ports as authenticators: `HP Switch(config)# aaa port-access authenticator e 10`
`Port <port-list> is not a supplicant.`	Occurs when there is an attempt to change the supplicant configuration on a port that is not currently enabled as a supplicant. Enable the port as a supplicant and then make the desired supplicant configuration changes. See **Enabling a switch port as a supplicant** on page 550.
`No server(s) responding.`	This message can appear if you configured the switch for EAP-RADIUS or CHAP-RADIUS authentication, but the switch does not receive a response from a RADIUS server. Ensure that the switch is configured to access at least one RADIUS server. (Use `show radius`.) If you also see the message `Can't reach RADIUS server <x.x.x.x>`, try the suggestions listed for that message (see **Messages related to RADIUS operation** on page 242).

`LACP has been disabled on 802.1X port(s).`
 	
`Error configuring port <port-list> LACP and 802.1X cannot be run together.`
	To maintain security, LACP is not allowed on ports configured for 802.1X authenticator operation. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables 802.1X on that port.Also, the switch will not allow you to configure LACP on a port on which port access (802.1X) is enabled.

Overview

• **Port security** on page 570: This feature enables you to configure each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch.

• **MAC Lockdown** on page 584: This feature, also known as "Static Addressing", is used to prevent station movement and MAC address "hijacking" by allowing a given MAC address to use only an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN.

• **MAC Lockout** on page 589: This feature enables you to block a specific MAC address so that the switch drops all traffic to or from the specified address.

> **NOTE**
>
> Port Security and MAC Lockdown do not prevent intruders from receiving broadcast and multicast traffic. Also, Port Security and MAC Lockdown are mutually exclusive on a switch. If one is enabled, then the other cannot be used.

Port security

Basic operation

Default port security operation: The default port security setting for each port is off, or "continuous". That is, any device can access a port without causing a security reaction.

Intruder protection: A port that detects an "intruder" blocks the intruding device from transmitting to the network through that port.

Eavesdrop protection: Using either the port-security command or the switch WebAgent to enable port security on a given port automatically enables eavesdrop prevention on that port.

General operation for port security: On a per-port basis, you can configure security measures to block unauthorized devices, and to send notice of security violations. Once port security is configured, you can then monitor the network for security violations through one or more of the following:

• Alert flags captured by network management tools.

• Alert Log entries in the WebAgent.

• Event Log entries in the console interface.

• Intrusion Log entries in the menu interface, CLI, or WebAgent.

For any port, you can configure the following:

• Action: Used when a port detects an intruder. Specifies whether to send an SNMP trap to a network management station and whether to disable the port.

• Address Limit: Sets the number of authorized MAC addresses allowed on the port.

• Learn-Mode: Specify how the port acquires authorized addresses.

- Limited-Continuous: Sets a finite limit (1-32) to the number of learned addresses allowed per port.

- Continuous: Allows the port to learn addresses from inbound traffic from any connected device. This is the default setting.

- Static: Enables you to set a fixed limit on the number of MAC addresses authorized for the port and to specify some or all the authorized addresses. (If you specify only some of the authorized addresses, the port learns the remaining authorized addresses from the traffic it receives from connected devices.)

- Configured: Requires that you specify all MAC addresses authorized for the port. The port is not allowed to learn addresses from inbound traffic.

- Authorized (MAC) Addresses: Specify up to eight devices (MAC addresses) that are allowed to send inbound traffic through the port. This feature:

  - Closes the port to inbound traffic from any unauthorized devices that are connected to the port.

  - Provides the option for sending an SNMP trap notifying of an attempted security violation to a network management station and, optionally, disables the port. (For more on configuring the switch for SNMP management, see "Trap receivers and authentication traps" in the management and configuration guide for your switch.)

- Port Access: Allows only the MAC address of a device authenticated through the switch 802.1X Port-Based access control.

Eavesdrop Prevention

Configuring port security on a given switch port automatically enables Eavesdrop Prevention for that port. This prevents use of the port to flood unicast packets addressed to MAC addresses unknown to the switch and blocks unauthorized users from eavesdropping on traffic intended for addresses that have aged-out of the switch address table. (Eavesdrop Prevention does not affect multicast and broadcast traffic; the switch floods these two traffic types out a given port regardless of whether port security is enabled on that port.)

Disabling Eavesdrop Prevention

Traffic with an unknown destination address is blocked when port security is configured and Eavesdrop Prevention is enabled. You can disable Eavesdrop Prevention on ports where it may cause problems, such as on ports that are configured to use limited-continuous learning mode. See **Configuring port security** on page 576 for more information on learning modes.

Feature interactions when Eavesdrop Prevention is disabled

The following table explains the various interactions between learning modes and Eavesdrop Prevention when Eavesdrop Prevention is disabled.

NOTE
When the learning mode is "port-access", Eavesdrop Prevention will not be applied to the port. However, it can still be configured or disabled for the port.

Table 46: *Learn – Effect*

Learn mode	Effect
Static	When Eavesdrop Prevention is disabled, the port transmits packets that have unknown destination addresses. The port is secured and only a limited number of static MAC addresses are learned. A device must generate traffic before the MAC address is learned and traffic is forwarded to it.
Continuous	The default. The Eavesdrop Prevention option does not apply because port security is disabled. Ports forward traffic with unknown destination addresses normally.
Port-access	Disabling Eavesdrop Prevention is not applied to the port. There is no change.
Limited-continuous	When Eavesdrop Prevention is disabled, the port transmits packets that have unknown destination addresses. The port is secured; MAC addresses age normally. Eavesdrop Prevention may cause difficulties in learning MAC addresses (as with static MAC addresses) and cause serious traffic issues when a MAC ages out.
Configured	When Eavesdrop Prevention is disabled, the port transmits packets that have unknown destination addresses. The port is secured by a static MAC address. Eavesdrop Prevention should not cause any issues because all valid MAC addresses have been configured.

Syntax:

```
[no] port-security <port-list> eavesdrop-prevention
```

When this option is enabled, the port is prevented from transmitting packets that have unknown destination addresses. Only devices attached to the port receive packets intended for them. This option does not apply to a learning mode of `port-access` or `continuous`. Default: Enabled

The show port-security command displaying Eavesdrop Prevention

```
HP Switch(config)# show port-security

 Port Security

  Port   Learn Mode    | Action           Eavesdrop Prevention
  ------ ------------- + ---------------- --------------------
  1      Continuous    | None             Enabled
  2      Continuous    | None             Enabled
  3      Continuous    | None             Enabled
  4      Continuous    | None             Enabled
  5      Continuous    | None             Enabled
```

MIB Support

The following MIB support is provided for Eavesdrop Prevention.

```
hpSecPtPreventEavesdrop OBJECT-TYPE
SYNTAX INTEGER {
enable (1),
disable (2)
}
MAX-ACCESS read-write
STATUS current
```

```
DESCRIPTION "If enabled on a switch, outbound unknown unicast
packets will not be forwarded out this port. If
enabled on a repeater, outbound unknown unicast
packets for this port will be scrambled.::= { hpSecurePortEntry 5 }
```

Blocked unauthorized traffic

Unless you configure the switch to disable a port on which a security violation is detected, the switch security measures block unauthorized traffic without disabling the port. This implementation enables you to apply the security configuration to ports on which hubs, switches, or other devices are connected, and to maintain security while also maintaining network access to authorized users. For example:

Figure 137: *How port security controls access*

NOTE

Broadcast and Multicast traffic is always allowed, and can be read by intruders connected to a port on which you have configured port security.

Trunk group exclusion

Port security does not operate on either a static or dynamic trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch will reset the port security parameters for those ports to the factory-default configuration. (Ports configured for either Active or Passive LACP, and which are not members of a trunk, can be configured for port security.)

Planning port security

1. Plan your port security configuration and monitoring according to the following:

a. On which ports do you want port security?

b. Which devices (MAC addresses) are authorized on each port?

c. For each port, what security actions do you want? (The switch automatically blocks intruders detected on that port from transmitting to the network.) You can configure the switch to (1) send intrusion alarms to an SNMP management station and to (2) optionally disable the port on which the intrusion was detected.

d. How do you want to learn of the security violation attempts the switch detects? You can use one or more of these methods:

- Through network management (That is, do you want an SNMP trap sent to a net management station when a port detects a security violation attempt?)

- Through the switch Intrusion Log, available through the CLI, menu, and WebAgent

- Through the Event Log (in the menu interface or through the CLI `show log` command)

2. Use the CLI or WebAgent to configure port security operating and address controls.

Port security command options and operation

This section describes the CLI port security command and how the switch acquires and maintains authorized addresses.

NOTE

Use the global configuration level to execute port-security configuration commands.

Displaying port security settings

Syntax:

```
show port-security
```

```
show port-security <port-number>
```

```
show port-security [<port-number>-<port-number>]...[,<port-number>]
```

The CLI uses the same command to provide two types of port security listings:

- All ports on the switch with their Learn Mode and (alarm) Action

- Only the specified ports with their Learn Mode, Address Limit, (alarm) Action, and Authorized Addresses

Without port parameters, `show port-security` displays Operating Control settings for all ports on a switch.

Port security listing (ports 7 and 8 show the default setting)

```
HP Switch(config)# show port-security

 Port Security

  Port Learn Mode  | Action                   Eavesdrop Prevention
  ---- ----------- + ------------------------ --------------------
```

```
1    Continuous  | Send Alarm, Disable Port Enabled
2    Continuous  | Send Alarm, Disable Port Enabled
3    Static      | Send Alarm              Enabled
4    Continuous  | Send Alarm, Disable Port Enabled
5    Continuous  | Send Alarm, Disable Port Enabled
6    Continuous  | Send Alarm, Disable Port Enabled
7    Continuous  | None                    Enabled
8    Continuous  | None                    Enabled
```

With port numbers included in the command, `show port-security` displays Learn Mode, Address Limit, (alarm) Action, and Authorized Addresses for the specified ports on a switch. The following example lists the full port security configuration for a single port:

The port security configuration display for a single port

```
HP Switch(config)# show port-security 3

 Port Security

  Port : 3
  Learn Mode [Continuous] : Static    Address Limit [1] : 1
  Action [None] : None
  Eavesdrop Prevention [Enabled] : Enabled

  Authorized Addresses
  --------------------
  00906d-fdcc00
```

The next example shows the option for entering a range of ports, including a series of non-contiguous ports. Note that no spaces are allowed in the port number portion of the command string:

Entering a range of ports

```
HP Switch(config)# show port-security 1-3,6,8
```

Listing authorized and detected MAC addresses

Syntax:

```
show mac-address [port-list|mac-address|vlan <vid>]
```

Without an optional parameter, `show mac-address` lists the authorized MAC addresses that the switch detects on all ports.

`mac-address`: Lists the specified MAC address with the port on which it is detected as an authorized address.

`port list`: Lists the authorized MAC addresses detected on the specified port(s).

`vlan <vid>` : Lists the authorized MAC addresses detected on ports belonging to the specified VLAN.

Show mac-address outputs

```
HP Switch(config)# show mac-address
 Status and Counters - Port Address Table

  MAC Address    Port  VLAN
  ------------- ----- ----
  00000c-07ac00 7     1
  0000aa-9c09cb 7     1
  000102-f215c7 5     100
        .
```

```
   0018fe-a5e504 1     222

HP Switch(config)# show mac-address 7
 Status and Counters - Port Address Table - 7

  MAC Address    VLANs
  ------------ ------------
  00000c-07ac00 1
  0000aa-9c09cb 1

HP Switch(config)# show mac-address 00000c-07ac00
 Status and Counters - Address Table - 00000c-07ac00

 Port  VLAN
 ----- ----
 5     100

HP Switch(config)# show mac-address vlan 1
 Status and Counters - Address Table - VLAN 1

  MAC Address    Port
  ------------ -----
  00000c-07ac00 1
  000050-53c774 1
        .
  0000aa-9c09cb 1
```

Configuring port security

Procedure

1. Using the CLI, you can do the following:

- Configure port security and edit security settings.

- Add or delete devices from the list of authorized addresses for one or more ports.

- Clear the Intrusion flag on specific ports.

    See **Port security commands**

Port security commands

Using the CLI, you can do the following:

- Configure port security and edit security settings.

- Add or delete devices from the list of authorized addresses for one or more ports.

- Clear the Intrusion flag on specific ports.

Syntax:

```
port-security
<port-list> <learn-mode|address-limit|mac-address|action|clear-intrusion-flag>
```

Options:

```
<port-list>
```

Specifies a list of one or more ports to which the port-security command applies.

`learn-mode <continuous|static|port-access|configured|limited-continuous>`

For the specified port:

- Identifies the method for acquiring authorized addresses.

- On switches covered in this guide, automatically invokes eavesdrop protection, see **Eavesdrop Prevention** on page 571.

`continuous`: (Default): Appears in the factory-default setting or when you execute `no port-security`. Allows the port to learn addresses from the device(s) to which it is connected. In this state, the port accepts traffic from any device(s) to which it is connected. Addresses learned in the learn continuous mode will "age out" and be automatically deleted if they are not used regularly. The default age time is five minutes.

Addresses learned this way appear in the switch and port address tables and age out according to the `MAC Age Interval` in the System Information configuration screen of the Menu interface or the `show system information` listing. You can set the MAC age out time using the CLI, SNMP, Web, or menu interfaces. For more information on the `mac-age-time` command see "Interface Access and System Information" in the management and configuration guide for your switch.

`static`: Enables you to use the `mac-address` parameter to specify the MAC addresses of the devices authorized for a port, and the `address-limit` parameter (explained below) to specify the number of MAC addresses authorized for the port. You can authorize specific devices for the port, while still allowing the port to accept other, non-specified devices until the device limit has been reached. That is, if you enter fewer MAC addresses than you authorized, the port authorizes the remaining addresses in the order in which it automatically learns them.

For example, if you use address-limit to specify three authorized devices, but use `mac-address` to specify only one authorized MAC address, the port adds the one specifically authorized MAC address to its authorized-devices list and the first two additional MAC addresses it detects.

If, for example:

You use `mac-address` to authorize MAC address 0060b0-880a80 for port A4.

You use `address-limit` to allow three devices on port A4 and the port detects these MAC addresses:

- 080090-1362f2

- 00f031-423fc1

- 080071-0c45a1

- 0060b0-880a80 (the address you authorized with the `mac-address` parameter)

In this example port A4 would assume the following list of authorized addresses:

080090-1362f2 (the first address the port detected)

00f031-423fc1 (the second address the port detected)

0060b0-880a80 (the address you authorized with the `mac-address` parameter)

The remaining MAC address detected by the port, 080071-0c45a1, is not allowed and is handled as an intruder. Learned addresses that become authorized do not age-out. See also **Retention of static addresses** on page 579.

CAUTION

Using the `static` parameter with a device limit greater than the number of MAC addresses specified with `mac-address` can allow an unwanted device to become "authorized". This is because the port, to fulfill the number of devices allowed by the `address-limit`

parameter (see below), automatically adds devices it detects until it reaches the specified limit.

NOTE

If 802.1X port-access is configured on a given port, then port-security learn-mode must be set to either `continuous` (the default) or `port-access`.

`port-access`: Enables you to use Port Security with (802.1X) Port-Based Access Control.

`configured`: Must specify which MAC addresses are allowed for this port. Range is 1 (default) to 64 and addresses are not ageable. Addresses are saved across reboots.

`limited-continuous`: Also known as MAC Secure, or "limited" mode. The limited parameter sets a finite limit to the number of learned addresses allowed per port. (You can set the range from 1, the default, to a maximum of 32 MAC addresses which may be learned by each port.)

