

AOS-CX 10.13.1000 Hardening Guide

All Switch Series



a Hewlett Packard
Enterprise company

Copyright Information

© Copyright 2024 Hewlett Packard Enterprise Development LP.

This product includes code licensed under certain open source licenses which require source compliance. The corresponding source for these components is available upon request. This offer is valid to anyone in receipt of this information and shall expire three years following the date of the final distribution of this product version by Hewlett Packard Enterprise Company. To obtain such source code, please check if the code is available in the HPE Software Center at <https://myenterpriselicense.hpe.com/cwp-ui/software> but, if not, send a written request for specific software version and product for which you want the open source code. Along with the request, please send a check or money order in the amount of US \$10.00 to:

Hewlett Packard Enterprise Company
Attn: General Counsel
WW Corporate Headquarters
1701 E Mossy Oaks Rd Spring, TX 77389
United States of America.

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise 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.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Contents	3
About this document	5
Latest version available online	5
About the examples	5
Identifying switch ports and interfaces	5
Identifying modular switch components	7
Overview	8
Hardening Objectives	8
Operational Assumptions	9
Syntax and Conventions	9
Software, Documentation, Security Advisories and Bug Bounty Program	9
Hardening the CX Management plane	10
Factory Defaults	10
Physical Security	12
Front Panel Security	12
USB Auxiliary Port	13
Firmware Validation	14
Enhanced security mode	14
ServiceOS password authentication	16
Securing Switch Management Access Control	16
User Management and Password Control	16
Security User Group	17
Hardening Password Rules	18
Password Complexity	18
Authentication, Authorization and Accounting (AAA)	20
Authentication	21
Authorization	25
Accounting	26
RadSec over RADIUS	27
Hardening SSH	28
Public Key Authentication	28
Allow List	28
Recommended Ciphers, MACs, and Algorithms	29
Server Port Customization	29
Two Factor Authentication and Authorization	30
Summary	31
Session Management	31
Limiting Shell Access	32
Securing SNMP Access	32
Control Plane ACLs	34
Time Synchronization	34
Secure Copy	35
Hardening PKI	35

Mandatory matching of peer device hostname	37
EST	38
TLS Enforcements	38
Secure Logging	38
Hardening the Control Plane	39
Control Plane Policing	39
Securing Spanning Tree	40
BPDU Protection	41
Root Protection	41
DHCP Security	41
DHCP Snooping	42
DHCPv6 Guard	43
Dynamic ARP Inspection	44
ND Snooping Attack Prevention	45
RA Guard	46
IPv6 Destination Guard	47
IP Source Lockdown	47
Securing Routing Protocols	48
OSPF Passive Interfaces	48
OSPF Neighbor Authentication	49
OSPFv3 Area Authentication and Encryption with IPsec	49
BGP	50
Control Plane ACL for BGP Peering Sessions	51
Authenticate BGP Peers Using MD5	51
BGP TTL Security	51
Multicast Security	52
SSDP	52
Hardening IGMP and MLD Snooping	53
Hardening PIM and PIMv6	54
PIM Accept-Register	54
PIM Accept-RP	54
PIM SSM	54
Securing MSDP	55
NAE Scripts	56
Trusted Supply Chain	57
Support and Other Resources	58
Accessing HPE Aruba Networking Support	58
Accessing Updates	59
Warranty Information	59
Regulatory Information	59
Documentation Feedback	59

This document describes features of the AOS-CX network operating system. It is intended for administrators responsible for installing, configuring, and managing Aruba switches on a network.

Latest version available online

Updates to this document can occur after initial publication. For the latest versions of product documentation, see the links provided in [Support and Other Resources](#).

About the examples

Examples in this document are representative and might not match your particular switch or environment.

The slot and port numbers in this document are for illustration only and might be unavailable on your switch.

Understanding the CLI prompts

When illustrating the prompts in the command line interface (CLI), this document uses the generic term **switch**, instead of the host name of the switch. For example:

```
switch>
```

The CLI prompt indicates the current command context. For example:

```
switch>
```

Indicates the operator command context.

```
switch#
```

Indicates the manager command context.

switch(CONTEXT-NAME)#

Indicates the configuration context for a feature. For example:

```
switch(config-if) #
```

Identifies the **interface** context.

Variable information in CLI prompts

In certain configuration contexts, the prompt may include variable information. For example, when in the VLAN configuration context, a VLAN number appears in the prompt:

```
switch(config-vlan-100) #
```

When referring to this context, this document uses the syntax:

```
switch(config-vlan-<VLAN-ID>) #
```

Where <VLAN-ID> is a variable representing the VLAN number.

Identifying switch ports and interfaces

Physical ports on the switch and their corresponding logical software interfaces are identified using the format:

member/slot/port

On the 4100i Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on the switch.

On the 6000 and 6100 Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on the switch.

On the 6200 Switch Series

- *member*: Member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 8. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 in slot 1 on member 1.

On the 6300 Switch Series

- *member*: Member number of the switch in a Virtual Switching Framework (VSF) stack. Range: 1 to 10. The primary switch is always member 1. If the switch is not a member of a VSF stack, then member is 1.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on member 1.

On the 6400 Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *slot*: Specifies physical location of a module in the switch chassis.
 - Management modules are on the front of the switch in slots 1/1 and 1/2.
 - Line modules are on the front of the switch starting in slot 1/3.
- *port*: Physical number of a port on a line module.

For example, the logical interface **1/3/4** in software is associated with physical port 4 in slot 3 on member 1.

On the 83xx, 9300, and 10000 Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *slot*: Always 1. This is not a modular switch, so there are no slots.
- *port*: Physical number of a port on the switch.

For example, the logical interface **1/1/4** in software is associated with physical port 4 on the switch.



If using breakout cables, the port designation changes to x:y, where x is the physical port and y is the lane when split to 4 x 10G or 4 x 25G. For example, the logical interface 1/1/4:2 in software is associated with lane 2 on physical port 4 in slot 1 on member 1.

On the 8400 Switch Series

- *member*: Always 1. VSF is not supported on this switch.
- *slot*: Specifies physical location of a module in the switch chassis.
 - Management modules are on the front of the switch in slots 1/5 and 1/6.
 - Line modules are on the front of the switch in slots 1/1 through 1/4, and 1/7 through 1/10.
- *port*: Physical number of a port on a line module

For example, the logical interface **1/1/4** in software is associated with physical port 4 in slot 1 on member 1.

Identifying modular switch components

- Power supplies are on the front of the switch behind the bezel above the management modules. Power supplies are labeled in software in the format: *member/power supply*:
 - *member*: 1.
 - *power supply*: 1 to 4.
- Fans are on the rear of the switch and are labeled in software as: *member/tray/fan*:
 - *member*: 1.
 - *tray*: 1 to 4.
 - *fan*: 1 to 4.
- Fabric modules are not labeled on the switch but are labeled in software in the format: *member/module*:
 - *member*: 1.
 - *member*: 1 or 2.
- The display module on the rear of the switch is not labeled with a member or slot number.

Security is a growing concern in today's all-digital enterprise infrastructure. Upper-level managers and IT administrators alike are held to higher accountability for the integrity and availability of their critical data and infrastructure. While clients and servers are often the focus of security discussions, the security of network devices such as switches, routers, and wireless access points should not be ignored. Critical enterprise data traverses these devices, and properly securing them is paramount to a stable and secure infrastructure.

The HPE Aruba Networking CX switching platform, powered by the AOS-CX network operating system, simplifies network operations by delivering automation, distributed analytics, security, and high availability to campus and data center networks. The microservices architecture around which AOS-CX is built delivers network-wide analytics and full programmability to enable complete network assurance.

The purpose of this document is to provide security guidelines and best practices for management features and protocols provided by the AOS-CX software, and to present sample configurations to illustrate these best practices in action. This document is not intended to be a comprehensive reference guide to the features and commands listed; for additional information on configuration syntax and advanced features referred to in this document, please obtain the latest software manual set from the [HPE Aruba Networking Support Portal](https://supportportal.arubanetworks.com).

Some of the features in this document are not be available on all switch platforms. Refer to <https://feature-navigator.arubanetworks.com> for a list of features supported on each platform.

Hardening Objectives

IETF BCP 61 points to a few definitions that help us define our goals, which we can summarize into three helpful points:

- **Authentication:** A security service that verifies an identity. This identity could be a user, a device, or a process.
- **Data Confidentiality:** A security service that protects data against unauthorized disclosure to unauthorized individuals or processes.
- **Data Integrity:** A security service that protects against unauthorized changes to data. Changes include intentional change and accidental change.

The applications and procedures we use in this document leverage these summarized definitions above and help to shape the following general guidelines:

- If there are methods, we can use to ensure the identities of the users and devices with which we interact, we should prefer these over insecure alternatives.
- We should limit the exposure to the equipment from sources we cannot trust, whenever possible. We should also make attempts to utilize encryption methods so that our data is not easily read by anyone besides the trusted receiver of the data.
- Assume that eventually, an event will occur that causes a need for reliable information we know we can trust. We need to make sure this data is safe, available for us to access, and unavailable to anyone else.

- This document helps you improve the overall network security by hardening the security of the management and control plane.

Operational Assumptions

- One or more authorized administrators are assigned who are competent to manage the device and the security of the information it contains, trained for the secure operation of the device, and who can be trusted not to deliberately abuse their privileges to undermine security.
- Authorized users are trusted to correctly install, configure, and operate the device according to the instructions provided by the device documentation.
- There will be no untrusted users and no untrusted software on component servers.
- The switch must be installed in a physically secure area where only authorized administrators have access to the physical device.
- Users will protect their authentication data.

Syntax and Conventions

This document provides examples for each configurable feature discussed. These examples follow a common format: commands and fixed options appear as fixed-width regular text, while configurable parameters appear in *italics*, as in the following:

```
switch(config)# ssh server vrf default
```

For more details on command syntax, refer to the documentation referenced for each feature, or use the built-in command syntax help on the switch by typing a partial command, then typing **?** (a question mark) to see possible options and parameters for that command.

Software, Documentation, Security Advisories and Bug Bounty Program

HPE Aruba Networking CX switch software, release notes, and user documentation can be found at the [HPE Networking Support Portal](#).

Security advisories are published on the [Aruba Security Advisory archive](#), and notification services are provided by a Security Alerts mailing list, with subscriptions offered via the [self-service portal](#). For more information, refer to the [Security Incident Response Policy](#).

HPE Aruba Networking also runs a Bug Bounty program for reporting security exploits and vulnerabilities. HPE Aruba Networking handles and discloses vulnerabilities in accordance with ISO/IEC 30111. Refer to <https://bugcrowd.com/aruba-product-public> for more information on the Bug Bounty program

The following section describes strategies to secure the switch management plane.