All addresses are ageable, meaning they are automatically removed from the authorized address list for that port after a certain amount of time. Limited mode and the address limit are saved across reboots, but addresses which had been learned are lost during the reboot process.

Addresses learned in the limited mode are normal addresses learned from the network until the limit is reached, but they are not configurable. (You cannot enter or remove these addresses manually if you are using learn-mode with the limited-continuous option.)

Addresses learned this way appear in the switch and port address tables and age out according to the MAC Age Interval in the System Information configuration screen of the Menu interface or the show system information listing. You can set the MAC age out time using the CLI, SNMP, Web, or menu interfaces. For more on the mac-age-time command, see "Interface Access and System Information" in the management and configuration guide for your switch. To set the learn-mode to limited use this Command syntax:

```
port-security <port-list> learn-mode limited address-limit <1..64> action <none|
send-alarm|send-disable>
```

The default address-limit is `1` but may be set for each port to learn up to 64 addresses.

The default action is none.

To see the list of learned addresses for a port use the command:

```
show mac <port-list>
```

`address-limit <integer>`

When `learn-mode` is set to `static`, `configured`, or `limited-continuous`, the `address-limit` parameter specifies how many authorized devices (MAC addresses) to allow. Range: 1 (the default) to 8 for static and configured modes. For `learn-mode` with the `limited-continuous` option, the range is 1-64 addresses.

Available for `learn-mode` with the, `static`, `configured`, or `limited-continuous` option. Allows up to eight authorized devices (MAC addresses) per port, depending on the value specified in the `address-limit` parameter. The `mac-address limited-continuous` mode allows up to 64 authorized MAC addresses per port.

If you use `mac-address` with `static`, but enter fewer devices than you specified in the `address-limit` field, the port accepts not only your specified devices, but also as many other devices as it takes to reach the device limit. For example, if you specify four devices, but enter only two MAC addresses, the port will accept the first two

non-specified devices it detects, along with the two specifically authorized devices. Learned addresses that become authorized do **not** age-out. See also **Retention of static addresses** on page 579.

```
action <none|send-alarm|send-disable>
```

Specifies whether an SNMP trap is sent to a network management station when Learn Mode is set to static and the port detects an unauthorized device, or when Learn Mode is set to continuous and there is an address change on a port.

`none`: Prevents an SNMP trap from being sent. `none` is the default value.

`send-alarm`: Sends an intrusion alarm. Causes the switch to send an SNMP trap to a network management station.

`send-disable`: Sends alarm and disables the port. Available only in the `static`, `port-access`, `configured`, or `limited learn` modes. Causes the switch to send an SNMP trap to a network management station and disable the port. If you subsequently re-enable the port without clearing the port's intrusion flag, the port blocks further intruders, but the switch will not disable the port again until you reset the intrusion flag. See the Note on **Keeping the intrusion log current by resetting alert flags** on page 594.

For information on configuring the switch for SNMP management, see the management and configuration guide for your switch.

```
clear-intrusion-flag
```

Clears the intrusion flag for a specific port, see **Reading intrusion alerts and resetting alert flags** on page 592.

```
no port-security <port-list> mac-address <mac-addr>[<mac-addr> <mac-addr>]
```

Removes any specified learned MAC addresses from the specified port.

Retention of static addresses

Static MAC addresses do not age-out. MAC addresses learned by using `learn-mode continuous` or `learn-mode limited-continuous` age out according to the currently configured MAC age time. For information on the `mac-age-time` command, see "Interface Access and System Information" in the management and configuration guide for your switch.

Learned addresses

In the following two cases, a port in Static learn mode retains a learned MAC address even if you later reboot the switch or disable port security for that port:

* The port learns a MAC address after you configure the port for Static learn mode in both the startup-config file and the running-config file (by executing the `write memory` command).

* The port learns a MAC address after you configure the port for Static learn mode in only the running-config file and, after the address is learned, you execute `write memory` to configure the startup-config file to match the running-config file.

To remove an address learned using either of the preceding methods, do one of the following:

* Delete the address by using `no port-security <port-number> mac-address <mac-addr>` .

* Download a configuration file that does not include the unwanted MAC address assignment.

* Reset the switch to its factory-default configuration.

Assigned/authorized addresses

If you manually assign a MAC address (using `port-security <port-number> address-list <mac-addr>`) and then execute write memory, the assigned MAC address remains in memory until you do one of the following:

- Delete it by using `no port-security <port-number> mac-address <mac-addr>`

- Download a configuration file that does not include the unwanted MAC address assignment.

- Reset the switch to its factory-default configuration.

Specifying authorized devices and intrusion responses

This example configures port A1 to automatically accept the first device (MAC address) it detects as the only authorized device for that port. (The default device limit is 1.) It also configures the port to send an alarm to a network management station and disable itself if an intruder is detected on the port.

```
HP Switch(config)# port-security a1 learn-mode static
action send-disable
```

The next example does the same as the preceding example, except that it specifies a MAC address of 0c0090-123456 as the authorized device instead of allowing the port to automatically assign the first device it detects as an authorized device.

```
HP Switch(config)# port-security a1 learn-mode static
mac-address 0c0090-123456 action send-disable
```

This example configures port A5 to:

- Allow two MAC addresses, 00c100-7fec00 and 0060b0-889e00, as the authorized devices.

- Send an alarm to a management station if an intruder is detected on the port, but allow the intruder access to the network.

```
HP Switch(config)# port-security a5 learn-mode static
address-limit 2 mac-address 00c100-7fec00 0060b0-889e00
action send-alarm
```

If you manually configure authorized devices (MAC addresses) and/or an alarm action on a port, those settings remain unless you either manually change them or the switch is reset to its factory-default configuration. You can "turn off" authorized devices on a port by configuring the port to continuous Learn Mode, but subsequently reconfiguring the port to static Learn Mode restores those authorized devices.

Adding an authorized device to a port

To simply add a device (MAC address) to a port's existing Authorized Addresses list, enter the port number with the `mac-address` parameter and the device's MAC address. This assumes that Learn Mode is set to `static`

and the Authorized Addresses list is not full (as determined by the current Address Limit value). For example, suppose port A1 allows two authorized devices, but has only one device in its Authorized Address list:

Figure 138: *Adding an authorized device to a port*

With the above configuration for port A1, the following command adds the 0c0090-456456 MAC address as the second authorized address.

```
HP Switch(config)# port-security a1 mac-address 0c0090-
456456
```

After executing the above command, the security configuration for port A1 would be:

Figure 139: *Adding a second authorized device to a port*

(The message `Inconsistent value` appears if the new MAC address exceeds the current Address Limit or specifies a device that is already on the list. Note that if you change a port from static to continuous learn mode, the port retains in memory any authorized addresses it had while in static mode. If you subsequently attempt to convert the port back to static mode with the same authorized address(es), the `Inconsistent value` message appears because the port already has the address(es) in its "Authorized" list.)

If you are adding a device (MAC address) to a port on which the Authorized Addresses list is already full (as controlled by the port's current Address Limit setting), then you must increase the Address Limit in order to add the device, even if you want to replace one device with another. Using the CLI, you can simultaneously increase the limit and add the MAC address with a single command. For example, suppose port A1 allows one authorized device and already has a device listed:

Port security on port A1 with an address limit of "1"

```
HP Switch(config)# show port-security 1
 Port Security

 Port : 1
 Learn Mode [Continuous] : Static Address Limit [1] : 2
 Action [None] : None
 Eavesdrop Prevention [Enabled] : Enabled

 Authorized Addresses
 --------------------
 0c0090-123456
 0c0090-456456
```

To add a second authorized device to port A1, execute a `port-security` command for port A1 that raises the address limit to 2 and specifies the additional device's MAC address. For example:

```
HP Switch(config)# port-security a1 mac-address 0c0090-
456456 address-limit 2
```

Removing a device from the "authorized" list for a port

This command option removes unwanted devices (MAC addresses) from the Authorized Addresses list. (An Authorized Address list is available for each port for which Learn Mode is currently set to "Static". See the Command syntax listing under **Configuring port security** on page 576.

> **CAUTION**
> When learn mode is set to static, the Address Limit (address-limit) parameter controls how many devices are allowed in the Authorized Addresses (`mac-address`) for a given port. If you remove a MAC address from the Authorized Addresses list without also reducing the Address Limit by 1, the port may subsequently detect and accept as authorized a MAC address that you do not intend to include in your Authorized Address list. Thus, if you use the CLI to remove a device that is no longer authorized, it is recommended that you first reduce the Address Limit (`address-limit`) integer by 1, as shown below. This prevents the possibility of the same device or another unauthorized device on the network from automatically being accepted as "authorized" for that port.

To remove a device (MAC address) from the "Authorized" list and when the current number of devices equals the Address Limit value, you should first reduce the Address Limit value by 1, then remove the unwanted device.

> **NOTE**
> You can reduce the address limit below the number of currently authorized addresses on a port. This enables you to subsequently remove a device from the "Authorized" list without opening the possibility for an unwanted device to automatically become authorized.

For example, suppose port A1 is configured as shown below and you want to remove 0c0090-123456 from the Authorized Address list:

Figure 140: *Two authorized addresses on port A1*

```
HP Switch(config)# show port-security 1
 Port Security
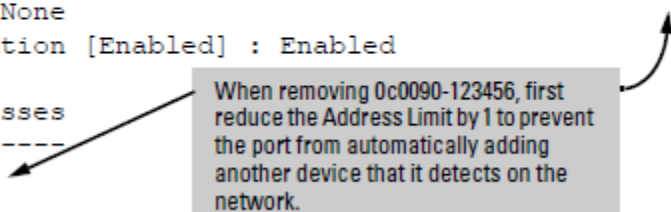
   Port : 1
   Learn Mode [Continuous] : Static          Address Limit [1] : 2
   Action [None] : None
   Eavesdrop Prevention [Enabled] : Enabled

   Authorized Addresses           When removing 0c0090-123456, first
   ------------------             reduce the Address Limit by 1 to prevent
   0c0090-123456                  the port from automatically adding
   0c0090-456456                  another device that it detects on the
                                  network.
```

The following command serves this purpose by removing 0c0090-123456 and reducing the Address Limit to 1:

```
HP Switch(config)# port-security a1 address-limit 1
HP Switch(config)# no port-security a1 mac-address
0c0090-123456
```

The above command sequence results in the following configuration for port A1:

Figure 141: *Port A1 after removing one MAC address*

```
HP Switch(config)# show port-security 1
 Port Security

   Port : 1
   Learn Mode [Continuous] : Static          Address Limit [1] : 1
   Action [None] : None
   Eavesdrop Prevention [Enabled] : Enabled

   Authorized Addresses
   ------------------
   0c0090-456456
```

Clear MAC address table

The following options allow learned MAC addresses to be removed from the MAC address table as follows:

• Remove all MAC addresses.

• Remove all MAC address on a specified VLAN

• Remove all MAC addresses on a port

• Remove a specific MAC address on a specific VLAN

This functionality is also supported by SNMP.

Configuring clearing of learned MAC addresses

Use the following commands to clear learned MAC addresses from a port or list of ports, a specific VLAN, or to clear a specific MAC address from a VLAN.

Syntax:

```
clear mac-address port < port-list >
```

Removes MAC addresses that were learned on the specified port or ports in <*port-list* > . Use `all` to remove all MAC addresses in the MAC address table.

```
HP Switch(config)# clear mac-address port 4-7
```

Syntax:

```
clear mac-address vlan < vid >
```

Removes all MAC addresses that were learned on the specified VLAN.

```
HP Switch(config)# clear mac-address vlan 2
```

Syntax:

```
clear mac-address vlan <vid> mac <mac-addr>
```

Removes the specified MAC address from the specified VLAN.

```
HP Switch(config)# clear mac-address vlan 2 mac 0001e6-b197a8
```

To view the results from clearing a MAC address, use the `show mac-address` command with the appropriate option.

A MAC address cleared from the MAC Address Table

```
HP Switch(config)# show mac-address vlan 2
 Status and Counters - Address Table - VLAN 2

  MAC Address     Located on Port
  ------------    ---------------
  00000c-07ac00   2
  000102-03db12   2
  0001e6-b197a8   2

HP Switch(config)# clear mac-address vlan 2 mac 0001e6-b197a8

HP Switch(config)# show mac-address vlan 2

 Status and Counters - Address Table - VLAN 2

  MAC Address     Located on Port
  ------------    ---------------
  00000c-07ac00   2
  000102-03db12   2
```

MAC Lockdown

Syntax:

```
[no] static-mac {< mac-addr > vlan < vid > interface < port-number >}
```

Locks down a given MAC address and VLAN to a specific port.

A separate command is necessary for each MAC/VLAN pair you wish to lock down. If not specifying a VID, the switch inserts "1".

MAC Lockdown, also known as "static addressing," is the permanent assignment of a given MAC address and VLAN to a specific port on the switch. MAC Lockdown is used to prevent station movement and MAC address hijacking. It also controls address learning on the switch.

Locking down a MAC address on a port and a specific VLAN only restricts the MAC address on that VLAN. The client device with that MAC address is allowed to access other VLANs on the same port or through other ports.

Port security and MAC Lockdown are mutually exclusive on a given port.

How MAC Lockdown works

When a device's MAC address is locked down to a port (typically in a pair with a VLAN) all information sent to that MAC address must go through the locked-down port. If the device is moved to another port it cannot receive data. Traffic to the designated MAC address goes only to the allowed port, whether the device is connected to it or not.

MAC Lockdown is useful for preventing an intruder from "hijacking" a MAC address from a known user in order to steal data. Without MAC Lockdown, this will cause the switch to learn the address on the malicious user's port, allowing the intruder to steal the traffic meant for the legitimate user.

MAC Lockdown ensures that traffic intended for a specific MAC address can only go through the one port which is supposed to be connected to that MAC address. It does not prevent intruders from transmitting packets with the locked MAC address, but it does prevent responses to those packets from going anywhere other than the locked-down port. Thus TCP connections cannot be established. Traffic sent to the locked address cannot be hijacked and directed out the port of the intruder.

If the device (computer, PDA, wireless device) is moved to a different port on the switch (by reconnecting the Ethernet cable or by moving the device to an area using a wireless access point connected to a different port on that same switch), the port will detect that the MAC Address is not on the appropriate port and will continue to send traffic out the port to which the address was locked.

Once a MAC address is configured for one port, you cannot perform port security using the same MAC address on any other port on that same switch.

You cannot lock down a single MAC Address/VLAN pair to more than one port; however you can lock down multiple different MAC Addresses to a single port on the same switch.

Stations can move from the port to which their MAC address is locked to other parts of the network. They can send, but will not receive data if that data must go through the locked down switch. Please note that if the device moves to a distant part of the network where data sent to its MAC address never goes through the locked down switch, it may be possible for the device to have full two-way communication. For full and complete lockdown network-wide all switches must be configured appropriately.

Once you lock down a MAC address/VLAN pair on one port that pair cannot be locked down on a different port.

You cannot perform MAC Lockdown and 802.1X authentication on the same port or on the same MAC address. MAC Lockdown and 802.1X authentication are mutually exclusive.

Lockdown is permitted on static trunks (manually configured link aggregations).

Differences between MAC Lockdown and port security

Because port-security relies upon MAC addresses, it is often confused with the MAC Lockdown feature. However, MAC Lockdown is a completely different feature and is implemented on a different architecture level. Port security maintains a list of allowed MAC addresses on a per-port basis. An address can exist on multiple ports of a switch.

Port security deals with MAC addresses only while MAC Lockdown specifies both a MAC address and a VLAN for lockdown.

MAC Lockdown, on the other hand, is not a "list." It is a global parameter on the switch that takes precedence over any other security mechanism. The MAC Address will only be allowed to communicate using one specific port on the switch.

MAC Lockdown is a good replacement for port security to create tighter control over MAC addresses and which ports they are allowed to use (only one port per MAC Address on the same switch in the case of MAC Lockdown). (You can still use the port for other MAC addresses, but you cannot use the locked down MAC address on other ports.)

Using only port security the MAC Address could still be used on another port on the same switch. MAC Lockdown, on the other hand, is a clear one-to-one relationship between the MAC Address and the port. Once a MAC address has been locked down to a port it cannot be used on another port on the same switch.

The switch does not allow MAC Lockdown and port security on the same port.

MAC Lockdown operating notes

Limits

There is a limit of 500 MAC Lockdowns that you can safely code per switch. To truly lock down a MAC address it would be necessary to use the MAC Lockdown command for every MAC Address and VLAN ID on every switch. In reality few network administrators will go to this length, but it is important to note that just because you have locked down the MAC address and VID for a single switch, the device (or a hacker "spoofing" the MAC address for the device) may still be able to use another switch which hasn't been locked down.

Event Log messages

If someone using a locked down MAC address is attempting to communicate using the wrong port the "move attempt" generates messages in the log file like this:

```
Move attempt (lockdown) logging:

W 10/30/03 21:33:43 maclock: module A: Move 0001e6-1f96c0
to A15 denied

W 10/30/03 21:33:48 maclock: module A: Move 0001e6-1f96c0
to A15 denied

W 10/30/03 21:33:48 maclock: module A: Ceasing move-denied
logs for 5m
```

These messages in the log file can be useful for troubleshooting problems. If you are trying to connect a device which has been locked down to the wrong port, it will not work but it will generate error messages like this to help you determine the problem.

Limiting the frequency of log messages

The first move attempt (or intrusion) is logged as you see in the example above. Subsequent move attempts send a message to the log file also, but message throttling is imposed on the logging on a per-module basis. What this means is that the logging system checks again after the first 5 minutes to see if another attempt has been made to move to the wrong port. If this is the case the log file registers the most recent attempt and then checks again after one hour. If there are no further attempts in that period then it will continue to check every 5 minutes. If another attempt was made during the one hour period then the log resets itself to check once a day. The purpose of rate-limiting the log messaging is to prevent the log file from becoming too full. You can also configure the switch to send the same messages to a server. See "Debug and Messaging Operation" in the management and configuration guide for your switch.

Deploying MAC Lockdown

When you deploy MAC Lockdown you need to consider how you use it within your network topology to ensure security. In some cases where you are using techniques such as "meshing" or Spanning Tree Protocol (STP) to

speed up network performance by providing multiple paths for devices, using MAC Lockdown either will not work or else it defeats the purpose of having multiple data paths.

The purpose of using MAC Lockdown is to prevent a malicious user from "hijacking" an approved MAC address so they can steal data traffic being sent to that address.

As we have seen, MAC Lockdown can help prevent this type of hijacking by making sure that all traffic to a specific MAC address goes only to the proper port on a switch which is supposed to be connected to the real device bearing that MAC address.

However, you can run into trouble if you incorrectly try to deploy MAC Lockdown in a network that uses multiple path technology, like Spanning Tree or "mesh networks."

Let's examine a good use of MAC Lockdown within a network to ensure security first.

Figure 142: *MAC Lockdown deployed at the network edge provides security*

Basic MAC Lockdown deployment

In the Model Network Topology shown above, the switches that are connected to the edge of the network each have one and only one connection to the core network. This means each switch has only one path by which data can travel to Server A. You can use MAC Lockdown to specify that all traffic intended for Server A's MAC Address must go through the one port on the edge switches. That way, users on the edge can still use other network resources, but they cannot "spoof" Server A and hijack data traffic which is intended for that server alone.

The key points for this Model Topology are:

- The Core Network is separated from the edge by the use of switches which have been "locked down" for security.

- All switches connected to the edge (outside users) each have only one port they can use to connect to the Core Network and then to Server A.

- Each switch has been configured with MAC Lockdown so that the MAC Address for Server A has been locked down to one port per switch that can connect to the Core and Server A.

Using this setup Server A can be moved around within the core network, and yet MAC Lockdown will still prevent a user at the edge from hijacking its address and stealing data.

Please note that in this scenario a user with bad intentions at the edge can still "spoof" the address for Server A and send out data packets that look as though they came from Server A. The good news is that because MAC Lockdown has been used on the switches on the edge, any traffic that is sent back to Server A will be sent to the proper MAC Address because MAC Lockdown has been used. The switches at the edge will not send Server A's data packets anywhere but the port connected to Server A. (Data would not be allowed to go beyond the edge switches.)

> **CAUTION**
>
> Using MAC Lockdown still does not protect against a hijacker **within the core**! In order to protect against someone spoofing the MAC Address for Server A inside the Core Network, you would have to lock down each and every switch inside the Core Network as well, not just on the edge.

Problems using MAC Lockdown in networks with multiple paths

Now let's take a look at a network topology in which the use of MAC Lockdown presents a problem. In the next figure, Switch 1 (on the bottom-left) is located at the edge of the network where there is a mixed audience that might contain hackers or other malicious users. Switch 1 has two paths it could use to connect to Server A. If you

try to use MAC Lockdown here to make sure that all data to Server A is "locked down" to one path, connectivity problems would be the result since both paths need to be usable in case one of them fails.

Figure 143: *Connectivity problems using MAC lockdown with multiple paths*

The resultant connectivity issues would prevent you from locking down Server A to Switch 1. And when you remove the MAC Lockdown from Switch 1 (to prevent broadcast storms or other connectivity issues), you then open the network to security problems. The use of MAC Lockdown as shown in the above figure would defeat the purpose of using MSTP or having an alternate path.

Technologies such as MSTP or "meshing" are primarily intended for an internal campus network environment in which all users are trusted. MSTP and "meshing" do not work well with MAC Lockdown.

If you deploy MAC Lockdown as shown in the Model Topology in **Deploying MAC Lockdown**, you should have no problems with either security or connectivity.

MAC Lockout

Syntax:

```
[no] lockout-mac <mac-address>
```

Locks a MAC address out on the switch and all VLANs.

MAC Lockout involves configuring a MAC address on all ports and VLANs for a switch, so that any traffic to or from the "locked-out" MAC address is dropped: all data packets addressed to or from the given address are stopped by the switch. MAC Lockout is like a simple blacklist.

MAC Lockout is implemented on a per switch assignment. To use it you must know the MAC Address to block. To fully lock out a MAC address from the network it is necessary to use the MAC Lockout command on all switches.

How MAC Lockout works

Let us say a customer knows there are unauthorized wireless clients who must not have access to the network. The network administrator "locks out" the MAC addresses for the wireless clients by using the MAC Lockout command (`lockout-mac <mac-address>`). When the wireless clients then attempt to use the network, the switch recognizes the intruding MAC addresses and prevents them from sending or receiving data on that network.

Unwanted MAC Addresses can be disallowed on all switch ports with a single command. You do not have to configure every port—perform the command on the switch and it is effective for all ports.

MAC Lockout overrides MAC Lockdown, port security, and 802.1X authentication.

You cannot use MAC Lockout to lock:

- Broadcast or Multicast Addresses (Switches do not learn these addresses)

- Switch Agents (The switch MAC Address)

A MAC address can exist on many different VLANs, so a lockout MAC address must be added to the MAC table as a drop. As many addresses can quickly fill the MAC table, restrictions are placed on the number of lockout MAC addresses based on the number of VLANs configured.

VLANs configured	Number of MAC lockout addresses	Total number of MAC addresses
1-8	200	1,600
9-16	100	1,600
17-256	64	16,384
257-1024	16	16,384
1025-2048	8	16,384

There are limits for the number of VLANs, Multicast Filters, and Lockout MACs that can be configured concurrently as all use MAC table entries.

Table 47: *Limits on Lockout MACs*

# VLANs	# Multicast filters	# Lockout MACs
<=1024	16	16
1025-2048	8	8

If someone using a locked out MAC address tries to send data through the switch, a message is generated in the log file:

```
Lockout logging format:

W 10/30/03 21:35:15 maclock: 0001e6-1f96c0 detected on port 15

W 10/30/03 21:35:18 maclock: 0001e6-1f96c0 detected on port 15

W 10/30/03 21:35:18 maclock: Ceasing lock-out logs for 5m
```

As with MAC Lockdown, a rate limiting algorithm is used on the log file so that it does not become clogged with error messages. See **Limiting the frequency of log messages** on page 586.

User-based lockout compliance

When a specified number of unsuccessful authentication attempts has occurred, remote users can be locked out. If the system-configured usernames (manager, operator, and local users) are not unique, the user-based lockout delay feature cannot be enabled. If lockout-delay with the user-based-lockout feature is enabled, the manager, operator, or local user creation fails for duplicate usernames. Locked out users are unlocked after reboot.

aaa authentication

Syntax

```
aaa authentication user-based-lockout

no aaa authentication user-based-lockout

aaa authentication lockout-delay <delay-time>
```

Description

Locks out users based on their usernames. RADIUS and TACACS users are locked out only if they logged in to the switch successfully at least once.

The no form unlocks currently locked-out users.

Context

```
config
```

Parameters

delay-time

 Delay time in seconds for user-based lockout to begin. Range: 0 to 3600.

Example

Initiate user-based lockout:

```
switch(config)# aaa authentication user-based-lockout
```

Example

Initiate user-based lockout after a 10-minute delay.

```
aaa authentication lockout-delay 600
```

aaa authentication unlock

Syntax

```
aaa authentication unlock user-name <USER-NAME>
```

Description

Unlock the specified locked user if user-based-lockout is enabled.

Parameters

USER-NAME

A valid username that is locked-out.

Example

Unlock a specific username.

```
switch(config)# aaa authentication unlock user-name manager
```

show authentication

Syntax

```
show authentication locked-out-users
```

Description

Show all users who are in a locked-out state.

Context

```
config
```

Example

Show all locked users:

```
switch(config)# show authentication locked-out-users
```

Port security and MAC Lockout

MAC Lockout is independent of port-security and in fact will override it. MAC Lockout is preferable to port-security to stop access from known devices because it can be configured for all ports on the switch with one command.

It is possible to use MAC Lockout in conjunction with port-security. You can use MAC Lockout to lock out a single address—deny access to a specific device—but still allow the switch some flexibility in learning other MAC Addresses. Be careful if you use both together, however:

- If a MAC Address is locked out and appears in a static learn table in port-security, the apparently "authorized" address will still be locked out anyway.

- MAC entry configurations set by port security will be kept even if MAC Lockout is configured and the original port security settings will be honored once the Lockout is removed.

- A port security static address is permitted to be a lockout address. In that case (MAC Lockout), the address will be locked out (SA/DA drop) even though it's an "authorized" address from the perspective of port security.

- When MAC Lockout entries are deleted, port security will then re-learn the address as needed later on.

Reading intrusion alerts and resetting alert flags

Notice of security violations

When the switch detects an intrusion on a port, it sets an "alert flag" for that port and makes the intrusion information available as described below. While the switch can detect additional intrusions for the same port, it does not list the next chronological intrusion for that port in the Intrusion Log until the alert flag for that port has been reset.

When a security violation occurs on a port configured for Port Security, the switch responds in the following ways to notify you:

- The switch sets an alert flag for that port. This flag remains set until:

  ◦ You use either the CLI, menu interface, or WebAgent to reset the flag.

  ◦ The switch is reset to its factory default configuration.

- The switch enables notification of the intrusion through the following means:

  ◦ In the CLI:

    – The `show port-security intrusion-log` command displays the Intrusion Log.

    – The `log` command displays the Event Log.

  ◦ In the menu interface:

    – The Port Status screen includes a per-port intrusion alert.

    – The Event Log includes per-port entries for security violations.

  ◦ In the WebAgent:

    – The Alert Log includes entries for per-port security violations.

    – The Intrusion Log lists per-port security violation entries.

  ◦ In network management applications using an SNMP trap sent to a network management station.

How the intrusion log operates

When the switch detects an intrusion attempt on a port, it enters a record of this event in the Intrusion Log. No further intrusion attempts on that port will appear in the Log until you acknowledge the earlier intrusion event by resetting the alert flag.

The Intrusion Log lists the 20 most recently detected security violation attempts, regardless of whether the alert flags for these attempts have been reset. This gives you a history of past intrusion attempts. Thus, for example, if there is an intrusion alert for port A1 and the Intrusion Log shows two or more entries for port 1, only the most recent entry has not been acknowledged (by resetting the alert flag). The other entries give you a history of past intrusions detected on port A1.

Multiple intrusion log entries for the same port

```
HP Switch(config)# show port-security intrusion-log
 Status and Counters - Intrusion Log

 Port  MAC Address      Date / Time
```

```
----- ------------- --------------------------
1      080009-e93d4f   03/07/11 21:09:34
1      080009-e93d4f   03/07/11 10:18:43
```

The log shows the most recent intrusion at the top of the listing. You cannot delete Intrusion Log entries (unless you reset the switch to its factory-default configuration). Instead, if the log is filled when the switch detects a new intrusion, the oldest entry is dropped off the listing and the newest entry appears at the top of the listing.

Keeping the intrusion log current by resetting alert flags

When a violation occurs on a port, an alert flag is set for that port and the violation is entered in the Intrusion Log. The switch can detect and handle subsequent intrusions on that port, but will not log another intrusion on the port until you reset the alert flag for either all ports or for the individual port.

NOTE

On a given port, if the intrusion action is to send an SNMP trap and then disable the port (`send-disable`), and an intruder is detected on the port, then the switch sends an SNMP trap, sets the port's alert flag, and disables the port. If you re-enable the port without resetting the port's alert flag, then the port operates as follows:

- The port comes up and will block traffic from unauthorized devices it detects.

- If the port detects another intruder, it will send another SNMP trap, but will not become disabled again unless you first reset the port's intrusion flag.

This operation enables the port to continue passing traffic for authorized devices while you take the time to locate and eliminate the intruder. Otherwise, the presence of an intruder could cause the switch to repeatedly disable the port.

Checking for intrusions, listing intrusion alerts, and resetting alert flags (Menu)

The menu interface indicates per-port intrusions in the Port Status screen, and provides details and the reset function in the Intrusion Log screen.

Procedure

1. From the Main Menu select the following:

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

a. `1. Status and Counters`

b. `4. Port Status`

Figure 144: *Port status screen with intrusion alert on port 3*

2. Type [I] (Intrusion log) to display the Intrusion Log.

Figure 145: *The Intrusion Log display*

This example shows two intrusions for port 3 and one intrusion for port 1. In this case, only the most recent intrusion at port 3 has not been acknowledged (reset). This is indicated by the following:

• Because the Port Status screen does not indicate an intrusion for port 1, the alert flag for the intrusion on port 1 has already been reset.

• Since the switch can show only one uncleared intrusion per port, the alert flag for the older intrusion for port A3 in this example has also been previously reset.

The intrusion log holds up to 20 intrusion records and deletes an intrusion record only when the log becomes full and a new intrusion is subsequently detected.

| NOTE | The "prior to " text in the record for the earliest intrusion means that a switch reset occurred at the indicated time and that the intrusion occurred prior to the reset. |

3. To acknowledge the most recent intrusion entry on port 3 and enable the switch to enter a subsequently detected intrusion on this port, type [R] For Reset alert flags.

Note that if there are unacknowledged intrusions on two or more ports, this step resets the alert flags for all such ports.

If you then re-display the port status screen, you will see that the Intrusion Alert entry for port 3 has changed to "No". That is, your evidence that the Intrusion Alert flag has been acknowledged (reset) is that the Intrusion Alert column in the port status display no longer shows "Yes" for the port on which the intrusion occurred (port 3 in this example). (Because the Intrusion Log provides a history of the last 20 intrusions detected by the switch, resetting the alert flags does not change its content.