Factory Defaults

Once the device boots, it is essential for an administrator to immediately connect to it and configure a password for the admin account. California signed into law bill SB-327 in 2018, requiring manufacturers of networking equipment to force users to create a password when they first connect to a device.

In a factory default state, AOS-CX devices are configured with the default user admin with no password. The user is prompted to create a password before access is given to the CLI, Web UI, and REST API:

```
Please configure the 'admin' user account password.  
Enter new password: *****  
Confirm new password: ****
```

The built-in management interface provides a way to access and manage the switch that is segregated from production traffic. Internal networks separated from production traffic are typically referred to as Out-Of-Band-Management networks (OOBM). By limiting the clients allowed to manage devices to only those who reside on the OOBM network, we sharply limit the large set of devices that can attempt to control the device.

In AOS-CX, the management interface is logically separated from the rest of the switch by means of a unique virtual routing and forwarding table (VRF), named the mgmt VRF. Please note that the mgmt VRF is unique in that it is permanently assigned to the physical management port and cannot be associated with any other switch interface; the management port itself cannot be associated with any other VRF.

The management interface is enabled by default to learn an IP address via DHCP. To configure the management interface with a static IP address, gateway, and DNS:

```
switch(config)# interface mgmt  
switch(config-if-mgmt)# ip static 10.1.1.5/24  
switch(config-if-mgmt)# default-gateway 10.1.1.1  
switch(config-if-mgmt)# nameserver 10.0.1.10 10.0.1.11
```

To show the status of the management interface:

```
switch# show interface mgmt  
Secondary Nameserver : 10.0.1.11  
Address Mode : static  
Admin State : up  
Mac Address : d0:67:26:11:11:11  
IPv4 address/subnet-mask : 10.1.1.5/24  
Default gateway IPv4 : 10.1.1.1  
IPv6 address/prefix :  
IPv6 link local address/prefix : fe80::d267:2611:1111:1111/64  
Default gateway IPv6 :
```

```
Primary Nameserver : 10.0.1.10
```

The other VRFs available on an AOS-CX device upon first boot is the default VRF. The default VRF is automatically associated with all non-management interfaces, including Layer 3 routed ports, non-routed ports, and switched VLAN interfaces (SVIs) created on the switch, unless the interface is explicitly attached to another VRF.

The following management services are enabled by default on an AOS-CX switch:

- SSH on TCP port 22
- WebUI and read/write REST API on TCP port 443. (Any Connections to TCP port 80 will be automatically redirected to TCP port 443)

The 10000, 8xxx and 9300 Switch series ship with these services enabled only on the mgmt VRF, while 6400, 6300 and 6200 switches ship with these services enabled on both the default and mgmt VRFs.

For optimal security, manage switches from a dedicated management network when possible, and disable management services on all other VRFs.

To disable management services on all other VRFs :

```
switch(config)# no ssh server vrf <vrf-name>
switch(config)# no https-server vrf <vrf-name>
```

To view the configuration change :

```
switch # show ssh server vrf vrf1
SSH server is not enabled on VRF vrf1.
```

As the 6100, 6000 and 4100i Switch series does not have a dedicated management port or the associated VRF, management services are enabled only on the default VRF. Therefore, disabling management services is not a feasible solution for these platforms. For these switches, use other alternatives, such as [Control Plane ACLs](#) , an [SSH Allow list](#) and [Per-User management interface enablement](#) features to protect the management services.

Restoring the Switch to Factory Defaults

The recommended method to return an AOS-CX switch to factory default settings is to zeroize it. The following occurs when the zeroization process is initiated:

- The switch reboots to ServiceOS
- Primary and secondary software image files are backed up to memory from flash storage
- The entire flash storage device is overwritten with zeroes to securely erase all stored data
- The flash storage device is reformatted with a factory default filesystem
- Backed up software image files are written to flash in their original locations
- The switch reboots to the primary software image with a default configuration

There are four methods that may be used to zeroize a switch. First, an admin user may use the erase all zeroize command from the AOS-CX CLI:

```
switch# erase all zeroize
This will securely erase all customer data and reset the switch to factory
```

```
defaults. This will initiate a reboot and render the switch unavailable until the
zeroization is complete.This should take several minutes to one hour to complete.
Continue (y/n)?
```

Second, an admin user may use the erase zeroize command from the ServiceOS CLI:

```
SVOS> erase zeroize
#####WARNIN
G#####
This will securely erase all customer data and reset the switch to factory
defaults. This will initiate a reboot and render the switch unavailable until the
zeroization is complete.This should take several minutes to one hour to complete.
#####WARNI
G#####
Continue (y/n)?
```

Third, a user with physical access to the switch front panel and a FAT32-formatted USB storage device may zeroize the switch from the ServiceOS login prompt by entering the username zeroize and following the provided instructions:

```
ServiceOS login: zeroize
This will securely erase all customer data, including passwords, and reset the
switch to factory defaults.
This action requires proof of physical access via a USB drive.
* Create a FAT32 formatted USB drive
* Create a file in the root directory of the USB drive named zeroize.txt
* Type the following serial number into the zeroize.txt file: xxxxxxxxxx
* Insert the USB drive into the target module
* Confirm the following prompt to continue
Continue (y/n)?
```

Finally, changing the switch security mode results in the switch being zeroized; see the [Enhanced security mode](#) section for more information.

Physical Security

The following sections describe physical security hardening workflows.

Front Panel Security

AOS-CX switches include a reset button on the front panel to allow users to perform the following reset operations:

Reset Type	Procedure	Outcome
Soft Reset	Press the reset button and release it before 5 seconds.	The switch operating system is cleared gracefully. The switch then reboots and runs self-tests.
Hard Reset	Press the reset button and release it between 5 to 20 seconds.	The switch reboots, like a power cycle. A hard reset is used, for example, when the

Reset Type	Procedure	Outcome
		switch CPU is in an unknown state or not responding.
Factory Reset	Press the reset button and release it between 20 to 25 seconds.	The switch will undergo the factory reset process.



Factory Reset functionality is available from 10.13.1000 release in HPE ANW CX 6000 and 6100 series of switches.

This factory reset capability creates a security and denial-of-service risk if the switch is in a location where it is impossible to prevent physical access to the front panel. It is disabled by default and recommended that administrators disable this feature after its usage to prevent malicious use by an attacker with physical access to the device.

```
switch(config)# front-panel-security factory-reset
This command will enable front-panel factory reset capability, where user can
trigger factory-reset via reset button. This feature will remain enabled until
it is disabled, or a factory-reset is performed.
Continue (y/n)?
```

To view the configuration change -

```
switch# show front-panel-security status
Front panel factory reset           : disabled
First occurrence of front-panel factory reset : N/A
```

USB Auxiliary Port

The AOS-CX switch front-panel includes an USB Auxiliary port for the following purposes –

- USB Mass storage - flash drive for deploying, troubleshooting, backing up configurations, or upgrading switches
- Bluetooth Adapter - allows Bluetooth enabled devices to connect to and manage the switch on a wireless Bluetooth Personal Area Network (PAN)

The Bluetooth feature has been enabled by default in AOS-CX switches and designed for operational simplicity. The switch provides an IP address to paired devices through DHCP when they join the Bluetooth Personal Area Network. Paired devices can then manage the switch through following methods

- SSH
- Web UI
- REST API
- Aruba CX Mobile App

Refer to [Securing Switch Management Access Control](#) for details on securing these management connections.

This USB Auxiliary port is enabled by default so recommended to be disabled when not in use, and only temporarily enabled when needed.

To disable the USB Auxiliary port entirely (USB Mass Storage and Bluetooth Adapter), use the following command:

```
switch(config)# no usb
```

To view the configuration change :

```
switch # show usb
Enabled: No
Mounted: No
```

AOS-CX switches also have the support to disable only the bluetooth feature rather than disabling the USB Auxiliary port completely, to perform the same following configuration can be executed which is enabled by default :

```
switch(config)# bluetooth disable
switch # show bluetooth
Enabled           : No
Device Name       : 6300-SG9ZKN002Z
Adapter State     : Absent
```

Firmware Validation

All AOS-CX switch firmware is signed by HPE at the time the firmware is created. The firmware signature is verified at the time of download and verified at every boot. The public keys used to verify the firmware is stored within the bootloader and firmware. The firmware is digitally signed with RSA-3072 and SHA-256.

If the switch firmware validation fails at boot, the switch will fail to boot with one of the following error messages and drop the user into the ServiceOS login screen:

```
Error: Signature verification failed
Error: Signature not found
Error: Invalid signature
```

Alternatively, after loading the firmware to the boot bank – primary or secondary , administrators can verify the firmware integrity using below show command before booting the switch.

The verify option performs an integrity check that the image has a valid signature and is compatible with this system. The switch prevents the download of firmware without a valid signature.

```
switch# show images verify primary
The primary image is valid
switch# show images verify secondary
The image does not contain a signature
```

Enhanced security mode

AOS-CX provides two security modes that control access to certain system management features — standard and enhanced. All AOS-CX switches operate in standard mode by default, with no system-level restrictions in place for any functionality. The enhanced security mode disables access to the start-shell command in the AOS-CX CLI, as well as the ServiceOS commands config-clear, password, sh, and update.

In a dual management module switch, both the management modules should be set to same secure mode.

Changing the switch security mode is performed either through CLI or from the ServiceOS shell. All changes to the switch security mode (**standard** to **enhanced** or **enhanced** to **standard**) result in zeroization of the filesystem and a reset to factory defaults.

```
6300-VSF(config)# secure-mode enhanced
This will set the switch into enhanced secure mode. Before enhanced secure mode is
enabled, the switch must securely erase all customer data and reset to factory
defaults. This will initiate a reboot and render the switch unavailable until the
zeroization is complete.
Continue (y/n)? y
```

ServiceOS be used a secondary method to boot the switch in enhanced secure-mode. Reboot the switch to ServiceOS using the following command:

```
switch# boot system serviceos
One time boot to ServiceOS initiated.
Checking if the configuration needs to be saved...
This will reboot the system to ServiceOS and render the entire switch unavailable.
Access to ServiceOS is only available through the serial console.
Continue (y/n)?
```

Once the switch has rebooted and the ServiceOS login prompt is displayed, login as admin (no password is set by default). Use the following command to enable enhanced security mode:

```
SVOS> secure-mode enhanced
#####WARNING#####
This will set the switch into enhanced secure mode. Before enhanced secure mode
is enabled, the switch must securely erase all customer data and reset the switch
to factory defaults. This will initiate a reboot and render the switch unavailable
until the zeroization is complete.
#####WARNING#####
Continue (y/n)?
```

Entering **y** will cause the switch to reboot, zeroize the filesystem, then reboot an additional time.

To revert to the standard security mode, reboot to ServiceOS as above, login as admin, then use the following command:

```
SVOS> secure-mode standard
#####WARNING#####
This will set the switch into standard secure mode. Before standard secure mode
is enabled, the switch must securely erase all customer data and reset the switch
to factory defaults. This will initiate a reboot and render the switch unavailable
until the zeroization is complete.
#####WARNING#####
Continue (y/n)?
```


ServiceOS shall default to standard secure mode if Zeroization fails while setting to standard or enhanced secure mode.

When this setting is enabled, logging in to the ServiceOS shell with the admin user requires the same password used to authenticate the admin user in the AOS-CX CLI or Web UI.

If this setting is enabled, a forgotten admin user password cannot be reset using ServiceOS; if there are no other local or RADIUS/TACACS user accounts with administrator-level access, the switch must be zeroized by entering the username zeroize at the ServiceOS login prompt to restore administrator access. See [password reset](#) for more information.

ServiceOS password authentication

By default, the ServiceOS shell (accessible only from the local switch console port) requires no password to login as admin; to enable password authentication for ServiceOS, use the following command from the configuration context:

```
switch(config)# system serviceos password-prompt
```

When this setting is enabled, logging in to the ServiceOS shell with the admin user requires the same password used to authenticate the admin user in the AOS-CX CLI or Web UI.

If this setting is enabled, a forgotten admin user password cannot be reset using ServiceOS; if there are no other local or RADIUS/TACACS user accounts with administrator-level access, the switch must be zeroized by entering the username zeroize at the ServiceOS login prompt to restore administrator access. See [password reset](#) for more information.

Securing Switch Management Access Control

Use the following console, SSH, Telnet, and HTTPS server strategies to secure the switch management access.

User Management and Password Control

User Groups

A factory-default switch comes with a single user named admin member of built-in administrators group. Up to 63 local users can be added, for a total of 64 users including the default user admin. A user can belong to only one group. The switch provides the following built-in user groups with corresponding roles. Each of these roles comes with a set of privileges.

- Administrators—full access (privilege level 15)
 - Perform firmware upgrades
 - Make configuration changes
 - View all switch configuration information, including sensitive data such as ciphertext passwords
 - Add and remove local user accounts, and change user passwords
 - All REST interface methods (GET, PUT, POST, PATCH, DELETE) can be used
- Operators – limited access (privilege level 1)

- Display-only CLI access
- View non-sensitive configuration information
- Only the REST interface GET method can be used
- Auditors – limited access (privilege level 19)
 - Access to Commands in “auditor” context only
 - Web-UI “system->Log Page” view only.
 - REST Interface GET method available only for following resources only
 - Audit log: /logs/audit
 - Event log: /logs/event

Apart from the built-in groups, the switch enables you to create up to 29 user-defined local user groups, for the purpose of configuring local authorization. Local authorization uses role-based access control (RBAC) to provide role-based privilege levels plus optional user-defined local user groups with command execution rules. Each of the 29 user-defined groups support up to 1024 CLI command authorization rules that define what CLI commands can be executed by members of the group.

Sample Configuration to create user-defined local user group:

```
switch(config)# user-group sample-group
switch(config-usr-grp-sample-group)# 10 comment Deny all show aaa commands
switch(config-usr-grp-sample-group)# 10 deny cli command "show aaa .*"
switch(config-usr-grp-sample-group)# 20 comment Permit all other show commands
switch(config-usr-grp-sample-group)# 20 permit cli command "show .*"
switch(config-usr-grp-sample-group)# exit
```

```
6200(config)# show user-group
```

GROUP NAME	GROUP TYPE	INCLUDED GROUP	NUMBER OF RULES
administrators	built-in	n/a	n/a
auditors	built-in	n/a	n/a
operators	built-in	n/a	n/a
sample-group	configuration	--	2

Security User Group

Security log commands for showing, clearing, and copying the security logs can be made available to a security user. To have a security user, the admin must create a security user group and add a user to the group. The admin must also grant permission to members of the security user group for the three security log commands. Only users who are members of the security group have permission to execute the security log commands. The admin user who created the security user group does not have permission to use the security log commands:

```
switch(config)# user-group security-group
switch(config-usr-grp-security-group)# permit cli command "show security-logs*"
switch(config-usr-grp-security-group)# permit cli command "clear security-logs"
switch(config-usr-grp-security-group)# permit cli command "copy security-log*"
switch(config-usr-grp-security-group)# exit
switch(config)# user security-user group security-group password
Adding user security-user
Enter password:*****
Confirm password:*****
```

Showing the security logs:

```
switch# show security-logs
```

Copying the security logs:

```
switch# copy security-log sftp://user1@99.99.99.99/coredump.xz vrf mgmt
```

Hardening Password Rules

When managing an AOS-CX Switches, setting up a secure network is essential. A crucial factor in security is the selection of a strong password. Passwords are never displayed in plaintext format in CLIs and config files. Passwords are encrypted when stored in the config file .

Passwords must:

- Contain only ASCII characters from decimal 33 to 126 (Hexadecimal 21 to 7E). Spaces are not allowed
- Contain at most 64 characters.

Passwords are portable to different switch using default or customer configured non-default export key. The password complexity feature will help organization to set password policy for their administrators

Password Complexity

The password complexity feature helps in enforcement of complexity rules when configuring local user account passwords. It is disabled by default. The password complexity feature will help organization to set password policy for their users. Remember to enable the password complexity feature after configuring it for the rules to be enforced. Enabling or changing password complexity settings affects password creation or password change after the password complexity feature is enabled or changed.

The following enforcement will apply to new user creation or a password update once the password complexity feature is enabled:

- User creation/Password update with `ciphertext-password` is not allowed, because password complexity check cannot be performed on ciphertext password.
- The following password complexity check will be enforced

```
switch(config)# password complexity
switch(config-pwd-cplx)# minimum-length 9
switch(config-pwd-cplx)# history-count 4
switch(config-pwd-cplx)# position-changes 5
switch(config-pwd-cplx)# enable
switch(config-pwd-cplx)# exit
switch # show password-complexity
Global password complexity checking criteria:
  Password complexity                : Enabled
  Previous passwords to check        : 4
  Minimum password length            : 9
  Minimum position changes           : 5
  Maximum adjacent characters count   : 0
  Password composition
    Minimum lowercase characters      : 1
    Minimum uppercase characters      : 1
    Minimum special characters        : 1
    Minimum numeric characters        : 1
```

Non-Default Export Password

The export password is used to transform critical sensitive information into ciphertext suitable for exporting and showing by commands such as show running-config. Transformation enables safe switch configuration import and export. All factory-default switches have identical default export passwords. For security, it is recommended that you set the same non-default export password on every switch in a group that will exchange sensitive configuration information. Only switches with identical export passwords can exchange sensitive configuration information.

```
switch# show service export-password
Export password: default
switch# config t
switch(config)# service export-password
Enter password: *****
Confirm password: *****
switch(config)# show service export-password
Export password: custom
```

Built-in Admin Account Password Reset

When administrators forget their switch console passwords, they must endure a time-consuming reset process, resulting in loss of productivity. If there are multiple administrators for the switch, it is recommended to reset the password using another administrator account. There are two ways to reset the password in case there is single admin user only for the switch:

Reverting the switch to factory defaults

1. At the manager command prompt, enter erase startup-config.


```
switch(config)# erase startup-config
```

2. Boot the switch without saving the current configuration

```
switch# boot system
Do you want to save the current configuration (y/n)? n
This will reboot the entire switch and render it unavailable until the
process is complete.
Continue (y/n)? y
The system is going down for reboot.
```

Resetting the switch admin password using the serviceOS console

Perform this task only when the switch admin user password has been forgotten:



Refer to the “Managing users and groups” section of the [Security Guide](#) for your switch model for more information.

Authentication, Authorization and Accounting (AAA)

AOS-CX have following management interfaces for accessing the switch for configuration and management -

- Console
- SSH
- Telnet
- https-server – Web UI and REST API



Telnet is not a recommended method to access the switch for configuration and management , as it is not a secure communication. It is not enabled by default on any VRF.

User accounts for accessing these management interfaces can be stored locally or managed on remote TACACS+ or RADIUS servers. AAA (Authentication, Authorization, and Accounting) is the security framework to manage user access, enforce privileges, and log the user access records.

The following table describes supported AAA services based on the user account management methods:

User Account Management	Local	TACACS+	RADIUS
Authentication	Yes	Yes	Yes

User Account Management	Local	TACACS+	RADIUS
Authorization	Yes, RBAC	Yes, Per Command Authorization and RBAC	No
Accounting	Yes	Yes	Yes

Authentication

Authentication is the process of identifying a user and granting them access to the network. Most of the time, this is done through traditional username and password credential , but it could be extended to SSH public key authentication. The following table describes supported authentication types based on their user account management methods.

User Account Management	Local	TACACS+	RADIUS
Authentication Type	■ Username/Password ■ SSH Public Key ■ SSH two-factor Authentication	Username/Password	■ Username/Password ■ SSH two-factor Authentication

Local Authentication

Local user names and passwords are configured on a per-switch basis and provide the most basic form of authentication. Local authentication is often used as the fallback login method. Local authentication can provide a minimum-security level should the primary method fail, but does not completely disable management access to the switch. To configure a local administrator-level user named localadmin with interactive password entry:

```
switch(config)# user localadmin group administrators password
Enter password: *****
Confirm password: *****
```

To create an operator-level user named **localoperator** with a plaintext password:

```
switch(config)# user localoperator group operators password plaintext abcdefghij
```

An administrator can also enter a password as a ciphertext string rather than being entered in plaintext. In AOS-CX, ciphertext passwords cannot be generated manually; they must be copied from another switch with the same export password configured. By default all the switches will have same export password. Refer to [Non-Default Export Password](#) for the configuration. Once the export passwords on the source and destination switches are the same, copy the ciphertext password from the source switch and apply it to the destination:

```
switch(config)# user localadmin group administrators password ciphertext
myCipherText
```



If password complexity is enabled, ciphertext password configurations are not allowed.

Local Authentication Configuration Task List

Task	Configuration	Show Commands
Enable local authentication for desired management interface Default – Includes all the management interfaces	aaa authentication login <console/default/ssh/telnet/https-server> local	show aaa authentication
Limit the login attempts	Console: aaa authentication console-login-attempts <1-10> console-lockout-time <1-3600s> SSH/Telnet/https-Server: aaa authentication login-attempts <1-10> lockout-time <1-3600s>	show aaa authentication show authentication locked-out-users

Remote Authentication

Remote Authentication involves the use of remote RADIUS , RadSec and TACACS+ servers for authenticating the management users. Remote AAA servers are used as single point of management to configure and store user accounts. They are often coupled with directories and management repositories, simplifying the setup and maintenance of the end-user accounts.