Checking for intrusions, listing intrusion alerts, and resetting alert flags (CLI)

The following commands display port status, including whether there are intrusion alerts for any ports, list the last 20 intrusions, and either reset the alert flag on all ports or for a specific port for which an intrusion was detected. The record of the intrusion remains in the log. For more information, see **Operating notes for port security**.

Syntax:

```
show interfaces brief
```

List intrusion alert status (and other port status information)'.

```
show port-security intrusion-log
```

List intrusion log content.

```
clear intrusion-flags
```

Clear intrusion flags on all ports.

```
port-security <port-number> clear-intrusion-flag
```

Clear the intrusion flag on one or more specific ports.

Example:

In the following example, executing `show interfaces brief` lists the switch port status, indicating an intrusion alert on port 1.

Figure 146: *An unacknowledged intrusion alert in a port status display*

```
HP Switch(config)# show int brief
                                          Intrusion Alert on port 5.
 Status and Counters - Port Status

              | Intrusion                               MDI  Flow Bcast
  Port  Type  | Alert      Enabled Status Mode          Mode Ctrl Limit
  ----- ------ + -------- ----- ------ ---------- ---- ---- -----
  1     100/1000T | Yes       Yes     Up     1000FDx       MDI  off  0
  2     100/1000T | No        Yes     Up     1000FDx       Auto off  0
  3     100/1000T | No        Yes     Up     1000FDx       Auto off  0
  4     100/1000T | No        Yes     Up     1000FDx       Auto off  0
```

To see the details of the intrusion, enter the `show port-security intrusion-log` command. For example:

Figure 147: *The intrusion log with multiple entries for the same port*

```
                    HP Switch(config)# show port-security intrusion-
MAC Address of latest   log                                          Dates and Times of
Intruder on Port 1       Status and Counters - Intrusion Log         Intrusions

                    Port  MAC Address          Date / Time
Earlier intrusions on    ---- ------------   ------------------------
port 1 that have already
been cleared (that is,   1     080009-e93d4f          03/07/11 21:09:34
the Alert Flag has been  1     080009-21ae84          03/07/11 17:26:27
reset at least twice     1     080009-e93d4f prior to 03/07/11 17:18:43
before the most recent
intrusion occurred.      0 secs
                         0 secs
                         35 mins
                         43 mins
```

The above example shows three intrusions for port 1. Since the switch can show only one uncleared intrusion per port, the older two intrusions in this example have already been cleared by earlier use of the `clear intrusion-log` or the `port-security <port-list> clear-intrusion-flag` command. The intrusion log holds up to 20 intrusion records, and deletes intrusion records only when the log becomes full and new intrusions are subsequently added. The "prior to " text in the record for the third intrusion means that a switch reset occurred at the indicated time and that the intrusion occurred prior to the reset.

To clear the intrusion from port 1 and enable the switch to enter any subsequent intrusion for port 1 in the Intrusion Log, execute the port-security `clear-intrusion-flag` command. If you then re-display the port status screen, you will see that the Intrusion Alert entry for port 1 has changed to "`No`". (Executing `show port-`

`security intrusion-log` again will result in the same display as above, and does not include the Intrusion Alert status.)

```
HP Switch(config)# port-security 1 clear-intrusion-flag
HP Switch(config)# show interfaces brief
```

Figure 148: *Port status screen after alert flags reset*

```
HP Switch(config)# show interfaces brief

  Status and Counters - Port Status                Intrusion Alert on port 1 is now cleared.

                   | Intrusion                           MDI   Flow Bcast
   Port  Type      | Alert     Enabled Status Mode        Mode  Ctrl Limit
   ----- --------- + --------- ------- ------ ----------  ----  ---- -----
   1     10/100TX  | No          Yes     Up     100FDx     MDI   off  0
   2     10/100TX  | No          Yes     Down   10FDx      MDI   off  0
   3     10/100TX  | No          Yes     Down   10FDx      MDIX  off  0
```

For more on clearing intrusions, see **Keeping the intrusion log current by resetting alert flags** on page 594.

Using the Event Log to find intrusion alerts (CLI)

The Event Log lists port security intrusions as:

`W MM/DD/YY HH:MM:SS FFI: port A3 — Security Violation`

where "W" is the severity level of the log entry and `FFI` is the system module that generated the entry. For further information, display the Intrusion Log, as shown below.

From the manager or Configuration level:

Syntax:

`log <search-text>`

For `<search-text>` , use `ffi`, `security`, or `violation`.

Figure 149: *Log listing with and without detected security violations*

```
Log Listing with
Security Violation        HP Switch(config)# log security                        Log Command
Detected                    Keys:    W=Warning    I=Information                  with
                                     M=Major      D=Debug E=Error                "security" for
                            ----  Event Log listing: Events Since Boot  ----     Search String
                            W 08/01/02 01:18:15 FFI: port 2 - Security Violation
                            W 08/01/02 04:28:08 FFI: port 2 - Security Violation
                            ----  Bottom of Log : Events Listed = 2   ----

Log Listing with No
Security Violation        HP Switch(config)# log security
Detected                    Keys:    W=Warning    I=Information
                                     M=Major      D=Debug E=Error
                            ----  Event Log listing: Events Since Boot  ----
                            ----  Bottom of Log : Events Listed = 0   ----
```

For more Event Log information, see "Using the Event Log to identify problem sources" in the management and configuration guide for your switch.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Operating notes for port security

Proxy Web servers

If you are using the WebAgent through a switch port configured for Static port security, and your browser access is through a proxy web server, then it is necessary to do the following:

- Enter your PC or workstation MAC address in the port's Authorized Addresses list.

- Enter your PC or workstation's IP address in the switch IP Authorized managers list. See **Using Authorized IP Managers** on page 600.

Without both of the above configured, the switch detects only the proxy server's MAC address, and not your PC or workstation MAC address, and interprets your connection as unauthorized.

"Prior to" entries in the Intrusion Log

If you reset the switch (using the Reset button, Device Reset, or Reboot Switch), the Intrusion Log will list the time of all currently logged intrusions as "prior to" the time of the reset.

Alert flag status for entries forced off of the Intrusion Log

If the Intrusion Log is full of entries for which the alert flags have not been reset, a new intrusion will cause the oldest entry to drop off the list, but will not change the alert flag status for the port referenced in the dropped entry. This means that, even if an entry is forced off of the Intrusion Log, no new intrusions can be logged on the port referenced in that entry until you reset the alert flags.

LACP not available on ports configured for port security

To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables port security on that port. For example:

```
HP Switch(config)# port-security e 7 learn-mode static
address-limit 2
LACP has been disabled on secured port(s).
HP Switch(config)#
```

The switch will not allow you to configure LACP on a port on which port security is enabled. For example:

```
HP Switch(config)# int e 7 lacp passive
Error configuring port 7: LACP and port security cannot
be run together.
HP Switch(config)#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.

Introduction

The Authorized IP Managers feature uses IP addresses and masks to determine which stations (PCs or workstations) can access the switch through the network. This covers access through the following means:

- Telnet and other terminal emulation applications

- The WebAgent

- SSH

- SNMP versions 1, 2 and 3 (with a correct community name)

- TFTP

When configured in the switch, the Authorized IP Managers feature takes precedence over local passwords, TACACS+, and RADIUS. This means that the IP address of a networked management device must be authorized before the switch will attempt to authenticate the device by invoking any other access security features. If the Authorized IP Managers feature disallows access to the device, then access is denied. Thus, with authorized IP managers configured, having the correct passwords is not sufficient for accessing the switch through the network unless the station attempting access is also included in the switch Authorized IP Managers configuration.

Use Authorized IP Managers along with other access security features to provide a more comprehensive security fabric than if you use only one or two security options.

NOTE	When no Authorized IP Manager rules are configured, the access method feature is disabled and access is not denied.

For each authorized manager address, you can configure either of these access levels:

- **Manager**

    Enables full access to all screens for viewing, configuration, and all other operations available.

- **Operator**

    Allows read-only access. (This is the same access that the switch allows for the operator-level password feature.)

Configure up to 100 authorized manager entries, where each entry applies to either a single management station or a group of stations.

CAUTION	Configuring Authorized IP Managers does not protect access to the switch through a modem or direct connection to the Console (RS-232) port. Also, if an unauthorized station "spoofs" an authorized IP address, it can gain management access to the switch even though a duplicate IP address condition exists. For these reasons, you should enhance your network security by keeping physical access to the switch restricted to authorized personnel, using the user name/password and other security features available in the switch, and preventing unauthorized access to data on your management stations.

Defining authorized management stations

- **Authorizing Single Stations:**

  The table entry authorizes a single management station to have IP access to the switch. To use this method, just enter the IP address of an authorized management station in the Authorized manager IP column, and leave the IP Mask set to `255.255.255.255`. This is the easiest way to use the Authorized managers feature. For more on this topic, see **Configuring one station per Authorized manager IP entry** on page 606.

- **Authorizing Multiple Stations:**

  The table entry uses the IP Mask to authorize access to the switch from a defined group of stations. This is useful if you want to easily authorize several stations to have access to the switch without having to type in an entry for every station. All stations in the group defined by the one Authorized manager IP table entry and its associated IP mask will have the same access level—manager or operator. For more on this topic, see **Configuring multiple stations per Authorized manager IP entry** on page 606.

To configure the switch for authorized manager access, enter the appropriate **Authorized manager IP** value, specify an **IP Mask**, and select either `manager` or `operator` for the **Access Level**. The IP Mask determines how the Authorized manager IP value is used to allow or deny access to the switch by a management station.

> **NOTE**
>
> If the management VLAN is configured, access can only be on that VLAN.

Overview of IP mask operation

The default IP Mask is 255.255.255.255 and allows switch access only to a station having an IP address that is identical to the Authorized manager IP parameter value. ("255" in an octet of the mask means that only the exact value in the corresponding octet of the Authorized manager IP parameter is allowed in the IP address of an authorized management station.) However, you can alter the mask and the Authorized manager IP parameter to specify ranges of authorized IP addresses. For example, a mask of `255.255.255.0` and any value for the Authorized manager IP parameter allows a range of 0 through 255 in the 4th octet of the authorized IP address, which enables a block of up to 254 IP addresses for IP management access (excluding 0 for the network and 255 for broadcasts). A mask of `255.255.255.252` uses the 4th octet of a given Authorized manager IP address to authorize four IP addresses for management station access. The details on how to use IP masks are provided under **Configuring one station per Authorized manager IP entry** on page 606.

> **NOTE**
>
> The IP Mask is a method for recognizing whether a given IP address is authorized for management access to the switch. This mask serves a different purpose than IP subnet masks and is applied in a different manner.

Viewing and configuring IP Authorized managers (Menu)

Only IPv4 is supported when using the menu to set the management access method.

From the console Main Menu, select:

```
2. Switch Configuration...
```

Figure 150: *Add an authorized manager entry*

```
==========================- CONSOLE - MANAGER MODE -==========================
                    Switch Configuration - IP Managers

   Authorized Manager IP          IP Mask            Access Level  Access Method
   ------------------------     ----------------------  ------------   -------------
   10.10.240.2                  255.255.255.255      Manager       all
   10.10.245.3                  255.255.255.255      Operator      ssh
   10.10.246.200                255.255.255.255      Operator      tftp
   10.10.245.30                 255.255.255.0        Operator      ssh


                                        1.  Select Add to add an authorized manager
                                            to the list.


   Actions->   Back      Add      Edit      Delete      Help
```

Figure 151: *Edit menu for authorized IP managers*

```
==========================- CONSOLE - MANAGER MODE -==========================
                    Switch Configuration - IP Managers

   Authorized Manager IP: 10.10.245.3        Enter an Authorized Manager IP address here.
   IP Mask [255.255.255.255]:255.255.255.255
   Access Level:Operator                     Use the default mask to allow access by one
   Access Method:ssh                         management device, or edit the mask to allow
                                             access by a block of management devices. See
                                             "Building IP Masks" on page 14-10.

                                             Use the Space bar to select Manager or Operator
                                             access.

                                             Press [Enter], then [S] (for Save) to configure the
   Actions->   Back      Add      Edit      Delete   IP Authorized Manager entry.
```

Editing or deleting an Authorized manager entry (Menu)

Go to the IP managers List screen **Figure 150: Add an authorized manager entry** on page 602, highlight the desired entry, and press `[E]` (for `Edit`) or `[D]` (for `Delete`).

Viewing and configuring IP Authorized managers (CLI)

Listing the switch's current IP Authorized manager(s)

Use the `show ip authorized-managers` command to list IP stations authorized to access the switch. For example:

The show authorized-managers command with access method configured

```
HP Switch(config)# show ip authorized-manager

IPV4 Authorized managers
------------------------

 Address : 10.10.10.10
 Mask    : 255.255.255.255
 Access  : manager
```

Configuring IP Authorized managers for the switch (CLI)

See the IPv6 configuration guide for information about Authorized IP manager configuration with IPv6 addresses.

Syntax:

```
[no] ip authorized-managers <ip-address> <ip-mask>> access [manager|operator]
```

```
access-method [all|ssh|telnet|web|snmp|tftp]
```

```
[no] ipv6 authorized-managers <ip-address> <ip-mask> access [manager|operator]
```

```
access-method [all|ssh|telnet|web|snmp|tftp]
```

Configures one or more authorized IP addresses.

```
access [manager|operator]
```

Configures the privilege level for `<ip-address>`. Applies only to access through telnet, SSH, SNMPv1, SNMPv2c, and SNMPv3.

Default: manager

```
access-method [all|ssh|telnet|web|snmp|tftp]
```

Configures access levels by access method and IP address. Each management method can have its own set of authorized managers.

Default: all

Configuring IP authorized manager access method SSH

```
HP Switch(config)# ip authorized-managers 10.10.10.2 255.255.255.255 manager
                   access-method ssh
```

To authorize manager access

This command authorizes manager-level access for any station with an IP address of 10.28.227.0 through 10.28.227.255:

```
HP Switch(config)# ip authorized-managers 10.28.227.101
255.255.255.0 access manager
```

Similarly, the next command authorizes manager-level access for any station having an IP address of 10.28.227.101 through 103:

```
HP Switch(config)# ip authorized-managers 10.28.227.101
255.255.255.252 access manager
```

If you omit the <*mask bits*> when adding a new authorized manager, the switch automatically uses `255.255.255.255`. If you do not specify either manager or operator access, the switch assigns the manager access.

To edit an existing manager access entry

To change the mask or access level for an existing entry, use the entry's IP address and enter the new value(s). (Notice that any parameters not included in the command will be set to their default.):

```
HP Switch(config)# ip authorized-managers
10.28.227.101 255.255.255.0 access operator
```

The above command replaces the existing mask and access level for IP address 10.28.227.101 with `255.255.255.0` and `operator`.

The following command replaces the existing mask and access level for IP address 10.28.227.101 with `255.255.255.255` and `manager` (the defaults) because the command does not specify either of these parameters.