Remote Authentication Configuration Task List

Task	Configuration	Show Commands
Configure the server	RADIUS Server: radius-server host <IPv4/IPv6/FQDN> key plaintext <secret-key> vrf <vrf-name> RadSec Server: radius-server host <IPv4/IPv6/FQDN> key plaintext <secret-key> tls vrf <vrf-name> TACACS+ Server: tacacs-server host <IPv4/IPv6/FQDN> key plaintext <secret-key> vrf <vrf-name>	show radius-server detail show tacacs-server detail
Server Group Creation and Association. The order in which servers are added to a	RADIUS Server : aaa group server radius <group-name> server <IPv4/IPv6/FQDN> vrf <vrf-name> RadSec Server:	show aaa server-groups

Task	Configuration	Show Commands
group is important. The server added first is accessed first, and if necessary, the second server is accessed second, and so on.	<pre>aaa group server radius <group-name> server <IPv4/IPv6/FQDN> tls vrf <vrf-name></pre> TACACS+ Server : <pre>aaa group server tacacs <group-name> server <IPv4/IPv6/FQDN> vrf <vrf-name></pre>	
Enable local authentication for desired management interface Default – Includes all the management interfaces	<pre>aaa authentication login <console/default/ssh/telnet/https-server> <group-name></pre>	show aaa authentication
Auth-Type “chap” By default CX switch uses “pap” as auth-type. “Chap” is stronger authentication method than pap	<pre>radius-server auth-type chap tacacs-server auth-type chap</pre>	show radius-server show tacacs-server
Source Interface To ensure that all traffic sent from the switch to the AAA server uses the same source IP address	<pre>ip source-interface <radius/tacacs> <ip-address> ipv6 source-interface <radius/tacacs> <ipv6-address></pre>	Show ip source-interface Show ipv6 source-interface
Role Assignment in RADIUS	■ Aruba-Admin-Role VSA - Map the user to the matching local user-group name. ■ Aruba-Priv-Admin-User VSA - Extract the privilege level (1, 15, or 19) and map the user to the local user-group corresponding to this privilege level (1=operators,15=administrators, 19=auditors). Privilege levels 2 to 14 may also be used with matching local user groups named 2 to 14. ■ RADIUS Service-Type - Map Administrative-User(6)to administrators and map NASPrompt-User(7) to operators.	
Role Assignment in TACACS+	■ Aruba-Admin-Role VSA - Map the user to the matching corresponding local usergroup Name ■ TACACS+ priv-lvl attribute - Extract the privilege level (1, 15, or 19) and map the user to the local user-group corresponding to this privilege level (1=operators,15=administrators, 19=auditors). ■ Privilege levels 2 to 14 may also be used with matching local user groups named 2 to 14.	

Authentication Fallback and Fail through

To prevent authentication failure because of Remote AAA Server failure, it is recommended to configure more than one remote AAA Server

When defining the server access sequence for authentication with a aaa authentication login default, there is an implied local included as the last item in the list. If no remote AAA servers can be reached, local authentication will be attempted.

Normally, authentication success or failure is performed by the first reachable AAA server. A rarely needed feature named "Authentication fail-through" is available. If authentication fail-through is enabled and authentication fails on the first reachable AAA server, authentication is attempted on the second AAA server, and so on, until successful authentication or the server list is exhausted. Enabling Authentication fail-through is typically unnecessary because the user credential databases should be consistent across all AAA servers. Authentication fail-through might be helpful if your AAA user credential databases are not quickly synchronized across all AAA servers

To configure and view the authentication fail-through feature:

```
switch(config)# aaa authentication allow-fail-through
switch# show aaa authentication
AAA Authentication:
Fail-through                : Enabled
Limit Login Attempts        : Not set
Lockout Time                 : 300
Console Login Attempts      : Not set
Console Lockout Time        : 300
Authentication for default channel:
-----
GROUP NAME                  | GROUP PRIORITY
-----
local                       | 0
```

Per-User Management Interface Enablement

By default, switch users are enabled for accessing the switch through all these available management interfaces: ssh, telnet, https-server, console. Additionally fine-grained command authorization can be performed using RBAC , but it is applicable only for the CLIs not for Web-UI/REST API requests . Hence HPE ANW recommends enabling the specific management interfaces for the users based on the user type using below ways -

Local Per-user Management Interface Enablement

Local per-user management interface enablement is performed with CLI command . Example of disabling the SSH management interface for local user admin1.

```
switch(config)# no user admin1 management-interface ssh
switch(config)# show user-list management-interface
USER ENABLED MANAGEMENT INTERFACE(S)
-----
admin  ssh,telnet,https-sever,console
admin1 telnet,https-server,console
```

Remote TACACS+ and RADIUS

For remote TACACS+ and RADIUS servers, per-user management interface enablement is performed by configuring the AOS-CX VSA Aruba-User-Mgmt-Interface. On the TACACS+ or RADIUS server, the AOS-CX

VSA Aruba-User-Mgmt-Interface must be set to a comma-separated list of management interface names for which login is permitted by the associated user. Management interfaces omitted from the list are disabled for the associated user. A maximum of four management interface names are allowed, with each management interface name given once. Permitted management interface names (always lowercase) are as follows:

- ssh
- telnet
- https-server
- console

The VSA has a maximum length of 32 characters. The VSA is ignored by the switch if longer than 32 characters. When a user login fails because of an attempt to use a management interface that is not allowed, an event log is available indicating the enabled management interfaces as received in the TACACS+ or RADIUS VSA.

When using a RADIUS server other than ClearPass Policy Manager (CPPM), before setting the Aruba-User-Mgmt-Interface VSA, you must first define the VSA on the RADIUS server in file

```
ATTRIBUTE Aruba-User-Mgmt-Interface 69 string
```

Example RADIUS server VSA value that enables the two named management interfaces (ssh, telnet) while disabling the two unnamed management interfaces (https-server, console):

```
Aruba-User-Mgmt-Interface = "ssh,telnet"
```

Example RADIUS server VSA value that enables all four management interfaces:

```
Aruba-User-Mgmt-Interface = "ssh,telnet,https-server,console"
```

Authorization

Authorization controls how authenticated users execute commands and interact with the switch. Authorization uses role-based access control (RBAC) to provide role-based privilege levels plus optional user-defined local user groups with command execution rules.

```
Switch(config)# aaa authorization commands <console | default | ssh | telnet >  
group <tacacs | local | none > <tacacs | local | none>
```

- TACACS+ Authorization - Upon successful user authentication, the user is assigned their role by the TACACS+ server. See also User role assignment using TACACS+ attributes. TACACS+ authorization provides command filtering to allow/disallow individual command or command set execution. Each command is sent to the TACACS+ server for approval, and the switch then allows/disallows command execution according to the server response. TACACS+ authorization applies only to the CLI interface.
- RADIUS Authorization - Command authorization is not supported by RADIUS servers, however, user-defined local user groups can be configured with command-authorization rules, providing locally configured per command authorization for members of such groups.
- Fallback - Local authorization can be used as a fallback for the situation in which communication is lost with all TACACS+ servers after a successful authentication.

- When defining the server access sequence for authorization with above aaa authorization commands, it is recommended to always include either local or none as the last item in the list
- Failthrough – Authorization fail-through is recommended only for deployments where there are potential synchronization issues, so authorization will be failing in one server but succeeding in other.

```
Switch(config)# aaa authorization allow-fail-through
switch# show aaa authorization
***** Command authorization *****
Fail-through                : Enabled
Authorization for default channel:
-----
GROUP NAME                  | GROUP PRIORITY
-----
local                       | 0
```

Accounting

Local Accounting records all the CLI and REST access activities by users from all channels. It logs and helps to track all the configuration changes and show command executions happened at the switch for auditing or accounting purposes. This accounting information is captured and made available locally (Enabled by default and always active) and, if desired, for sending to remote AAA servers:

- Exec Accounting: user login/logout events.
- Command accounting: commands executed by users.
- System accounting: remote accounting On/Off events.
- Interactions on the non-CLI interfaces: REST and WebUI.

The following is not captured or made available as accounting information:

- CLI commands that reboot the switch.
- Interactions in the bash shell. (On the other hand, logging of “start-shell” CLI is supported . It helps in auditing)

Sample accounting information:

```
Switch# show aaa accounting log all
-----
Command logs from previous boots
-----
2023-06-09T05:50:27.765+00:00 acctsyslogd[2788]: AUDIT|CLI "enable" executed
by user 'admin' from address '0.0.0.0' through CONSOLE session which resulted
in success at timezone UTC.
```

Remote Accounting

For remote accounting, the information is sent to the first reachable remote server that was configured with this command for remote accounting. If no remote server is reachable, local accounting remains available by default

To enable and view the accounting configuration:

```
Switch(config)# aaa accounting all-mgmt <console|default|https-server|ssh|telnet>
start-stop <group|local>
switch# show aaa accounting
AAA Accounting:
  Accounting Type           : all
  Accounting Mode           : start-stop
  Accounting Fail-through   : Disabled

Accounting for default channel:
-----
GROUP NAME                  | GROUP PRIORITY
-----
local                       | 0
```

RadSec over RADIUS

The RADIUS protocol uses UDP as underlying transport layer protocol. RadSec is a protocol that supports RADIUS over TCP and TLS. In conventional RADIUS requests, security is a concern as the confidential data is sent using weak encryption algorithms. The access requests are in plain text includes information such as user name, IP address and so on. The user password is an encrypted shared secret. As a result, eavesdroppers can listen to these RADIUS requests and collect confidential information. Data protection is necessary in roaming environments where the RADIUS packets travel across multiple administrative domains and untrusted networks. The RadSec module secures the communication between the switch and RADIUS server using a TLS connection. Using RADIUS over TLS provides users with the flexibility to host RADIUS servers across geographies and WAN networks. HPE Aruba Networking recommends the usage of RadSec over RADIUS. Both IPv4 and IPv6 RadSec servers are supported.

To enable RADIUS security, use the **tls** parameter with the following command.



Refer to the Security Guide for your switch for detailed steps to associate the TLS certificate for mutual authentication.

```
Switch(config)# radius-server host <FQDN/ipv4/ipv6> tls
```

To view the RadSec server configuration:

```
switch# show radius-server
Unreachable servers are preceded by *
***** Global RADIUS Configuration *****
Shared-Secret: None
Timeout: 10
Auth-Type: pap
Retries: 3
Initial TLS Connection Timeout: 30
TLS Timeout: 5
Tracking Time Interval (seconds): 300
Tracking Retries: 3
Tracking User-name: radius-tracking-user
Tracking Password: None
Status-Server Time Interval (seconds): 300
Number of Servers: 1
```

```
AAA Server Status Trap: Disabled
```

SERVER NAME	TLS	PORT	VRF
cppm.abcd.net	Yes	1812	mgmt

Hardening SSH

The following sections describe security and hardening workflows for SSH.

Public Key Authentication

Passwords are easy to use and remember, but they are vulnerable to attacks and human errors. Keys are more secure and efficient compared to passwords. SSH Public key authentication is enabled by default and takes precedence over password-based authentication. Validate users identified with SSH public keys stored in the local user database using the following commands.

```
Switch(config)# user admin authorized-key ecdsa-sha2-nistp256 E2VjZH...QUiCAk=
root@switch
Switch# Show user <username> authorized-key
E2VjZH...QUiCAk= root@switch
```

Allow List

The SSH server access control can be implemented with an ACL applied to the control plane per VRF. A mistake in the configuration of the control-plane ACL applied to the **default** VRF might block other network protocols since the ACL involves rule ordering and can deny incoming packets. The SSH allow-list feature enhancement simplifies the configuration and protects against unauthorized SSH access. To use this feature, configure a list of addresses or prefixes that will be the only hosts allowed to connect to the SSH servers running on all VRFs of the switch. By default, the allow-list is disabled and any host is allowed to connect given the correct authentication criteria. When the allow-list is enabled, only the hosts that fall under one of the list entries may connect with the correct authentication criteria; all other hosts will be denied to attempt authentication.

```
switch(config)# ssh server allow-list
switch(config-ssh-al)# ip 10.10.0.0/16
switch(config-ssh-al)# ipv6 fd10::0/64
switch(config-ssh-al)# enable
Active SSH sessions will be terminated.
Do you want to continue (y/n)? y
switch(config-ssh-al)#exit
```

```
switch(config)# show ssh server allow-list
SSH server allow-list:
Status: Enabled
Allowed host IPs:
10.10.0.0/16
fd10::0/64
```



If the ACL is applied to the control-plane and the SSH allow-list is also enabled, the control-plane ACL has pre-emption over the SSH allow-list.

Recommended Ciphers, MACs, and Algorithms

AOS-CX switches by default supports the following SSH Ciphers, MACs, and Algorithms:

```
switch # show ssh server
SSH server configuration on VRF default :
IP Version      : IPv4 and IPv6      SSH Version      : 2.0
TCP Port        : 22                  Grace Timeout (sec) : 60
Max Auth Attempts : 6                  Server Status     : running
Allow-list: disabled

Ciphers:
chacha20-poly1305@openssh.com, aes128-ctr, aes192-ctr, aes256-ctr,
aes128-gcm@openssh.com, aes256-gcm@openssh.co

Host Key Algorithms

ecdsa-sha2-nistp256, ecdsa-sha2-nistp384, ecdsa-sha2-nistp521,
ssh-ed25519, rsa-sha2-256, rsa-sha2-512, ssh-rsa

Key Exchange Algorithms:
curve25519-sha256, curve25519-sha256@libssh.org, ecdh-sha2-nistp256,
ecdh-sha2-nistp384, ecdh-sha2-nistp521
MACs:
hmac-sha2-256-etm@openssh.com, hmac-sha2-512-etm@openssh.com,
hmac-sha1-etm@openssh.com, hmac-sha2-256, hmac-sha2-512, hmac-sha1

Public Key Algorithms:
rsa-sha2-256, rsa-sha2-512, ssh-rsa, ecdsa-sha2-nistp256,
ecdsa-sha2-nistp384, ecdsa-sha2-nistp521, ssh-ed25519,
x509v3-rsa2048-sha256, x509v3-ssh-rsa, x509v3-sign-rsa,
x509v3-ecdsa-sha2-nistp256, x509v3-ecdsa-sha2-nistp384,
x509v3-ecdsa-sha2-nistp521
```

The previously mentioned default ciphers , message authentication codes (MACs), and algorithms are based on OpenSSH's default settings and are deemed secure by the community.

For highly secure deployments like Federal Accounts which mandates the compliance of NDcPP (Common Criteria Protection Profile), it is recommended to configure the following list of ciphers , MACs, and algorithms as per the NDcPP evaluation criteria.

```
switch(config)# ssh ciphers aes128-ctr, aes256-ctr, aes128-cbc, aes256-cbc
switch(config)# ssh macs hmac-sha2-256, hmac-sha2-512, hmac-sha1
switch(config)# ssh key-exchange-algorithms ecdh-sha2-nistp256, ecdh-sha2-
nistp384, diffie-hellman-group14-sha1
switch(config)# ssh host-key-algorithms ecdsa-sha2-nistp256, ecdsa-sha2-
nistp384,ecdsa-sha2-nistp521
switch(config)# ssh public-key-algorithms ecdsa-sha2-nistp256, ecdsa-sha2-
nistp384, ecdsa-sha2-nistp521
```

Individual algorithms are ordered and advertised to the peer SSH device as configured. Please order the algorithms appropriately to ensure that desired preference of algorithms

Server Port Customization

By default, SSH server listens on TCP port 22. This port will be used for all VRFs that have SSH server enabled. Optionally AOS-CX switches provides the ability to modify the default SSH server port to add extra protection to the server. Supported Port number range from 1 to 65535. Although it is possible to use all ports, it might cause a network conflict. Thus, it is safer to choose a port number which is not

reserved for any other service. Additionally ensure the firewall is not blocking the port you want to use for SSH.

Sample configuration to modify the SSH server port :

```
switch(config)# ssh server port 19222
```



This port will be used for all VRFs that have an SSH server enabled. If the new port is currently opened by another service on a VRF, the SSH server will go into an error state for that VRF, and an event log message will be logged.

Two Factor Authentication and Authorization

Two factor Authentication is an extra layer of protection used to guarantee the secure access of switch management interfaces like SSH and HTTPS-Server. In two-factor authentication, X.509 certificate-based authentication is combined with RadSec authentication. Two-factor authentication can be performed locally or remote RadSec server. Refer security guide for detailed steps.

Following table summarizes the Two factor Authentication and Authorization support across different authentication methods – Local and Remote -

Task/Methods	Local Only	Local + Remote	Remote Only
Supported Management Interfaces	SSH	SSH	SSH and HTTPS-Server
X.509 Certificate Authentication	Validated using locally configured TA profile in switch	Validated using locally configured TA profile in switch	Validated using locally configured TA profile in switch
Validation of Username present in certificate's Common Name or Subject Alternative Name - User Principal Name	Local user Accounts	Local user Accounts	Remote RADIUS Authorize only request
Validation of - Username and Password	Local user Accounts	Remote Server	No Validation.
Authorization	Local user Account	Remote Server	Remote Server

- Authorization requests are sent over TLS and therefore RADIUS authorize-only requires a RadSec RADIUS server. It should be supporting **Authorize** only request.
- For Remote only authentication , password is not required at the time of authentication.
- Your switch management computer has access to the REST API using an appropriate HTTPS client. This can be done with a web browser, using the WebUI, or other HTTPs request tools such as

Postman. Usage of Firefox is not recommended, as it requires additional configuration to work with this feature.

Summary

The following table summarizes the management access methods available on an AOS-CX Switch, how they are secured by default, and the ways in which they can be secured.

Access Method	Secured by Default	Ways to Secure	Other Hardening Recommendations
Console	No	Enable AAA through External TACACS+/RADIUS/RadSec server or Local (Mandatory Fallback)	• Limiting Shell Access • Session Management
Telnet	No	Enable AAA through External TACACS+/RADIUS/RadSec server or Local	• Limiting Shell Access • Session Management • Control Plane ACLs
SSH	No	Enable AAA through External TACACS+/RADIUS/RadSec server or Local	• Limiting Shell Access • Session Management • Hardening SSH • Control Plane ACLs
Web UI	No	Enable Authentication and Accounting through External TACACS+/RADIUS/RadSec server or Local. Authorization is supported via RBAC.	• Hardening PKI • TLS Enforcements • Control Plane ACLs
REST API	No	Enable Authentication and Accounting through External TACACS+/RADIUS/RadSec server or Local. Authorization is supported via RBAC	Control Plane ACLs
SNMP	No	Refer to Securing SNMP Access	Control Plane ACLs

Session Management

Session management enhances security by enforcing specific CLI user session requirements for console, SSH and telnet connections. The following information is provided at the time of a successful login:

- When applicable, the number of failed login attempts since the most recent successful login.
- The date, time, and location (console or IP address or hostname) of the most recent previous successful login.
- The count of successful logins within the past (configurable) time period.

The following example configures CLI user session settings for a maximum of one concurrent session with a 15-minute timeout, and tracking for a maximum of 25 days

```
switch(config)# cli-session
switch(config-cli-session)# max-per-user 1
switch(config-cli-session)# timeout 15
switch(config-cli-session)# tracking-range 25
switch# exit
```

It is recommended to configure at least five to ten minutes of timeout for sensitive networks. For non-sensitive networks, a 15 minute timeout is recommended.

When the same user name is used for both local and remote authentication, both users, regardless of privilege level, are considered to be the same user for the purpose of counting concurrent CLI sessions. For example, with **max-per-user** value set to **1** and user **admin1** configured for local and remote authentication, only the local user **admin1** or the remote user **admin1** can be logged in at any given moment. Both admin1 users cannot be logged in simultaneously unless the **max-per-user** value is increased to at least **2**.

Limiting Shell Access

The AOS-CX operating system provides access to the underlying Linux system, allowing administrators to launch a bash shell session from the switch command-line interface. Misuse of shell access could expose sensitive network traffic to an unauthorized third party via packet mirroring to a remote device or could cause a denial of service by modifying or removing system files. This file modification could render the device unbootable, and require software restoration through the ServiceOS console..

The following are best practices for limiting shell access:

- Disable access to the Bash shell by changing the switch security mode to [enhanced from ServiceOS](#).
- Limit shell access by using RBAC or an external TACACS+ authorization server to deny access to the start-shell command to all users except those who specifically require it.

Securing SNMP Access

SNMP is used to manage and monitor networked devices from a centralized platform. There are three versions of the SNMP protocol: v1, v2c, and v3. SNMPv1 and v2c use community names for read and write access. Much like passwords are used for authentication, these community names are sent across the wire as clear text. If a malicious user were to capture these community names, they could potentially issue SNMP **set** commands to make unauthorized and potentially harmful configuration changes to a network device. SNMPv3, by comparison, utilizes a user-based security model with both authentication and privacy protocols to prevent unauthorized access or eavesdropping of management traffic.

SNMP is disabled by default on all AOS-CX devices. When enabled, SNMP provides limited write support in addition to read-only access and trap support for SNMP v1, v2c, and v3.