```
HP Switch(config)# ip authorized-managers 10.28.227.101
```

To delete an authorized manager entry

This command uses the IP address of the authorized manager you want to delete:

```
HP Switch(config)# no ip authorized-managers
10.28.227.101
```

Configuring IP Authorized managers (WebAgent)

In the WebAgent you can configure IP Authorized managers as described below.

To Add, Modify, or Delete an IP Authorized manager address:

1. In the navigation tree, click **Security**.

2. Click **IP Authorization**.

3. Click **Add Address** to add an IP Authorized manager. Enter the appropriate parameter settings for the operation you want.

4. To delete an IP authorized manager, select the Authorized Address and click **Delete**.

5. To change IP Authorization parameters, click **Change** in the IP Authorization Details box. Enter the information and click **Save**.

Figure 152: *Configuring Authorized manager access method in the WebAgent*

Web proxy servers

If you use the WebAgent to access the switch from an authorized IP manager station, it is highly recommended that you avoid using a web proxy server in the path between the station and the switch. This is because switch access through a web proxy server requires that you first add the web proxy server to the Authorized manager IP list. This reduces security by opening switch access to anyone who uses the web proxy server.

How to eliminate the web proxy server

There are two ways to eliminate a web proxy server from the path between a station and the switch:

1. Add the IP address or DNS name of the switch to the non-proxy or Exceptions list in the web browser interface used on the authorized station.

2. If you don't need proxy server access on the authorized station, disable the proxy server feature in the station's web browser interface.

> **NOTE**
>
> IP or MAC authentication can be used without a web proxy server.

Using a web proxy server to access the WebAgent

> **CAUTION**
>
> This is NOT recommended. Using a web proxy server between the stations and the switch poses a security risk. If the station uses a web proxy server to connect to the switch, any proxy user can access the switch.

If it is necessary to use the WebAgent and your browser access is through a web proxy server, perform these steps:

1. Enter the web proxy server's MAC address in the port's Authorized Addresses list.

2. Enter the web proxy server's IP address in the switch's IP Authorized managers list.

You must perform both of these steps or the switch only detects the proxy server's MAC address and IP address instead of your workstation addresses, and your connection is considered unauthorized.

Building IP Masks

The IP Mask parameter controls how the switch uses an Authorized manager IP value to recognize the IP addresses of authorized manager stations on your network.

Configuring one station per Authorized manager IP entry

This is the easiest way to apply a mask. If you have ten or fewer management and/or operator stations, you can configure them by adding the address of each to the Authorized manager IP list with `255.255.255.255` for the corresponding mask. For example, as shown in **The show authorized-managers command with access method configured** on page 603, if you configure an IP address of `10.28.227.125` with an IP mask of `255.255.255.255`, only a station having an IP address of `10.28.227.125` has management access to the switch.

Table 48: *Analysis of IP Mask for Single-Station Entries*

	1st Octet	2nd Octet	3rd Octet	4th Octet	manager-Level or operator-Level Device Access
IP Mask	255	255	255	255	The "255" in each octet of the mask specifies that only the exact value in that octet of the corresponding IP address is allowed. This mask allows management access only to a station having an IP address of 10.33.248.5.
Authorized manager IP	10	28	227	125	

Configuring multiple stations per Authorized manager IP entry

The mask determines whether the IP address of a station on the network meets the criteria you specify. That is, for a given Authorized manager entry, the switch applies the IP mask to the IP address you specify to determine a range of authorized IP addresses for management access. As described above, that range can be as small as one IP address (if `255` is set for all octets in the mask), or can include multiple IP addresses (if one or more octets in the mask are set to less than `255`).

If a bit in an octet of the mask is "on" (set to 1), then the corresponding bit in the IP address of a potentially authorized station must match the same bit in the IP address you entered in the Authorized manager IP list. Conversely, if a bit in an octet of the mask is "off" (set to 0), then the corresponding bit in the IP address of a potentially authorized station on the network does not have to match its counterpart in the IP address you entered in the Authorized manager IP list. Thus, in the example shown above, a "255" in an IP Mask octet (all bits in the octet are "on") means only one value is allowed for that octet—the value you specify in the corresponding octet of the Authorized manager IP list. A "0" (all bits in the octet are "off") means that any value from 0 to 255 is allowed in the corresponding octet in the IP address of an authorized station. You can also specify a series of values that are a subset of the 0-255 range by using a value that is greater than 0, but less than 255.

Table 49: *Analysis of IP Mask for Multiple-Station Entries*

	1st Octet	2nd Octet	3rd Octet	4th Octet	manager-Level or operator-Level Device Access
IP Mask	255	255	255	0	The "255" in the first three octets of the mask specify that only the exact value in the octet of the corresponding IP address is allowed. However, the zero (0) in the 4th octet of the mask allows any value between 0 and 255 in that octet of the corresponding IP address. This mask allows switch access to any device having an IP address of 10.28.227.*xxx*, where *xxx* is any value from 0 to 255.
Authorized manager IP	10	28	227	125	

Table Continued

	1st Octet	2nd Octet	3rd Octet	4th Octet	manager-Level or operator-Level Device Access
IP Mask	255	255	255	249	In this example (**Figure 153: How the Bitmap in the IP Mask defines Authorized manager addresses** on page 608, below), the IP mask allows a group of up to 4 management stations to access the switch. This is useful if the only devices in the IP address group allowed by the mask are management stations. The "249" in the 4th octet means that bits 0 and 3 - 7 of the 4th octet are fixed. Conversely, bits 1 and 2 of the 4th octet are variable. Any value that matches the authorized IP address settings for the fixed bits is allowed for the purposes of IP management station access to the switch. Thus, any management station having an IP address of 10.28.227.121, 123, 125, or 127 can access the switch.
Authorized manager IP	10	28	227	125	

Figure 153: *How the Bitmap in the IP Mask defines Authorized manager addresses*

4th Octet of IP Mask: 249
4th Octet of Authorized IP Address: 5

Bit Numbers	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
Bit Values	128	64	32	16	8	4	2	1

4th Octet of IP Mask (249)

4th Octet of IP Authorized Address (125)

Bits 1 and 2 in the mask are "off", and bits 0 and 3 - 7 are "on", creating a value of 249 in the 4th octet. Where a mask bit is "on", the corresponding bit setting in the address of a potentially authorized station must match the IP Authorized Address setting for that same bit. Where a mask bit is "off" the corresponding bit setting in the address can be either "on" or "off". In this example, in order for a station to be authorized to access the switch:

- The first three octets of the station's IP address must match the Authorized IP Address.
- Bit 0 and Bits 3 through 6 of the 4th octet in the station's address must be "on" (value = 1).
- Bit 7 of the 4th octet in the station's address must be "off" (value = 0).
- Bits 1 and 2 can be either "on" or "off".

This means that stations with the IP address 13.28.227.*X* (where *X* is 121, 123, 125, or 127) are authorized.

Table 50: *Additional examples for authorizing multiple stations*

	Entries for authorized manager list				Results
IP mask	255	255	0	255	This combination specifies an authorized IP address of 10.33.xxx.1. It could be applied, for example, to a subnetted network where each subnet is defined by the third octet and includes a management station defined by the value of "1" in the fourth octet of the station's IP address.
Authorized manager IP	10	33	248	1	
IP mask	255	238	255	250	Allows 230, 231, 246, and 247 in the 2nd octet, and 194, 195, 198, 199 in the 4th octet.
Authorized manager IP	10	247	100	195	

Operating notes

- **Network Security Precautions:**

  Enhance your network's security by keeping physical access to the switch restricted to authorized personnel, using the password features built into the switch, using the additional security features described in this manual, and preventing unauthorized access to data on your management stations.

- **Modem and Direct Console Access:**

  Configuring authorized IP managers does not protect against access to the switch through a modem or direct Console (RS-232) port connection.

- **Duplicate IP Addresses:**

  If the IP address configured in an authorized management station is also configured (or "spoofed") in another station, the other station can gain management access to the switch even though a duplicate IP address condition exists.

- **Web Proxy Servers:**

  If you use the WebAgent to access the switch from an authorized IP manager station, it is recommended that you avoid the use of a web proxy server in the path between the station and the switch. This is because switch access through a web proxy server requires that you first add the web proxy server to the Authorized manager IP list. **This reduces security by opening switch access to anyone who uses the web proxy server.** The following two options outline how to eliminate a web proxy server from the path between a station and the switch:

  ◦ Even if you need proxy server access enabled in order to use other applications, you can still eliminate proxy service for web access to the switch. To do so, add the IP address or DNS name of the switch to the non-proxy, or "Exceptions" list in the web browser interface you are using on the authorized station.

  ◦ If you don't need proxy server access at all on the authorized station, then just disable the proxy server feature in the station's web browser interface.

Overview

The switches covered in this guide provide support for advanced routing capabilities. Security is extremely important as complex networks and the internet grow and become a part of our daily life and business. This fact forces protocol developers to improve security mechanisms employed by their protocols, which in turn becomes an extra burden for system administrators who have to set up and maintain them. One solution to this is centralizing the mechanisms used to configure and maintain security information for all routing protocols. The Key Management System (KMS) can carry this burden.

KMS is designed to configure and maintain key chains. A key chain is a set of keys with a timing mechanism for activating and deactivating individual keys. KMS provides specific instances of routing protocols with one or more Send or Accept keys that must be active at the time of a request. A protocol instance is usually an interface on which the protocol is running.

Configuring key chain management

KMS has three configuration steps:

Procedure

1. Create a key chain entry.

2. Assign a time-independent key or set of time-dependent keys to the Key Chain entry. The choice of key type is based on the level of security required for the protocol to which the key entry will be assigned.

3. Assign the key chain to a KMS-enabled protocol.

This procedure is protocol-dependent. For information on a specific protocol, see the management and configuration guide for your switch.

Creating and deleting key chain entries

To use KMS, you must create one or more key chain entries. An entry can be the pointer to a single time-independent key or a chain of time-dependent keys.

NOTE	The key chain information is copied to the standby management module (if redundancy is enabled and the standby module has passed self-test).

Syntax:

```
[no] key-chain <chain_name>
```

Generate or delete a key chain entry. Using the optional `no` form of the command deletes the key chain. The `<chain_name >` parameter can include up to 32 characters.

```
show key-chain
```

Displays the current key chains on the switch and their overall status.

For example, to generate a new key chain entry:

Figure 154: *Adding a new key chain entry*

```
HP Switch(config)# key-chain Networking1        ◄──────  Add new key chain
HP Switch(config)# show key-chain                        Entry

 Key Chains                                              Display key chain
                                                         entries.
   Chain Name                       Keys Active Expired
   ------------------------------   ---- ------ -------
   Networking1                       0    0      0
```

After adding an entry, assign keys to it for use by a KMS-enabled protocol.

Assigning a time-independent key to a chain

A time-independent key has no Accept or Send time constraints. It is valid from boot-up until you change it. If you use a time-independent key, then it is the only key needed for a key chain entry.

Syntax:

```
[no] key-chain <chain_name> key <key_id>
```

Generates or deletes a key in the key chain entry <*chain_name* >. Using the optional `no` form of the command deletes the key. The <*key_id* > is any number from 0-255.

```
[key-string <key_str>]
```

This option lets you specify the key value for the protocol using the key. The <*key_str* > can be any string of up to 14 characters in length.

```
[accept-lifetime infinite][send-lifetime infinite]
```

`accept-lifetime infinite`: Allows packets with this key to be accepted at any time from boot-up until the key is removed.

`send-lifetime infinite`: Allows the switch to send this key as authorization, from boot-up until the key is removed.

```
show key-chain <chain_name>
```

Displays the detail information about the keys used in the key chain named <*chain_name* >.

Example:

To generate a new time-independent key for the switch key chain entry:

Figure 155: *Adding and displaying a time-independent key to a key chain entry*

```
HP Switch(config)# key-chain Networking1 key 1      ◄──────  Adds a new Time-Independent
HP Switch(config)# show key-chain Networking1                key to the "Networking1" chain.

 Chain - Networking1                                         Displays keys in the key chain
                                                             entry.
   Key | Accept Start GMT  Accept Stop GMT   Send Start GMT    Send Stop GMT
   --- + ----------------  ----------------  ----------------  -------------
    1  | Bootup            Infinite          Bootup            Infinite
```

Assigning time-dependent keys to a chain

A time-dependent key has Accept or Send time constraints. It is valid only during the times that are defined for the key . If a time-dependent key is used, there is usually more than one key in the key chain entry.

Syntax:

```
[no]key-chain <chain_name> key <key_id>
```

Generates or deletes a key in the key chain entry `<chain_name>` . Using the optional `no` form of the command deletes the key. The `<key_id>` is any number from 0-255.

```
[key-string <key_str>]
```

This option specifies the key value referenced by the protocol using the key. The `<key_str>` can be any string up to 14 characters in length.

```
[encrypted-key <key_str>]
```

Set key string using a base64-encoded aes-256 encrypted string.

```
accept-lifetime <mm/dd/yy[yy]hh:mm:ss|now>
```

Specifies the **start** date and time of the valid period in which the switch can use this key to authenticate inbound packets.

```
duration <mm/dd/yy[yy]hh:mm:ss|seconds>
```

Specifies the **time period** during which the switch can use this key to authenticate inbound packets. Duration is either an end date and time or the number of seconds to allow after the start date and time (which is the `accept-lifetime` setting).

```
send-lifetime <mm/dd/yy[yy]hh:mm:ss|now>
```

Specifies the **start** date and time of the valid period in which the switch can transmit this key as authentication for outbound packets.

```
show key-chain <chain_name>
```

Displays the detail information about the keys used in the key chain named `<chain_name >`.

NOTE	Using time-dependent keys requires that all the switches have accurate, synchronized time settings. You can manually set the time or use the Time protocol feature included in the switches. See time protocols in the management and configuration guide for your switch.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Example:

Figure 156: *Adding time-dependent keys to a key chain entry*

```
HP Switch(config)# key-chain Networking2 key 1 accept-1ifetime now 06/17/11
8:00:00
                                                    Adds a key with
                                                    time and date

HP Switch(config)# key-chain Networking2 key 2 accept-lifetime 06/18/11
8:00:00 duration 87000 send-lifetime 06/18/11 8:00:00 duration 86400
HP Switch(config)# key-chain Networking2 key 3 accept-lifetime 06/19/11
8:00:00 duration 87000 send-lifetime 06/19/11 8:00:00 duration 86400
HP Switch(config)# key-chain Networking2 key 4 accept-lifetime 06/20/11
8:00:00 duration 87000 send-lifetime 06/20/11 8:00:00 duration 86400
HP Switch(config)# key-chain Networking2 key 5 accept-lifetime 06/21/11
8:00:00 duration 87000 send-lifetime 06/21/11 8:00:00 duration 86400
                                                    Adds a key with
                                                    duration expressed
                                                    in seconds.
```

NOTE

Given transmission delays and the variations in the time value from switch to switch, it is advisable to include some flexibility in the Accept lifetime of the keys you configure. Otherwise, the switch may disregard some packets because either their key has expired while in transport or there are significant time variations between switches.

To list the result of the commands in **Figure 156: Adding time-dependent keys to a key chain entry** on page 613:

Figure 157: *Display of time-dependent keys in the key chain entry*

```
HP Switch(config)# show key-chain Networking2

 Chain - Networking2

  Key | Accept Start GMT  Accept Stop GMT   Send Start GMT    Send Stop GMT
  --- + ----------------  ----------------  ----------------  ----------------
   1  | 01/03/90 13:59:20 06/17/11 08:00:00 01/03/90 13:59:20 06/17/11 08:00:00
   2  | 06/18/11 08:00:00 06/19/11 08:10:00 06/18/11 08:00:00 06/19/11 08:00:00
   3  | 06/19/11 08:00:00 06/20/11 08:10:00 06/19/11 08:00:00 06/20/11 08:00:00
   4  | 06/20/11 08:00:00 06/21/11 08:10:00 06/20/11 08:00:00 06/21/11 08:00:00
   5  | 06/21/11 08:00:00 06/22/11 08:10:00 06/21/11 08:00:00 06/22/11 08:00:00
```

Use `show key-chain` to display the key status at the time the command is issued. Using the information from the example configuration in **Figure 156: Adding time-dependent keys to a key chain entry** on page 613 and **Figure 157: Display of time-dependent keys in the key chain entry** on page 613 if you execute `show key-chain` at 8:05 on 01/19/03, the display would appear as follows:

Figure 158: *Status of keys in key chain entry "Networking2"*

```
HP Switch(config)# show key-chain

 Key Chains

  Chain Name                       Keys Active Expired
  -------------------------------- ---- ------ -------
  Networking1                       1    0      1
  Networking2                       5    1      0
```

The "Networking1" key chain entry is a time-independent key and will not expire. "Networking2" uses time-dependent keys, which result in this data:

| Expired=1 | Key 1 has expired because its lifetime ended at 8:10 on 01/18/03, the previous day. |
| Active=2 | Key 2 and 3 are both active for 10 minutes from 8:00 to 8:10 on 1/19/03. |

Keys 4 and 5 are either not yet active or expired. The total number of keys is 5.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

Certificate Manager enables Public Key Infrastructure (PKI) capability on the switch providing authentication of network entities. This feature enables configuration and management of digital certificates on Networking switches, a key component of establishing digital identity in PKI.

Each entity in the PKI has their identity validated by a CA/RA. The CA issues a digital certificate as part of enrolling each entity into the PKI. This digital certificate is used by replying parties (e.g., network connection peers) to set up secure communication. Based on the information present in the certificate of the sender, the receiving entity can validate the authenticity of the sender and subsequently establish a secure communication channel.

Configuration support

The certificate manager CLI provides configuration support for integrating the switch into a customer's PKI.

Trust anchor profile (crypto pki ta-profile)

The profile defines required Anchor Trust for several certificate-specific operations, such as certificate enrollment and certificate validations. A trust anchor may be a Root CA certificate or an Intermediate CA certificate. The following command creates a trust anchor profile.

Syntax

```
(config) # [no] crypto pki ta-profile <profile-name> ssh-username <ssh-username>
```

Description

Create a Trust Anchor profile and associate it with an SSH username.

Options

profile-name

A name (maximum 100 characters) with a unique identifier for the Trust Anchor Profile. Ten TA profiles are supported: one for each allowed trust anchor (Root CA certificate.)

Profile number 2 is always reserved for self-signed certificate. For example, you can only create 9 TA profiles (Root CA certificates) per switch.

ssh-username

Set the username whose certificate will be validated with the TA profile for two-factor authentication.

Validation rules

Validation	Error/Warning/Prompt
If the maximum number of <username : TA profile> associations is reached for a given TA profile, a message displays.	Maximum number of username associations with a TA profile is 10.

Web User's Interface

When permitted by the existing configuration, the Web UI creates a "default" Trust Anchor profile (the profile name is "default") when a TA certificate is installed. The Web UI may only manage the TA certificate installed against the

---

"default" profile—no other certificates are visible or installed via Web UI. An administrator may create this same "default" TA profile. Restrictions on the "default" profile are described in Local Certificate Installation.

The Web UI manages a TA profile implicitly and only under the following conditions:

- If a TA Profile with the name "default" exists.

- If a TA Profile with the name "default" does not exist but one of the TA Profiles is not configured.

In these cases the Web UI may configure the "default" TA Profile.

When a default profile does not exist and both TA Profiles have been configured by the CLI (i.e., they both have a name that is not 'default'), the Web UI may not alter either TA profile and the usage web certificate to be installed must fit within a certificate chain belonging to an existing TA Profile.

Switch identity profile

The switch can have multiple certificates using the same base identity but with different protocol usage. This profile captures the common identity data for use in multiple certificates. The switch identity profile is a configuration aid that configures default values used when creating multiple certificates. This profile is not used for any other purpose and is therefore optional. The user can enter both subject information and one or more IP addresses when creating an Identity Profile. There is no constraint to have either subject or IP addresses, they are not mutually exclusive although at least one must be present.

Syntax:

```
(config)# [no] crypto pki [identity-profile][profile-name] subject[CommonName <cn-
value>] [Org <org-value> ] [OrgUnit <org-unit value>] [Locality <location-value>]
[State <state-value>] [Country <country-code>]
```

Subject fields

If the certificate `subject` fields are not provided on the command line, they are interactively prompted as follows:

identity-profile

  Creates an identity profile.

profile-name

  Specify the Switch Id Profile name.

cn-value

  Common Name (CN) – must be present, max length 90.

org-value

  Organization Name (O) – preferred, max length 100.

org-unit value

  Organizational Unit Name (OU) – preferred, max length 100.

location-value

  Locality (L) – optional, max length 100.

state-value

  State (ST) – optional, max length 100.

country-code

  To specify the two letter ISO 3166-1 country code. Max length 2.

Local certificate enrollment – manual mode

To enroll a Local certificate using a manual copy and paste method, a key size and the relevant TA profile details are required. The following command manually creates a certificate signature request. Including the subject overrides the configured Identity Profile:

Syntax:

```
crypto pki create-csr certificate-name <CERT-NAME>
  ta-profile <Profile-name>
  [usage <openflow | web | all>]
  [key-type rsa key-size <1024|2048>]
  [key-type ecdsa curve <256|384>]
  subject [common-name <CN-Value>]
  [org <Org-Value>] [org-unit <Org-unit-value>]
  [locality <Location-Value>] [state <state-Value>]
  [country <Country-Code>]
  [valid-start <date>] [valid-end <date>]
```

Options

key-size [1024|2048]

The length of the RSA key, default is 1024 bits.

Definitions:

certificate-name

Name of the certificate.

ta-profile

The Trust Anchor Profile associated with the certificate. A profile named 'default' is updateable from the web UI.

ta-profile-name

Specify the Switch Id TA profile name.

usage <openflow|web|all>

Intended application for the certificate, the default is web.

valid-start

Certificate validity start date (MM/DD/YY).

valid-end

Certificate validity end date (MM/DD/YY).

Subject fields

cn-value

Common Name (CN) – must be present, max length 90.

org-value

Organization Name (O) – preferred, max length 100.

org-unit value

Organizational Unit Name (OU) – preferred, max length 100.

location-value

Locality (L) – optional, max length 100.

state-value

State (ST) – optional, max length 100.

country-code

To specify the two letter ISO 3166-1 country code. Max length 2.

NOTE A CSR created with TA profile name of 'default' MUST include usage of either "web" or "all".

Example: of PEM format output

This command creates a certificate signing request in realtime and then output the result to the console:

```
-----BEGIN CERTIFICATE REQUEST-----
MIIBpDCCAQ0CAQAwZDELMAkGA1UEBhMCVVMxCzAJBgNVBAgTAkNBMRIwEAYDVQQH
EwlSb3NldmlsbGUxCzAJBgNVBAoTAkhQMQ0wCwYDVQQLEwRFVlBHMRgwFgYDVQQD
Ew9UZXN0IE1hY2hpbmUgMDEwgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJAoGBAN7i
w3x2gi3tZf4LnXltSicl7RNcVggxYHcZQySWFtCXFTb5uaJ6vA3RdBIThgUKZSpc
rgtc7jQmRDUdKAbWLPrqC7wBxMlXbnQYegubvOfzf/dT1CYJXxdUZh5BMN5ob/00
t60m9cM7Odsu0a0dBoQQRI8315KJ0AuHDE6VOe4dAgMBAAGgADANBgkqhkiG9w0B
AQUFAAOBgQBQCZar2ox6RXm7F/vVhyrrp0E0YrPimxDvg40jnwqtwOgpQAvns4pt
o5RVx4/Q6hzF2QivYqLl3+K8WOVVJ7XLDcHNea8RJgx13t45uMYrsMKWdbhR9+jQ
KFzmffQJXRXOnH6rfQSNYBXndg0azhc8saORrOqrTn3Yw3psYSNMbA==
-----END CERTIFICATE REQUEST-----
```

You must manually copy the certificate signing request (CSR) created with the "create-csr" command (above) and have it signed by a CA. The local certificate status is updated to "CSR" after the CSR is created. A pending certificate request is not persistent across a power cycle or reboot. Once the CA-signed certificate response is received, the user executes the following command and pastes the signed certificate provided by CA on the command line.

The switch retains the name of the certificate used when creating the CSR in memory while waiting for the signed certificate to be installed. When the signed certificate is pasted to the command line, the switch matches the certificate to the CSR by matching the public key and then saves the signed certificate to flash. The signed certificate will not be accepted if a CSR does not exist or if the trust chain cannot be verified (for example, if the CA's root certificate is not installed in the Trust Anchor Profile.)

Syntax:

```
(config)# crypto pki install-signed-certificate <data>
```

Intermediate certificate installation is similar to the local certificate installation. When intermediate certificates are to be individually installed, the local-certificate name is used and certificate manager uses this name to build the certificate chain between the root and the leaf certificate of the specified name. Intermediate certificates must be presented in order from the trust anchor to the local (leaf) certificate. The user is prompted to paste the new certificate (PEM-encoded PKCS#7) to the command line. The provided data is parsed internally by Certificate Manager and stored in DER format thus requiring no additional parsing in CLI. The following text appears.

NOTE To install a signed certificate, the certificate must match a previously created signing request.

With the cursor at the start of a blank line, when the user presses the Enter key, the user operation is done. Usage of word pad is suggested to copy the certificate and paste it to this command.

To check the CSR status, enter:

```
show crypto pki local-certificate
```

Local enrollment is implemented in the web UI; specifically the security — SSL page is updated for the Web UI SSL server application, with web usage. The Web UI does not provide general PKI configurability for all applications (Web UI does not allow creation or management of other device certificates add.)

	Self-signed certificate for a specific application (along with the key-pair) is removed once a CA signed local-certificate is installed for that application.

Self-signed certificate enrollment

This certificate installation method may be used when a Certificate Authority is not available. A self-signed certificate provides the relying party no assurance of identity, so this is not as secure as using a CA-signed certificate. A self-signed certificate may be useful, but its use is not recommended.

A self-signed certificate many only be installed on the "default" TA-Profile, so the ta-profile-name parameter is not present in the command.

To enroll a local certificate in self-signed mode, the user must specify the subject information and key-size. The details specific to the certificate "subject" are obtained from id-profile if not specified here.

Syntax:

```
[no] crypto pki enroll-self-signed certificate-name <CERT-NAME> subject [common-
name <CN-Value>] [org <Org-Value>][org-unit <Org-unit-value>] [locality <Location-
Value>] [state <state-Value>][country <Country-Code>][valid-start <date>][valid-
end <date>] [usage <openflow | web | all>] [key-type rsa key-size <1024|2048>]
[key-type ecdsa curve <256|384>]
```

Options

key-size [1024|2048]

The length of the key; default is 1024 bits.

usage [<openflow|web|all>]

Intended application for the certificate; the default is web. The openflow option is not supported for self-signed certificate enrollment.

Subject Fields

The following prompts appear if these required fields are not given as arguments.

```
Enter  Common Name(CN) :
Enter  Org Unit(OU) :
Enter  Org Name(O) :
Enter  Locality(L) :
Enter  State(ST) :
Enter  Country(C) :
```

Self-signed certificate

A self-signed certificate uses the "default" TA profile, which is created automatically if it does not already exist and one of the ten available TA Profiles is not yet assigned.

Syntax:

```
[no] crypto pki create-self-signed certificate-name [name] subject [common-name cn-
```

value] [org *org-value*] [org-unit *org-unit-value*] [locality *location-value*] [state *state-value*] [country *country-code*]]

To create and install a self-signed local certificate the certificate subject may be configured with the `crypto pki identity-profile` command.

Options

key-size [1024|2048]

The length of the key; default is 1024 bits.

subject [field <field value>]

Subject fields of the certificate; the default values are specified in the identity profile.

`usage [<openflow|web|all>]`

Intended application for the certificate; the default is `web`.

valid-start *date*

Start date of the certificate.

valid-end *date*

End date of the certificate.

Subject Fields

Following are the prompts appear if these required fields are not given as arguments.

```
Enter Common Name(CN) :
Enter Org Unit(OU) :
Enter Org Name(O) :
Enter Locality(L) :
Enter State(ST) :
Enter Country(C) :
```

Definitions:

certificate-name

Name of the certificate.

ta-profile

The Trust Anchor Profile associated with the certificate. A profile named 'default' is updateable from the web UI.

ta-profile-name

Specify the Switch Id TA profile name.

cn-value

Common Name (CN) – must be present, max length 90.

org-value

Organization Name (O) – preferred, max length 100.

org-unit value

Organizational Unit Name (OU) – preferred, max length 100.

location-value

Locality (L) – optional, max length 100.

state-value

> State (ST) – optional, max length 100.

country-code

> To specify the two letter ISO 3166-1 country code. Max length 2.

valid-start

> Certificate validity start date (MM/DD/YYYY).

valid-end

> Certificate validity end date (MM/DD/YYYY).

The default value for start date is the current date and the default value for the end date is the current date plus one year.

Local enrollment is implemented in the web UI and the security — SSL page is updated for the web UI SSL server application. The Web UI does not provide general PKI configurability for all applications creation or management of other device certificates.

Removal of certificates/CSRs

To remove the certificates/CSRs, use the following command:

Syntax:

```
(config)# crypto pki clear certificate-name [Cert-Name]
```

Clears the CSR or certificate and its related private key.

Definitions:

certificate-name

> Name of the local certificate.

Zeroization

Certificate and key removal is discussed as part of the [no] form of each certificate installation command above. The [no] forms described above delete certificates and keys. The "Zeroize" command simply deletes (unlinks) key files. Full file system zeroization is performed by following with FIPS/Secure Mode commands.

The `[no]` form is supported only for TA profile and identity profile. It is not supported for local certificate. Zeroization erases keys and related PKI data such as CSRs and TA profiles from the file system.

Syntax:

```
crypto pki zeroize
```

This command returns crypto pki configuration to the factory default state by deleting all certificates and related private keys. The Trust Anchor profile and switch identity profile configurations are also removed.

`zeroize`

> Removes all pki configuration, including profiles, certificates and keys.

> **NOTE**
>
> The `[no]` form is not available for the certificate command. To remove a certificate from the switch, use the `clear` command.

File transfer

To load a Trust Anchor Certificate against a TA Profile, execute the following command.

NOTE   The TA profile must exist for the command to succeed.

Syntax:

```
(config#)copy tftp ta-certificate ta-profile-name <ip-addr/ipv6-addr> filename
```

or

```
copy sftp ta-certificate ta-profile-name ip-addr/ipv6-addr|host-name-str user <user-name>|username@ip-strport <TCP-port> FILE-NAME
```

Definitions:

ta-certificate

Copy a Trust Anchor certificate to the device.

ta-profile-name

The Trust Anchor Profile associated with the certificate.

ip-addr

IP address of the server.

file-name

Name of the certificate file.

ipv6-addr

Specify TFTP server IPv6 address.

host-name-str

Specify hostname of the SFTP server.

user

Specify the username on the remote system.

username@ip-str

Specify the username along with remote system information (hostname, IPv4 or IPv6 address.)

port

TCP port of the SSH server on the remote system.

Syntax:

Copy a Trust Anchor (TA) certificate to the device using TFTP:

```
copy tftp local-certificate [Cert-Name] [<ip-addr/ipv6-addr>]
```

Copy a Trust Anchor (TA) certificate to the device using SFTP:

```
copy sftp local-certificate [Cert-Name] [<ip-addr/ipv6-addr/host-name-str>] [user <user-name>] [username@ip-str <filename>]
```

The file is checked immediately upon completion of transfer and results written to the CLI. The file can be in PEM-encoded or DER-encoded (binary) PKCS#7 format. If the certificate subject matches an existing TA certificate associated with the specified TA profile, then the new certificate updates the existing certificate.

Any certificate which is a root or intermediate certificate will be accepted as a TA certificate. There is no check for the subject.

Loading a local certificate

To load a local certificate (single certificate/certificate chain), execute the following command.

Syntax:

```
(Switch_Name#)copy tftp local-certificate <ip-addr> <file-name>
```

```
(Switch_Name#)copy sftp local-certificate [user <user-name>] [<ip-addr/ipv6-addr/
host-name-str>] [<username@ip-str>] <filename> [port <1-65535>]
```

Definitions:

ta-certificate

 Copy a Trust Anchor certificate to the device.

ta-profile-name

 The Trust Anchor Profile associated with the certificate.

local certificate

 Local Certificate to be copied.

ip-addr

 IP address of the server.

file-name

 Name of the certificate file.

ipv6-addr

 Specify TFTP server IPv6 address.

host-name-str

 Specify hostname of the SFTP server.

user

 Specify the username on the remote system.

username@ip-str

 Specify the username along with remote system information (hostname, IPv4 or IPv6 address.)

port

 TCP port of the SSH server on the remote system.

| NOTE | The loaded certificate is validated against a pending Certificate Request and the maximum number of intermediate certificates allowed is 3 (with a maximum of 5 certificates in a chain). Intermediate certificates must either be loaded before the local certificate or are included as part of the local certificate. |

Debug logging

The following command enables/disables debug logs.

Syntax:

```
(config)# [no] debug security <crypto>
```

Definitions:

crypto

Display all Crypto messages.

Certificate specific

This command displays summary information for all certificates or detailed information for the named certificate.