The default SNMP community string is public, a common setting for SNMP-capable devices. Replace the public community string with another value that is hard to guess , but note that this doesn't fully prevent against attacks as this string is in clear text format in packet captures:

```
Switch(config)# snmp-server community zerotrust
```

The default access level for SNMP communities is read-only; if read-write support is required, set the access level for the community to **rw** from the community context. IPv4 and/or IPv6 ACLs may be used to limit access to allowed management stations or subnets; only one ACL (IPv4 or IPv6) may be applied to a community at a time. Apply an IPv4 or IPv6 ACL from the SNMP **config-community** context.

```
switch(config)# snmp-server community zerotrust
switch(config-community)# access-level rw
switch(config-community)# access-list snmp_acl
```

```
switch # show snmp community
```

Community	Access-level	ACL Name	ACL Type	View
zerotrust	rw	snmp_acl	ipv4	none

Best practices is to use SNMPv3 instead of older versions of SNMP. Older versions of SNMP are unauthenticated and unencrypted, with the community string acting as a password, transmitted in plaintext. SNMPv3, offers support for different users, authentication, and strong encryption. AOS-CX supports stronger authentication protocols (SHA224, SHA256, SHA384, and SHA 512) and privacy protocols (AES192 and AES256).

To create an SNMPv3 user using SHA for authentication and DES for privacy:

```
switch(config)# snmpv3 user myUser auth sha auth-pass plaintext myAuthPswrd priv
des priv-pass plaintext myPrivPswrd
```

The following example creates an SNMPv3 context with the community name created above and assigned to the mgmt VRF:

```
switch(config)# snmpv3 context snmpv3mgmt vrf mgmt community zerotrust
```

Disable support for SNMPv1 and SNMPv2c and only accept SNMPv3 messages using the following command:

```
switch(config)# snmp-server snmpv3-only
```

To enable SNMP on the **mgmt** VRF:

```
switch(config)# snmp-server vrf mgmt
```

```
switch# show snmpv3 context
```

Name	vrf	Community	Type[Instance_id]
snmpv3mgmt	mgmt	zerotrust	vrf

```
switch# show snmpv3 users
```

User	AuthMode	PrivMode	Status	Context	Access-level	View
------	----------	----------	--------	---------	--------------	------

myUser	sha	des	Enabled	none	ro	none
--------	-----	-----	---------	------	----	------

Control Plane ACLs

Once an IP address is bound to an interface associated with a VRF, the switch may become exposed to management access from untrusted users or devices. This potential point of vulnerability can be mitigated by binding an Access Control List (ACL) to the control plane for that VRF. The control plane handles the device's management and routing functionality.

Once a control plane ACL is applied to a VRF, it filters packets to all IPv4/IPv6 addresses bound to the device on that VRF. It is possible to create a control plane ACL for each existing VRF, including the **mgmt** VRF.

The following commands are an example of an ACL an administrator can apply that limits SSH and SNMP control plane access to source devices with IP addresses in the 10.10.0.0/24 subnet, with counters for denied SSH and SNMP packets.

```
switch(config)# access-list ip CONTROLPLANE
switch(config-acl-ip)# 05 comment ALLOW SSH AND SNMP ON ADMIN SUBNET, BLOCK ALL OTHERS
switch(config-acl-ip)# 10 permit tcp 10.10.0.0/24 any eq 22
switch(config-acl-ip)# 20 permit udp 10.10.0.0/24 any eq 161
switch(config-acl-ip)# 30 permit udp 10.10.0.0/24 any eq 162
switch(config-acl-ip)# 40 deny tcp any any eq 22 log count
switch(config-acl-ip)# 50 deny udp any any eq 161 log count
switch(config-acl-ip)# 60 deny udp any any eq 162 log count
switch(config-acl-ip)# 990 comment ALLOW ANYTHING ELSE
switch(config-acl-ip)# 1000 permit any any any
```



Event logs for Control Plane ACE is supported using the **log** keyword. This option offers better troubleshooting and visibility of an ACL applied to the control plane.

To apply this ACL to the default VRF:

```
switch(config)# apply access-list ip CONTROLPLANE control-plane vrf default
```

All ACLs in AOS-CX have an implicit **deny any** rule at the end of the rules list. This requires that allowed traffic be explicitly permitted to pass through an applied ACL. In the above example, SSH and SNMP traffic on ports 22 is allowed from 10.10.0.0/24. The SSH and SNMP traffic is then blocked from any other subnets. The final ACL entry (**permit any any any**) permits all other traffic.

Time Synchronization

Many secure protocols and auditing functions rely on system times being synchronized with a reliable time source, either within or (where security considerations permit) external to the managed network. One of the most commonly-used protocols to accomplish this is the Network Time Protocol (NTP), which can use both local and Internet-hosted servers to synchronize system time across a network. Recommendation - NTP should be configured and enabled on the device prior to enabling secure management protocols.

A common practice among organizations that span multiple time zones is to use NTP to synchronize time clocks and set the local time zone on all equipment to UTC. This practice aids in troubleshooting and security audits for devices that might be continents apart. Both IPv4 and IPv6 Servers are supported.

To configure a switch to use NTP authentication and connect to a local NTP server at 10.100.1.254 using the switch management port:

```
switch(config)# ntp authentication
switch(config)# ntp authentication-key 1 md5 ntpauthkey
switch(config)# ntp server 10.100.1.254 prefer
switch(config)# ntp vrf mgmt
```

```
switch# show ntp servers
```

NTP SERVER	KEYID	MINPOLL	MAXPOLL	OPTION	VER
10.100.1.254	--	6	10	none	4 prefer
10.80.2.219	--	6	10	iburst	4 prefer (auto)
pool.ntp.org	--	4	4	iburst	4

```
switch# show ntp authentication-keys
```

Key ID	Trusted	Type	Encrypted Key
1	No	MD5	AQBapUtt1YTjZS2PH4+J7G5OKJG0GuZ2WxmD0339TNg6nfGXY=

```
switch# show ntp status
```

```
NTP Status Information
NTP : Enabled
NTP DHCP : Enabled
NTP Authentication : Enabled
NTP Server Connections : Using the mgmt VRF
System time : Fri Mar 8 03:51:46 PST 2024
NTP uptime : 8 days, 15 hours, 24 minutes, 37 seconds
Not synchronized with an NTP server.
```

Secure Copy

The **copy** command is widely used in AOS-CX switches to transfer files, configurations and log messages. The commonly used file transfer protocol TFTP transfers files in plaintext, so attackers can easily capture transferred packets. To protect the device against security threats, it is recommended to use **SFTP** and **SCP** to perform the copy operations.

Hardening PKI

The Public Key Infrastructure (PKI) feature enables administrators to manage digital certificates on the switch. The switch uses certificates to validate clients when acting as a server, and when communicating with servers when TLS encryption is used.

The AOS-CX Switch Series supports the installation of certificate authority (CA) certificates and the generation and installation of leaf certificates. The switch supports 64 trust anchor (TA) profiles. Each TA

profile stores a trusted CA certificate. The certificate can be either a root CA certificate, which must be self-signed, or an intermediate CA certificate that is issued by another CA. The TA profile also enables configuration of real-time checking of certificate revocation (through OCSP).

Leaf certificates can be installed on the switch for use by applications such as:

- RadSec Client
- dot1x-suppliant
- EST Client
- captive-portal
- syslog client
- https-server - Web UI or REST API

AOS-CX switches by default supports the following preinstalled leaf certificates:

- **local-cert:** A self-signed certificate that switch automatically generated at first boot, as the default certificate for any application when the application's associated certificate is not configured
- **device-identity:** A device-identity certificate built into a switch at manufacturing and resident for the life of the product. The *identity* is a combination of an RSA key pair with physical information such as the unit's model, chassis/PCA serial number, and base MAC ID. Device Identity will be used for following purposes:
 - Allow 801.2X-2010 to perform peer authentication without the need for certificate or pre-shared key installation to automate the formation of MACsec secure channels between neighbor devices.
 - Authentication with HPE Aruba Networking cloud services.

```
switch # show crypto pki certificate
Certificate Name      Cert Status      EST Status      Associated Applications
-----
local-cert           installed        n/a             captive-portal, dot1x
suppliant, est-client, https-server, radsec-client, syslog-client
device-identity      installed        n/a             none
```



AOS-CX recommends the usage of trusted CA signed certificate over the self-signed certificate for all the applications to avoid potential security risks.

If you are purchasing a certificate from a trusted CA, the switch can generate the certificate signing request (CSR) that is used to request the certificate. The switch can also directly generate self-signed certificates. Alternatively, the certificate and private key can be generated outside the switch and then imported. X509 certificate management software such as OpenSSL can be used to generate the private key and CSR and then combine the certificate and private key into one PEM or PKCS#12 file suitable for import into the switch.

The following procedure describes how to create and install an X.509 leaf certificate that is initiated inside the switch but signed outside the switch by a Certificate Authority.

```
switch(config)# crypto pki ta-profile root-cert
switch(config-ta-root-cert)# revocation-check ocsp
switch(config-ta-root-cert)# ocsp url primary http://ocsp-server.site.com
switch(config-ta-root-cert)# ocsp url secondary http://ocsp-server2.site.com
switch(config-ta-root-cert)# ta-certificate import terminal
Paste the certificate in PEM format below, then hit enter and ctrl-D:
switch(config-ta-cert)# -----BEGIN CERTIFICATE-----
```

```

switch(config-ta-cert)# MIIDuTCCAqECCQCuoxeJ2ZNYcjANBgkqhkiG9w0BAQsFADCBqzELMAEBh
switch(config-ta-cert)# VVMxEzARBgNVBAGMCKNhbgGlm3JuaWExEDAOBgNVBAcMB1JvY2tsDAKBg
switch(config-ta-cert)# BAoMA0hQTjEVMBMGA1UECwwMSFBOUm9zZXZpbGx1MSowKAYDVQocG5zdz
switch(config-ta-cert)# x3WFF3dFZ8o9sd5LVAHneH/ztb9MP34z+le1V346r12L2kpxmTOVJVyTO
switch(config-ta-cert)# BIzD/ST/HaWI+OS+S80rm93PSscEbb9GWk7vshh5EnW/moehBKCE40lzy
switch(config-ta-cert)# 3LvMLZcSSse5J2Ca2XIhfDme8UaNZ7syGYMsAW0nG7yYHWkEOQu9s
switch(config-ta-cert)# -----END CERTIFICATE-----
switch(config-ta-cert)#
The certificate you are importing has the following attributes:
Issuer: C=US, ST=CA, L=Rocklin, O=Company, OU=Site,
CN=site.com/emailAddress=test.ca@site.com
Subject: C=US, ST=CA, L=Rocklin, O=Company, OU=Site,
CN=8400/emailAddress=test.ca@site.com
Serial Number: 12121221634631568498 (0xae51217d5945772)
TA certificate import is allowed only once for a TA profile
Do you want to accept this certificate (y/n)? y
TA certificate accepted.
switch(config-ta-root-cert)# exit