```
show crypto pki local-certificate [summary | <Cert-Name>]
```

Example of displaying certificate summary information:

```
switch# show crypto pki local-certificate summary
  Name                    Usage         Expiration      Parent / Profile
  --------------------    ------------  --------------  ----------------
  IDEVID_CERT             IDEVID        2031/01/26      IDEVID_INTER_1
  IDEVID_INTER_1          IDEVID        2031/01/26      IDEVID_INTER_2
  IDEVID_INTER_2          IDEVID        2031/01/26      IDEVID_ROOT
```

Summary mode lists all certificates below a TA profile, including both local certificates and installed intermediates. The names of intermediate certificates are transitory and can change after local certificates are added or removed. In detailed mode the "certificate name" can be provided as an argument and details specific to the certificate are displayed. If the "expiration" displays CSR, then detailed mode re-displays the CSR as described with the `crypto pki create-csr local-certificate` commands.

All installed certificates are shown in the same way, provided that the fields exist in the certificate. For example, a CA signed certificate has an "Issuer:" field with a different value from the "Subject" field. In a self-signed certificate, these fields are set to the same value. Since the fields are present in either type of certificate, they are always shown. Similarly, a Root certificate is a self-signed certificate. A trust anchor certificate can be either a Root certificate or an Intermediate certificate. The same fields are present in the certificate—just set to different values.

When working in the summary mode:

• An installed certificate can or can not have a subject key identifier.

• An installed certificate can or can not contain an authority key identifier.

• An installed certificate can or can not contain key usage constraints, which can or can not be marked critical.

• When an extension is critical, the keyword "critical" is displayed; when the extension is not critical, no additional wording is displayed.

While address ranges can be encoded in a certificate, this usage is not consistent with identifying a switch (or switch interface), so CIDR format is not expected. However, if present it must be displayed for diagnostic purposes. (CIDR format display can be eliminated by adding tests to reject certificates with a range at the time of certificate installation.) IP addresses are listed in lexicographical order, except that all IPv4 addresses are shown as a group before IPv6 addresses are displayed. IPv6 addresses are shown in full, without the "zeroes removed" notation.

> **NOTE**
> Per RFC-5280: "Certificate users MUST be able to handle serial Number values up to 20 octets." Thus, the serial number can take 40 hex characters to print. The serial number is printed in hex to limit string length and to allow easier manual decoding of UUID type serial numbers.

The detail form of the certificate specific `show` command is available from the web UI. The web UI allows display of those configured certificates related to the web server only. This includes the SSL server certificate, trust anchor certificate and any other certificates configured as part of the certificate chain. All the certificates in the trust chain are also displayed.

Profile specific—TA profile

Two forms of output are available for this command, summary and detailed. If no argument is provided, a brief about all profiles is printed as shown below.

Show profile specific

Show Trust Anchor profile specific details.

Syntax:

```
(Switch_Name#)show crypto pki ta-profile [ta-profile-name]
```

ta-certificate

    Copy a Trust Anchor certificate to the device.

ta-profile-name

    Trust Anchor Profile for the certificate.

Example:

```
Profile Name         Status
-------------------- -------------- -----------------------------
HP Procurve          Root Certificate Installed
Microsoft_Inc        Pending Root Certificate Installation
```

NOTE: This command is not available on the web UI.

Example:

If a TA Profile name is specified as an argument, the Trust Anchor details for the specified TA Profile name is displayed.

```
Profile Name         Status
-------------------- ------------------------------------------------
HP Procurve          3 certificates installed

Trust Anchor: <print_cert for Trust Anchor>
```

The output format for the TA certificate is same as the format for "Certificate details" above. The "Status" field lists the total number of certificates, including intermediates and local, that references this trust anchor. Intermediate certificates are shown with local certificates, as certificates under an anchor form a tree not a list.

NOTE: This command is not available on the web UI.

Certificate details

Show the configured switch identity.

Syntax:

```
(Switch_Name#) show crypto pki <identity-profile>
```

Definitions:

identity-profile

   Displays the configured Id profile related information.

Example:

```
Sample output:
Switch Identity:
Common Name (CN) : <max 90 chars>
Org Unit (OU)    : <max 100 chars>
Org Name (O)     : <max 100 chars>
Locality (L)     : <max 100 chars>
State (ST)       : <max 100 chars>
Country (C)      : <max 2 chars>
```

> **NOTE**
>
> Blank fields display as empty, except for `Common Name`. Since `Common Name` is a required field, it displays as `not configured` when empty.
>
> This command is not available with the web UI. If configured, the ID profile is displayed in the Web UI.

Display PKI certificate

Syntax:

```
display pki certificate <ca|local>
```

Definitions:

pki

   Display PKI information.

certificate

   Display certificate configuration information.

ca

   CA certificate.

local

   Local certificate.

Example:

```
<Sysname> display pki certificate local
Certificate:
  Data:
    Version: 3 (0x2)
    Serial Number: 10B7D4E3 00010000 0086...
  Issuer:
    emailAddress=myca@aabbcc.net
    C=CN
    ST=Country A
    L=City X
    O=abc
    OU=bjs
```

```
        CN=new-ca
Validity
    Not Before:  Jan 13 08:57:21 2004 GMT
    Not After :  Jan 20 09:07:21 2005 GMT
Subject:
    C=CN
    ST=Country B
    L=City Y
    CN=PKI test
Subject Public Key Info:
    Public Key Algorithm:  rsaEncryption
    RSA Public Key:  (512 bit)
    Modulus (512 bit):
     00D41D1F ...
     Exponent:  65537 (0x10001)
    X509v3 extensions:
     X509v3 Subject Alternative Name:
     DNS: hyf.xxyyzz.net
    X509v3 CRL Distribution Points: URI:http://1.1.1.1:447/myca.crl
...          ...
Signature Algorithm: md5WithRSAEncryption
A3A5A447 4D08387D ...
```

All data needed to display the certificate as shown above comes from the certificate so is available. Note that the X509 extension display values will change depending on what extensions are present in the certificate. When no extensions are present, the "X509v3 extensions:" line is not present.

> The X509 extension display values change depending on what extensions are present in the certificate. When no extensions are present, the "X509v3 extensions:" line is not present.

Web support

The current security—SSL page configures web UI SSL servers only. The Suite B features are not supported on the web UI. The following are requirements for a web UI design:

- The web UI implicitly uses a TA profile named "default". If the TA certificate installed on the switch is associated with a profile of another name, the TA certificate is read-only to the web UI. See **Trust anchor profile (crypto pki ta-profile)** on page 615.

- The web UI supports local certificate enrollment with an implicit usage of 'web'. See **Local certificate enrollment – manual mode** on page 617.

- The web UI supports self-signed local certificate enrollment with an implicit usage of 'web'. See **Local certificate enrollment – manual mode** on page 617.

- The web UI shows the TA certificate and the configured SSL server certificate with 'web' usage with any intermediate certificates in the chain. The display will match the Certificate Detail format as described in **Show profile specific** on page 625.

- The web UI must be able to replace an SSL server certificate (as it currently does.)

- The web UI does not need to provide 'zeroization' of any certificates. See **Zeroization** on page 621.

SSL screen

The following figure shows the results if the user clicks Security — SSL. Details about each panel/widow in the SSL page are given below.

Figure 159: *SSL Screen*

Panel hierarchy

The SSL panel displays Certificate Management features.

TA certificates panel

The **Trust Anchor (TA) Certificates** Panel displays information and status for TA profiles. Buttons, Install and Remove, install new TA profiles or remove existing ones.

Aruba 2930F / 2930M Access Security Guide for ArubaOS-Switch 16.05

To install a new TA certificate, click **Install**. The install screen appears and prompts for certificate location. Click the **Upload** button to upload the new TA certificate to the switch. Click **Cancel** to abort the installation.

Figure 160: *Install TA profile*

Trust Anchor (TA) Certificates			Upload	Cancel	? ▲
TA Profile:	default				
Certificate File:	c:\root.cert	Browse			

A **default** TA profile is automatically created when the conditions explained in section have been satisfied.

The **install** option is not available if:

Procedure

1. All ten TA profiles are used and none are named "default". The TA profile number 2 is always reserved for self-signed certificate.

2. The current certificate with 'usage=web' is linked to a TA profile whose name is not "default"

Switch identity profile panel

Switch Identity Profile displays the details of switch identity profile, if already configured with the CLI. Otherwise displays `Switch Identity Profile is not configured`.

Installed certificates panel

The **Installed Certificates** panel displays the certificate profile, usage, key size, status, type, beginning and end date for currently installed certificates.

Figure 161: *Installed certificates*

Installed Certificates						Remove	
Certificate Name	**TA Profile**	**Usage**	**Key Size**	**Status**	**Start Date**	**End Date**	**Type**
locCertSS1	default	web	1024	Installed	########	1971-01-01	CA Signed
View Certificate	TA Profile Cert						

View Certificate displays all certificates in the certificate chain. The view certificate list displays the local certificate, up to three intermediate certificates and one TA certificate.

When a certificate is selected, a detailed view of the certificate is displayed in a popup window.

Figure 162: *TA certificate*

```
Model J1234A                                                              ?      X

Serial Number:    75A5A501ABCDEF12345675A5A501ABCDEF123456
Sig. Algorithm:   SHA1 with RSA encryption
Issuer:       CN=HP Networking Platform Certificate Authority 01, OU=HP Networking, O=Hewlett-Packard
Company, L=Roseville, ST=California, C=US