```

```

switch(config)# crypto pki certificate lcert
switch(config-cert-lcert)# subject common-name Leaf country US state CA
locality Rocklin org Company org-unit Site
switch(config-cert-lcert)# key-type rsa key-size 3072
switch(config-cert-lcert)# enroll terminal
You are enrolling a certificate with the following attributes:
Subject: C=US, ST=CA, L=Rocklin, O=Company, OU=Site
CN=Leaf
Key Type: RSA (2048)
Continue (y/n)? y
-----BEGIN CERTIFICATE REQUEST-----
MIIBozCCAQAwYzEVMBMGA1UEAxMMcG9kMDEtODQwMCM0xMQ4wDAYDV
nViYTEMMAoGA1UEChMDSFBFMRlWEAYDVQQHEWlSb3Nldm1sbGUxYzAjbG
NBMQswCQYDVQQGEWJVUzCBnzANBgkqhkiG9w0BAQEFAAOBjQAwgYkCgYE
...
GBAJ4L3LFFFWBEL+KAKpOGJZcVmw1BMqSKFtOFNF9nzmUmONmU3SKy6dz
7Au22mf3lWDxztCC/dj5RtWJeJekxp2LCIK/3eRXUwbYveQDKcxH7j9Z
ace+2tA68F2vlgRCQ/hcQH0YmNuaq4Ne3w0dhm7HlUrx
-----END CERTIFICATE REQUEST-----

```

```

switch(config-cert-lcert)# import terminal ta-profile root-cert
Paste the certificate in PEM format below, then hit enter and ctrl-D:
switch(config-cert-import)# -----BEGIN CERTIFICATE-----
switch(config-cert-import)# MIIFRDCCAYygwIBAgIQPnnS2Vp5u07XMdktDJzANBgkqhkiG9w0Bv
switch(config-cert-import)# MQswCQYDVQGEWJVEOMAwG1UECgwFJ1YmxDAOgNBMMB1Jvb3QgQ0Ew
switch(config-cert-import)# HhcNMTkNDEwMjIwNTWcjIwMTA0MjwNE1WBzQswCQYDVQQGEWJVUzEL
...
switch(config-cert-import)# 1fIYZYGQyla0AwFuTTxBXYwRxPbUYU5tumrfwRPmE40VY8S9DQgcr
switch(config-cert-import)# 1NGNm3NG03GqPcs/T9bVyF5BOrS5lmm7kNfRYl8D/kMTfRreSdxis
switch(config-cert-import)# YQ1ulNqShps=
switch(config-cert-import)# -----END CERTIFICATE-----
switch(config-cert-import)#
Leaf certificate is validated with root-cert and imported successfully.
switch(config-cert-lcert)# exit
switch(config)# crypto pki application syslog-client certificate lcert

```

Mandatory matching of peer device hostname

To enhance the server-side certificate verification, the AOS-CX switch checks that the peer device configured hostname matches either the Subject Alternative Name (SAN) field or the Common Name

(CN) within the certificate Subject field. If the SAN field is present and matches the hostname, validation succeeds, otherwise it fails. If the SAN field is not present, and the CN matches the hostname, validation succeeds, otherwise it fails.

EST

EST stands for Enrolment over Secure Transport; An EST client is implemented as a part of the PKI infrastructure in the AOS-CX switches. Switches can be configured to request the trusted CA certificates and to request enrolment/reenrolment of leaf certificates automatically, without the need for administrator intervention, while maintaining the security and integrity of the whole enrolment process. Refer the **PKI EST** section in Security Guide for more information.

TLS Enforcements

Minimum TLS version supported in AOS-CX switches is TLS1.2. The following are recommended cipher suites for TLS Applications/Protocols

- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384

The Extended Key Usage X.509 v3 extension defines one or more purposes for which the public key can be used. This is in addition to or in place of the basic purposes specified by the Key Usage extension. As per NDcPP recommendation, that a peer certificate being used to establish TLS connection must have its extended key usage field set as client-auth or server-auth, depending on its role of the peer device. This configuration enables the checking of key usage during TLS handshake. It is disabled by default

```
switch(config)# tls check-key-usage
```

```
switch# show tls
TLS crypto algorithms state: default
TLS key usage checking      : on
```

Secure Logging

AOS-CX Switch provides both locally stored event and security logs, as well as using the syslog protocol to forward events to a remote IPv4/IPv6 syslog server for auditing purposes. Logged events can be filtered by severity level, originating system modules, or using regular expressions to match against message text.

When configuring AOS-CX to send logs to a remote server, it is common practice to set a facility value. This value acts as a label that the remote server can use to determine which file the syslog message should get appended.

Below is an example of how to configure AOS-CX to send event log messages via syslog to a remote server. This example uses the default facility of local7 and sends event messages marked informational and higher:

```
switch(config)# logging 10.100.1.250 vrf mgmt
```


To include security-related accounting logs in addition to the event logs, then add the include-auditable-events option to the configuration:

```
switch(config)# logging 10.100.1.250 include-auditable-events vrf mgmt
```

The syslog client can connect to a server using UDP (default), TCP, or TLS protocols. TLS is the recommended protocol, as it provides an encrypted connection to the syslog receiver. This requires the switch to possess a signed TLS client certificate, and the receiver to possess a signed TLS server certificate. The process of requesting and installing a signed TLS client certificate for syslog is similar to that for requesting and installing an SSL/TLS certificate for web-management.

Hardening the Control Plane

The following sections describe strategies for securing and hardening the switch control plane

Control Plane Policing

Control Plane Policing prevents flooding of certain types of packets from overloading the switch or module CPU by either rate-limiting or dropping packets.

The switch software provides several configurable classes of packets that can be rate-limited, including (but not limited to) ARP broadcasts, multicast, routing protocols (BGP, OSPF), and spanning tree. CoPP is always active and cannot be disabled.

The following default CoPP policy applies the following traffic class definitions and rate limits (in packets per second) on 6300 series switch series:

```
switch# show copp-policy default
```

class	drop	priority	rate pps	burst pkts	hardware rate pps
acl-logging	0	25	25	25	
arp-broadcast	2	1250	1250	1250	
arp-protect	2	2075	2075	2075	
arp-unicast	3	825	825	825	
bfd-control	5	850	850	850	
bgp	5	750	750	750	
captive-portal	2	2075	259	2075	
client-onboard	5	1024	1024	1000	
dfp-collector	0	512	512	500	
dhcp	2	750	750	750	
erps	6	225	225	225	
fib-optimization	0	100	200	100	
icmp-broadcast-ipv4	2	325	325	325	
icmp-multicast-ipv6	2	475	475	475	
icmp-security-ipv6	2	325	325	325	
icmp-unicast-ipv4	3	225	225	225	
icmp-unicast-ipv6	3	400	400	400	
ieee-8021x	2	2075	259	2075	
igmp	4	1600	450	1600	
ip-exceptions	0	100	100	100	
ip-lockdown	0	100	100	100	
ip-tracker	0	256	256	250	
ipfix	0	2500	2500	2500	
ipsec	5	1025	128	1025	
ipv4-options	1	100	100	100	
lACP	5	2050	2050	2050	

lldp	5	100	100	100
loop-protect	6	225	225	225
mac-lockout	0	100	100	100
manageability	4	4218	4218	4200
mdns	2	150	150	150
mirror-to-cpu	0	100	100	100
mld	4	1600	450	1600
mvrp	5	225	225	225
nae-packet-monitor	0	100	200	0
ntp	4	100	100	100
ospf-multicast	5	1025	1025	1025
ospf-unicast	5	1025	1025	1025
pim	5	1700	1700	1700
ptp	5	1000	250	1000
secure-learn	2	2075	259	2075
sflow	1	1000	125	1000
stp	6	2000	2000	2000
udld	6	450	450	450
unknown-multicast	1	1025	128	1025
unresolved-ip-unicast	1	325	325	325
vrrp	4	400	400	400
default	2	4225	528	4225

The default CoPP policy can be modified but cannot be deleted.

To revert a modified default CoPP policy to factory default settings:

```
switch(config)# copp-policy default revert
```

Administrators may create up to 32 custom CoPP policies, though only one can be active at any given time.

To create and apply a simple custom CoPP policy:

```
switch(config)# copp-policy copp_policy_01
switch(config-copp)# class arp-broadcast priority 2 rate 1000 burst 1000
switch(config-copp)# class unknown-multicast priority 2 rate 1000 burst 1000
switch(config-copp)# class unresolved-ip-unicast priority 2 rate 1000 burst 1000
switch(config-copp)# default-class priority 1 rate 3000 burst 3000
switch(config-copp)# exit
switch(config)# apply copp-policy copp_policy_01
```

To remove a custom CoPP policy from service and automatically apply the default policy

```
switch(config)# no apply copp-policy copp_policy_01
```

To delete a custom CoPP policy:

```
switch(config)# no copp-policy copp_policy_01
```

An active custom CoPP policy cannot be deleted; it must first be removed from service using the above command.

Securing Spanning Tree

The following sections describe security and hardening workflows for Spanning Tree.

BPDU Protection

Various security mechanisms are in place to protect spanning tree configurations from interference and rogue devices or unwarranted changes to the network. BPDU protection secures the active topology by preventing spoofed BPDU packets from entering the network. Typically, BPDU protection is applied on edge ports connected to end user devices that do not run STP. If STP BPDU packets are received on a protected port, BPDU guard disables the port and an alert is sent. Hence recommended to enable BPDU guard on end user/device connected ports to prevent any inadvertent spanning tree or malicious attack.

```
switch(config)# interface 1/1/8
switch(config-if)# no shutdown
switch(config-if)# no routing
switch(config-if)# vlan access 10
switch(config-if)# spanning-tree bpduguard
switch(config-if)# exit
```

Root Protection

Root protection secures the active topology by preventing other switches from declaring their ability to propagate superior BPDUs, containing both better information on the root bridge and path cost to the root bridge which would normally replace the current root bridge selection. This is typically carried out between the core that is required to be the root and access switches to prevent ports that are not expected to originate root information such as server ports and access switch ports.

```
switch(config)# interface 1/1/8
switch(config-if)# no shutdown
switch(config-if)# no routing
switch(config-if)# vlan access 10
switch(config-if)# spanning-tree rootguard
switch(config-if)# exit
```

Viewing the configuration change:

```
switch# show spanning-tree interface 1/1/3
Port : 1/1/3
Admin State : up
BPDU Guard : enabled
BPDU Filter : disabled
RPVST Guard : disabled
RPVST Filter : disabled
Loop Guard : disabled
Root Guard : enabled
TCN Guard : disabled
Admin Edge Port : admin-network
Link Type : Point to Point
BPDU Tx Count : 31
BPDU Rx Count : 0
TCN Tx Count : 0
TCN Rx Count : 0
```

DHCP Security

The following sections describe security and hardening workflows for DHCP.

DHCP Snooping

DHCP snooping protects the network from common DHCP attacks, including address spoofing resulting from a rogue DHCP server operating on the network or exhaustion of addresses on a DHCP server caused by mass address requests by an attacker on the network. DHCP snooping designates trusted DHCP servers and ports on which DHCP requests and responses are accepted.