Validity From: Mar 11 23:56:35 2010 GMT
Validity To: Mar 8 23:56:38 2030 GMT

Subject: CN=Model J1234A/serialNumber=SW123456780A, BaseMAC 010203-040506, OU=HP Networking
EVDC  O=Hewlett Packard Company
```

Certificate requests panel

The Certificate Requests panel displays the status of currently requested certificate.

Within the panel the **Create Self-Signed Certificate** link is available.

Procedure

1. Ensure that there are no pending requests.

2. Create and install a self-signed certificate.

Any existing certificate will be replaced with one of the same name. A non-default TA Profile with a certificate configured with usage of **web** will not be allowed.

Create Self-Signed Certificate

Creates a self-signed certificate. Upon selection of this link, an edit request form becomes available which provides all required information for the creation of the certificate.

> **NOTE**
>
> The default TA Profile is called **Default**.

Figure 163: *Certificate requests form*

The **Certificate Request** field have the following constraints:

Common Name (CN) – must be present, max length 90. Common Name should be preset with value from Switch ID profile if one exists.

Organizational Unit Name (OU) – preferred, max length 100.

Organization Name (O) – preferred, max length 100.

Locality (L) – optional, max length 100.

State (ST) – optional, max length 100.

Country (C) – preferred, max length 2.

Start Date — Preset with current date.

End Date — Preset with current date + 1 year.

Select **Install** when the form has been completed to Install this certificate to the switch.

Select **Cancel** to cancel the user request.

Error messages

Table 51: *Error messages*

Error Message	Explanation
The TA profile %s does not exist.	Fail the revocation-check command when a given ta-profile is not already configured on the switch. And Fail the clear crl command if an invalid ta-profile is given.
The TA profile %s has no certificate configured.	Fail the revocation-check command if TA certificate associated to the given profile is not already downloaded.
The URL length exceeds the maximum allowed length of 255 characters.	Restrict length of the URL revocation URL to 255 characters.

Table Continued

Error Message	Explanation
`Either the TA certificate is not installed or the revocation check is not set to CRL.`	A warning message will be displayed if any of the following is not completed when this command is run. 1. If revocation check is not CRL. 2. If TA certificate is not installed.
`Do you want to delete all the CRLs? Continue (y/n):`	A prompt is given to user to select y/n before deleting all CRLs.
`Do you want to delete the CRL of the TA profile %s? Continue (y/n):`	A prompt is given to user to select y/n before deleting CRL of the TA profile.
`The SuiteB-minLoS command in strict mode is mutually exclusive, with minimum TLS configured for an application.`	Following is the command which needs to be mutually exclusive with the SuiteB-minLOS command if its configured in strict mode for SSL: `[no] tls application {web-ssl | openflow | syslog | tr69 | cloud | all} lowest-version {tls1.0 | tls 1.1| tls 1.2 | default } [cipher {aes256-sha256 | aes256-sha | aes128-sha256 | aes128-sha | des3-cbc-sha | ecdh-rsa-aes128-gcm-sha256}]`
`Do you want to terminate the existing SSL/TLS sessions? Continue (y/n):`	When minLOS is configured for TLS, prompt the user to kill the existing SSL/TLS sessions.
`All manager level %s public keys will be deleted. Continue (y/n):`	A warning message when user tries to clear all the manager keys with key type (rsa/dsa) specified.
`All operator level %s public keys will be deleted. Continue (y/n):`	A warning message when user tries to clear all the operator keys with key type (rsa/dsa) specified.
`All operator level %s public keys will be deleted. Continue (y/n):`	A warning message when user tries to clear all the keys with key type (rsa/dsa) specified.
`The CRL is not downloaded or not available in the switch.`	An error message when CRL is not available in the switch and when user checks the revocation status of a certificate using the command `show crypto pki ta-profile ta1 crl certificate-serial-num 0x3535456889ccce2e.`
`A DNS server must be configured before configuring the named URL for CRL/OCSP.`	An error message when a named URL is configured without DNS is on switch configured.
`A CRL URL must begin with ldap://.`	The URL scheme should 'ldap' for CRL.
`An OCSP URL must begin with http:// or https://.`	The URL scheme should be 'http' or 'https' for OCSP anything else should be given an error.

Table Continued

Error Message	Explanation
A self-signed certificate requires a TA profile. One of the existing TA profiles must be removed before a self-signed certificate may be created.	A default profile does not already exist and the ten TA profiles are already in use (preventing creation of the 'default' TA profile.) The maximum number of TA profiles is ten; and one of those ten TA profiles is reserved for self-signed (the default profile).
No matching CSR found. Certificate validation failed.	If there is no CSR present on the switch, then the certificate cannot be matched.
Certificate "<cert name>" already exists".	The specified certificate name is already used.
Configuration failed. Incomplete certificate chain.	The certificate chain is incomplete.
The certificate has expired or is not yet valid.	The certificate is invalid.
Configuration failed. The specified key is not available. Please wait and try again.	The crypto keys are not available.
Certificate name is too long. The maximum length is 20.	The certificate name exceeds the maximum length allowed.
Certificate subject does not match the existing certificate associated with Trust Anchor profile <TA-Profile-Name>.	Overwriting a configured TA certificate.
File format not recognized or file is corrupted. Certificate validation failed.	There is a problem with the file, such a corruption.
Key generation in progress, try again later.	The key is not generated yet during a request for a CSR.
No matching CSR found. Certificate validation failed.	There is no CSR present on the switch, the certificate cannot be matched.
Profile was not added. The maximum number of profiles is %d	User tried to add the tenth TA profile.
Removing this TA profile will also remove all associated certificates. Continue (y/n) ?	Removing a TA profile removes all associated certificates. The following warning/query appears.
TA Profile%s does not exist.	During deletion, the mentioned TA profile is not existing / not configured.
The certificate cannot be verified because the associated Trust Anchor profile has no certificate configured.	The Trust Anchor certificate is not installed on the switch against the specified TA profile.

Table Continued

Error Message	Explanation
`The default value for start date is the current date and the default value for the end date is the current date plus one year.`	When enrolled for self-signing without having a switch identity profile or subject fields in the command line.
`The existing certificate for this TA profile [%s] will be replaced. Continue (y/n)?`	The mentioned TA profile certificate already exists and the user attempts to install a new certificate over the existing one.

`Enter Common Name(CN) :`
`Enter Org Unit(OU) :`
`Enter Org Name(O) :`
`Enter Locality(L) :`
`Enter State(ST) :`
`Enter Country(C) :`
	Prompts appear if the required fields are not given as arguments.

Suite B is a set of cryptographic algorithms used for encryption, key exchange, digital signature, and hashing. As per RFC 6460, the Fact Sheet on Suite B Cryptography requires key establishment and authentication algorithms based on Elliptic Curve Cryptography and encryption using AES.

In particular, Suite B includes the following:

- Advanced Encryption Standard (AES) – FIPS 197 (with key sizes of 128 and 256 bits)

- Elliptic Curve Digital Signature Algorithm (ECDSA) using 256 and 384 bit prime module curves – digital signatures

- Elliptic Curve Diffie-Hellman (ECDH) using 256 and 384 bit prime module curves – key exchange

- Secure Hash Algorithm 2 (SHA-256 and SHA-384) – message digest

- Additional PKI / Certificate management requirements: Certificate Revocation List (CRL) and Online Certificate Status Protocol (OCSP)

Suite B algorithms are defined to support two minimum levels of security, `minLoS`, with security strengths of 128 and 192 bits:

- `minLOS-128`

- `minLOS-192`

The level of security is determined by the strength of the keys.

Configuration support

CRL configuration facts

- When a certificate is presented while a CRL download is in progress and that the cached CRL has become stale or is not present, the acceptation or rejection of the certificate is subject to the policy enforcement of CRL configuration.

- When a CRL becomes stale, for example if the current time is ahead of the `nextUpdateTime` of the CRL, the CRL is deleted immediately.

- Once a successful TLS connection is established, even if the server certificate is revoked at a later time, the connection continues to exist until a renegotiation happens.

- If a CRL download fails due to any reason (for example, the server is not reachable or the memory is not available), an event is recorded in the system log with the failure reason. Once you have resolved the failure issue, you must initiate a download.

- You can download only one CRL at a time. If you initiate a request to fetch a CRL while a CRL download is already in progress, your request will be rejected.

- The Cumulative Maximum storage allowed for CRLs in flash is 1 MB.

- Only two CRL files are allowed in the system. Any fetch request beyond this limit is rejected and logged appropriately.

- CRL fetch is supported only via LDAP. The CRL downloaded is of DER (binary) format.

- If you delete an installed root-certificate when a CRL download for that profile is already in progress, the download will be uninterrupted. The downloaded CRL thereafter will be deleted once its lifetime expires (becomes stale).

- When you configure a CRL URL for a given TA profile, it takes priority over the CDP server settings mentioned in the certificate.

- You can configure two URLs per CRL/CDP LDAP servers and OCSP responders.

- Standard TCP timeouts are applicable during CRL fetch or OCSP status fetch.

- CRLs are also written into the non-volatile memory so that when a device reboots or failover and previously had a valid CRL, it will automatically be loaded from the non-volatile memory avoiding a re-fetch of the CRL. In addition, for every 24 hour period (per CRL file), a given CRL file is updated into the flash memory if there is any recent update to the last written state.

OCSP configuration facts

- If you delete an installed root-certificate at the same time that an OCSP handshake is in progress, the revocation status o/p will be based on the deleted root-certificate.

- If you configure an OCSP responder URL for a given TA profile, it takes priority over the OCSP server settings specified in the AIA field of the client certificate.

- You can configure at most two OCSP responder URLs.

- If the revocation-check is configured as both OCSP and CRL, OCSP takes precedence. For example, the switch tries to retrieve the revocation status using OCSP first followed by CRL.

Configure CRL for revocation check

Configures the parameters for the Certificate Revocation list (CRL) revocation check mode.

Syntax

```
crypto pki ta-profile <profile-name>revocation-check [crl] [[strict|optional]
[url1 <REVOC-URL> | url2 <REVOC-URL>|[refresh-interval <hours>]
```

Options

profile-name

A name (maximum 100 characters) with a unique identifier for the Trust Anchor Profile. Ten TA profiles are supported: one for each allowed trust anchor (Root CA certificate.)

revocation-check

Applies revocation check on a TA profile.

crl

Uses CRL for revocation.

Options

You can only specify one of these options.

strict

Sets the enforcement as strict.

optional

Sets enforcement as optional.

url1

Configures the first URL.

url2

Configures the second URL.

refresh-interval

Sets the periodic update interval in hours, default is 24.

Configure OCSP for revocation check

Configures the parameters for the OCSP revocation check mode.

Syntax

```
crypto pki ta-profile profile-name revocation-check ocsp [[strict|optional] | [url1 REVOC-URL] | [url2 REVOC-URL] | [disable-nonce]]
```

Definitions

profile-name

A name (maximum 100 characters) with a unique identifier for the Trust Anchor Profile. Ten TA profiles are supported: one for each allowed trust anchor (Root CA certificate.)

revocation-check

Applies revocation check on a TA profile.

ocsp

Uses OCSP for revocation.

Options

You can only specify one of these options.

strict

Sets the enforcement as strict.

optional

Sets enforcement as optional.

url1

Configure the first URL.

url2

Configures the second URL.

disable-nonce

Disables the nonce.

Retrieve CRL

Retrieves the CRL of the TA profile.

Syntax

```
crypto pki ta-profile profile-name retrieve-crl
```

Definitions

profile-name

> A name (maximum 100 characters) with a unique identifier for the Trust Anchor Profile. Ten TA profiles are supported: one for each allowed trust anchor (Root CA certificate.)

retrieve-crl

> Retrieves the CRL of the TA profile. You must configure the CRL URLs before you can perform this command. See **Configure CRL for revocation check** .

Set TA profile to validate CRL and OCSP

Sets the TA profile that contains root certificate for validating the CRL file.

Syntax

```
crypto pki ta-profile profile-name crl-root-profile ta-profile-name
```

Definitions

ta-profile-name

> Name of the TA profile that contains root-certificate to validate revocation response.

crl-root-profile

> Sets the TA profile that contains root certificate for validating the CRL file.

Sets the TA profile that contains root certificate for validating the OCSP response.

Syntax

```
crypto pki ta-profile profile-name ocsp-root-profile ta-profile-name
```

Definitions

ta-profile-name

> Name of the TA profile that contains root-certificate to validate revocation response.

ocsp-root-profile

> Sets the TA profile that contains root certificate for validating the OCSP response.

Clear CRL

Clears the CRL associated with the TA profiles.

Syntax

```
crypto pki clear crl [all | ta-profile profile-name]
```

Definitions

crl

> Clears all the CRLs associated with the TA profiles.

all

> Clears all the CRLs associated with the all the TA profiles.

ta-profile

Clears the CRL of the specified TA profile.

profile-name

A name (maximum 100 characters) with a unique identifier for the Trust Anchor Profile. Ten TA profiles are supported: one for each allowed trust anchor (Root CA certificate.)

Create a certificate signing request

Creates a certificate signing request on the switch. Including the subject will override the configured identify profile.

Syntax

```
crypto pki create-csr certificate-name CERT-NAME ta-profile Profile-Name [usage
<openflow | web | all>][key-type rsa key-size <1024|2048>] [key-type ecdsa curve
<256|384>] [subject [command-name CN-Value] [org Org-Value] [org-unit Org-unit-
value] [locality Location-Value] [state state-Value] [countryCountry-Code]] [valid-
start date valid-end date]
```

Definitions

profile-name

A name (maximum 100 characters) with a unique identifier for the Trust Anchor Profile. Ten TA profiles are supported: one for each allowed trust anchor (Root CA certificate.)

usage

When `usage` is set to `all`, it includes the OpenFlow and web applications, as well as other applications such syslog.

Options

rsa

Uses the RSA key. You must specify the size of the key, `key-size`. Default is 1024.

ecdsa

Uses the ECDSA key. You must specify the elliptic curve size, `curve`. Default is 256.

> **NOTE**
> Attempting to install a CA signed ecdsa 256/384 bit certificate fails with an error similar to `Invalid certificate`.

Create and enroll a self-signed certificate

Creates and enrolls a self-signed local certificate. Including the subject will override the configured identity profile.

Syntax

```
crypto pki enroll-self-signed certificate-name CERT-NAME [subject [command-name CN-
Value] [org Org-Value] [org-unit Org-unit-value] [locality Location-Value] [state
state-Value] [countryCountry-Code] [valid-start date valid-end date] [usage
<openflow | web | all>][key-type rsa key-size <1024|2048>] [key-type ecdsa curve
<256|384>]
```

Definitions

---

profile-name

A name (maximum 100 characters) with a unique identifier for the Trust Anchor Profile. Ten TA profiles are supported: one for each allowed trust anchor (Root CA certificate.)

usage

When `usage` is set to `all`, it includes OpenFlow and web applications, as well as other applications such as syslog.

Options

rsa

Uses the RSA key. You must specify the size of the key, `key-size`. Default is 1024.

ecdsa

Uses the ECDSA key. You must specify the elliptic curve size, `curve`. Default is 256.

Configure or remove the minimum levels of security minLos for TLS

Configures the minimum levels of security for TLS to comply with Suite B. If strict mode is configured, only TLS 1.2 connections with ciphers compatible with 128 and 192 bits are accepted. For non-strict mode (which is the default option), the TLS 1.0 and later connections are supported.

Syntax

```
[no] crypto SuiteB-MinLoS <128|192> tls [strict]
```

Definitions

no

If you specify the strict mode, the strict configuration is removed. If you do not specify the strict mode, the minLoS and the strict configuration on the TLS are removed.

SuiteB-MinLoS

Configures Suite B minimum levels of security for TLS.

128

This security level matches Elliptic Curve P-256.

192

This security level matches Elliptic Curve P-384.

tls

Configures Suite B compliant minimum levels of security for TLS.

Options

strict

Enforces the strict mode.

Install authentication files

Install authentication files for the SSH server as follows:

Syntax

```
crypto key generate ssh [rsa|dsa]
```

Definitions

key

Operation on an SSH key file.

generate

Installs a new key.

	Installing a new key might be very slow in the first few minutes after booting the device.

Options

rsa

Specifies a RSA key type.

dsa

Specifies a DSA key type.

Install authentication files for Autorun as follows:

Syntax

```
crypto key generate autorun-key [rsa]
```

Definitions

key

Operation on an SSH key file.

generate

Installs a new key.

	Installing a new key might be very slow in the first few minutes after booting the device.

Options

rsa

Specifies a RSA key type.

Remove authentication files

Removes authentication files for the SSH server or for Autorun.

Syntax

```
crypto key zeroize [ssh | autorun | ssh-client-key | ssh-client-known-hosts]
```

Definitions

key

Operation on an SSH key file.

zeroize

Removes the existing key.

autorun

Installs or removes the RSA key for Autorun.

ssh-client-key

Deletes SSH client key pair.

ssh-client-known-hosts

Removes the SSH client known hosts file.

show crypto client-public-key

Syntax

```
show crypto client-public-key [babble] [fingerprint] [manager] [operator]
```

Description

View the client public keys configured on the switch.

Options

babble

Display phonetic hash.

fingerprint

Display hexadecimal hash.

manager

Select manager public keys.

operator

Select operator public keys.

Remove the client public keys from configuration

Removes the currently loaded authorized client public keys from the active configuration. By default, the operator client public keys are removed.

Syntax

```
clear crypto client-public-key [<operator|manager> key-type <dsa|rsa>]
```

Options

dsa

Removes the DSA key.

rsa

Removes the RSA key.

Show details of TA profile

Shows the details of the Trust Anchor profile specified.

Syntax

```
show crypto pki ta-profile TA-Profile-Name detail
```

Full syntax example

```
show crypto pki ta-profile crl TA-profile-name crl certificate-serial-numserial-num
```

Displays all CRLs available in all TA-profiles. The option `certificate-serial-num` is only used when `crl` option is used in the `show` CLI.

Syntax

```
show crypto pki ta-profile [TA-PROFILE-NAME] [crl | detail] [certificate-serial-num <SERIAL-NUM>]
```

Options

show crypto pki ta-profile

Displays all the CRLs available in all TA-profiles.

TA-PROFILE-NAME

A name (maximum 100 characters) with a unique identifier for the Trust Anchor Profile. Ten TA profiles are supported: one for each allowed trust anchor (Root CA certificate.)

crl

Shows the CRL details of the TA profile

detail

Shows the configuration details of the TA profile.

SERIAL-NUM

Serial number of the certificate whose revocation information is required.

NOTE

The option `certificate-serial-num` is only used when the `crl` option is used in the `show` CLI; for example:

```
show crypto pki ta-profile [TA-profile-name] crl certificate-serial-num <serial-num>
```

Networking Websites

Hewlett Packard Enterprise Networking Information Library

> **www.hpe.com/networking/resourcefinder**

Hewlett Packard Enterprise Networking Software

> **www.hpe.com/networking/software**

Hewlett Packard Enterprise Networking website

> **www.hpe.com/info/networking**

Hewlett Packard Enterprise My Networking website

> **www.hpe.com/networking/support**

Hewlett Packard Enterprise My Networking Portal

> **www.hpe.com/networking/mynetworking**

Hewlett Packard Enterprise Networking Warranty

> **www.hpe.com/networking/warranty**

General websites

Hewlett Packard Enterprise Information Library

> **www.hpe.com/info/EIL**

For additional websites, see **Support and other resources**.

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:

  http://www.hpe.com/assistance

- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:

  http://www.hpe.com/support/hpesc

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates:

  Hewlett Packard Enterprise Support Center
      www.hpe.com/support/hpesc
  Hewlett Packard Enterprise Support Center: Software downloads
      www.hpe.com/support/downloads
  Software Depot
      www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

  www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

www.hpe.com/support/AccessToSupportMaterials

ⓘ   Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

http://www.hpe.com/support/selfrepair

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information
HPE Get Connected
    www.hpe.com/services/getconnected
HPE Proactive Care services
    www.hpe.com/services/proactivecare
HPE Proactive Care service: Supported products list
    www.hpe.com/services/proactivecaresupportedproducts
HPE Proactive Care advanced service: Supported products list
    www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information
Proactive Care central
    www.hpe.com/services/proactivecarecentral
Proactive Care service activation
    www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty for your product or to view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* reference document, go to the Enterprise Safety and Compliance website:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional warranty information
HPE ProLiant and x86 Servers and Options
    www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (**docsfeedback@hpe.com**). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.