Refer to the IP Services Guide for more information.

The following is a DHCPv4-snooping sample configuration:

```
switch(config)# dhcpv4-snooping
switch(config)# vlan 100
switch(config-vlan-100)# dhcpv4-snooping
switch(config-vlan-100)# exit
switch(config)#
switch(config)# interface 1/1/1
switch(config-if)# dhcpv4-snooping trust
switch(config-if)# exit
switch(config)# dhcpv4-snooping authorized-server 192.168.2.10 vrf default
```

To view the configuration change with DHCPv4-snooping:

```
switch# show dhcpv4-snooping

DHCPv4-Snooping Information
  DHCPv4-Snooping           : Yes           Verify MAC Address : Yes
  Allow Overwrite Binding   : No             Enabled VLANs      : 100
  IP Binding Disabled VLANs :
  Static Attributes         : No
  Client Event Logs         : No
  Trust VxLAN Tunnels       : Yes

Option 82 Configurations
  Untrusted Policy          : drop            Insertion           : Yes
  Option 82 Remote-id       : mac

External Storage Information
  Volume Name               : --
  File Name                  : --
  Inactive Since             : --
  Error                      : --

Flash Storage Information
  File Write Delay           : --

Active Storage : --

Authorized Server Configurations
  VRF                        Authorized Servers
  -----
  default                    192.168.2.10

Port Information
  Max      Static      Dynamic
```

Port	Trust	Bindings	Bindings	Bindings
-----	-----	-----	-----	-----
1/1/1	Yes	0	0	0

The following is a DHCPv6-snooping sample configuration:

```
switch(config)# dhcpv6-snooping
switch(config)# vlan 100
switch(config-vlan-100)# dhcpv6-snooping
switch(config-vlan-100)# exit
switch(config)#
switch(config)# interface 1/1/1
switch(config-if)# dhcpv6-snooping trust
switch(config-if)# exit
switch(config)# dhcpv6-snooping authorized-server ABCD:5ACD::2000 vrf default
```

To view the configuration change with DHCPv6-snooping:

```
6200(config)# show dhcpv6-snooping

DHCPv6-Snooping Information

    DHCPv6-Snooping      : Yes    Enabled VLANs      : 100
    IP Binding Disabled VLANs :
    Trusted Port Bindings Enabled VLANs :
    Client Event Logs      : No
    Trust VxLAN Tunnels    : Yes

External Storage Information

    Volume Name          : --
    File Name             : --
    Inactive Since        : --
    Error                  : --

Flash Storage Information
    File Write Delay : --

Active Storage : --

Authorized Server Configurations
    VRF                Authorized Servers
    -----
    default             ABCD:5ACD::2000

Port Information

    Port      Trust  Max   Static  Dynamic
    -----  -----  -----  -----  -----
    1/1/1     Yes    0      0        0
```

DHCPv6 Guard

DHCPv6 guard is an extension of DHCPv6 snooping. When the DHCPv6 snooping feature is configured globally and on the VLAN, the ports are configured as trusted and untrusted ports. DHCPv6 guard enhances this by creating a policy and applying it on a port and on the VLAN. This policy contains multiple attributes which are compared against the packet that is received on trusted ports. If the

packet complies with the attributes of the policy, it is forwarded to the destination port; otherwise the packet is dropped.

The following are sample configurations of DHCPv6 guard:

```
switch(config)# dhcpv6-snooping guard-policy poll
switch(config-dhcpv6-guard-policy)# match server access-list acl1
switch(config-dhcpv6-guard-policy)# preference min 6
switch(config-dhcpv6-guard-policy)# preference max 250
switch(config-dhcpv6-guard-policy)# match client prefix-list pref1
```

```
switch(config)# vlan 5
switch(config-vlan-100)# dhcpv6-snooping guard-policy poll
```

To view the configuration change with DHCPv6 guard:

```
switch# show dhcpv6-snooping guard-policy
DHCPv6-Snooping guard-policy Information
DHCPV6 Guard Policy name : POL1
Attached Access List : ACL1
Attached Prefix List : PRF1
Preference Range : 6-250
Applied on VLAN : 5
Applied on Port
```

Dynamic ARP Inspection

Dynamic ARP Inspection provides additional security for ARP. Dynamic ARP resolves IP addresses against MAC addresses on a broadcast network segment such as Ethernet, originally defined by Internet Standard RFC 826. ARP does not support any inherent security mechanism and as such, depends on simple datagram exchanges for the resolution, with many of these being broadcast. Because it is an unreliable and non-secure protocol, ARP is vulnerable to attacks.

Some attacks may be targeted toward the networks whereas other attacks may be targeted toward the switch itself. The attacks primarily intend to create denial of service (DoS) for the other entities present in the network. Most of the attacks are carried out in one of the following three forms:

- Overwhelming the switch control plane with too many ARP packets.
- Overwhelming the switch control plane with too many unresolved data packets.
- Posing as a trusted gateway/server by wrongly advertising ARPs.

The following defense mechanisms can be put in place on a switch to protect against attacks:

- Limiting the amount of ARP activity allowed from a host or on a port.
- Ensuring that all ARP packets are consistent with one or more binding databases.
- Enforcing integrity checks on the ARP packets to check against different MAC or IP addresses in the Ethernet or IP and ARP header.

The following are supported:

- Enabling and disabling of Dynamic ARP Inspection on a VLAN level (it does not have to be SVI).
- Defining the member ports of a VLAN as either trusted or untrusted. Only ARP traffic on untrusted ports subjected to checks.

- Routed ports (RoPs) always treated as trusted.
- Listening to the DHCP Bindings table and checking every ARP packet to match against the binding.

Prerequisites

Dynamic ARP Inspection is enforced using **DHCP Snooping binding** and **Static IP Binding**. Refer to the [DHCP Snooping](#) section for the DHCP snooping configuration to enable the **Static IP Binding**.

```
switch(config)# interface vlan 10
switch(config-if-vlan)# arp ipv4 2.2.2.2 mac 01:00:5e:00:00:01
```

```
switch# configure terminal
switch(config)# vlan 1
switch(config-vlan)# arp inspection
```

To configure the interface as trusted:

```
switch# configure terminal
switch(config)# interface 1/1/1
switch(config-if)# arp inspection trust
```

```
switch# show arp inspection interface 1/1/1
```

```
-----
Interface Trust-State
```

```
-----
1/1/1 Trusted
-----
```



All interfaces are untrusted by default.

ND Snooping Attack Prevention

ND snooping is used in Layer 2 switching networks and prevents ND attacks. ND snooping drops invalid ND packets, and combined with DIPLDv6 (Dynamic IP Lockdown for IPv6), blocks data traffic from invalid hosts. ND snooping learns the source MAC addresses, source IPv6 addresses, input interfaces, and VLANs of incoming ND messages and data packets to build IP binding entries.

ND snooping drops ND packets for the following reasons:

- If the Ethernet source MAC address does not match the address in the ICMPv6 Target link layer address field of the ND packet.
- If the global IPv6 address in the source address field does not match the ND snooping prefix filter table.
- If the global IPv6 address or the link-local IPv6 address in the source IP address field does not match the ND snooping binding table

The following are sample VLAN configurations of globally enabled ND snooping:

```
switch# configure terminal
switch(config)# nd-snooping
switch(config)# vlan 100
switch(config-vlan-100)# nd-snooping
switch(config-vlan-100)# exit
switch(config)#
```

```
switch(config)# show nd-snooping
ND Snooping Information
=====
ND Snooping : Enabled
ND Snooping Enabled VLANs : 100
Trusted Port Bindings Enabled VLANs : 100
ND Guard Enabled VLANs : 100
RA Guard Enabled VLANs : 100
RA Drop Enabled VLANs :
MAC Address Check : Disabled
PORT TRUST MAX-BINDINGS CURRENT-BINDINGS
-----
1/1/1 Yes
1/1/2 Yes
```



For more information on ND snooping refer the **AOS-CX IP Services Guide**.

RA Guard

Router Advertisement (RA) guard blocks unwanted, forged RA messages on a Layer 2 access device. ND snooping drops both RA and RR packets on untrusted ports. To block only RA packets on VLANs with ND snooping enabled, use **nd-snooping ra-drop**.

RA drop is disabled by default on VLANs. When enabled (with **nd-snooping ra-drop**), ND snooping blocks RA packets on both trusted and untrusted ports. When RA drop is disabled, ND snooping allows RA packets on trusted ports and blocks them on untrusted ports.

When RA guard policy is enabled (with **ipv6 nd-snooping ra-guard policy**), RA packets received on trusted ports are validated against a set of parameters configured on the policy and assigned to a port or VLAN. RA Guard policy options include:

- Hop Limit
- Managed Config Flag
- Other Config Flag
- Router Preference
- ACL
- Advertised Prefix Lists

```
switch(config)# ipv6 nd-snooping ra-guard policy <POLICY-NAME>
switch(config-ra-guard-policy)#-----Policy Parameters-----
```

```
switch(config)# vlan 10
switch(config-vlan-10)# nd-snooping ra-guard attach-policy POLICY_NAME
```



```

switch# show nd-snooping ra-guard interface 1/1/1
RA Guard Policy Counters
=====
RA Guard Policy Applied : POLICY_2
RA Packets Received : 10
RA Packets Forwarded : 5
RA Packets Dropped : 5 [Total]
reason : Managed flag error [0]
Other flag error [0]
Access list error [0]
Prefix list error [0]
Router preference error[0]
Hop limit error [5]

```

IPv6 Destination Guard

Enabling IPv6 destination guard on a switch prevents ND cache depletion issues and helps in minimizing Denial-of-Service (DoS) attacks. When IPv6 destination guard is enabled, address resolution is performed only for the destination addresses that are active on the link. This feature requires the binding table to be populated with the help of DHCPv6 snooping, ND snooping, or static-ip-bindings.

Destination guard enables the destination address based filtering of IPv6 traffic and blocks the Neighbor Discovery (ND) protocol resolution for destination addresses that are not found in the binding table.

```

switch(config)# vlan 10
switch(config-vlan-10)# ipv6 destination-guard

```

```

switch# show ipv6 destination-guard statistics
Packets dropped for VLAN 10 : 25467
Packets dropped for VLAN 30 : 434
Packets dropped for VLAN 50 : 8767

```

IP Source Lockdown

IP source lockdown provides added security by preventing IP source address spoofing on a per-port basis. Every packet is inspected for this purpose in hardware. When IP source lockdown is enabled, IP traffic received on an interface (port) is forwarded only if the VLAN, IP address, MAC address, and interface (port) match the IP binding database entry.

To use IP source lockdown, the IP binding database must be populated. The binding database is dynamically populated by DHCP snooping that learns and saves the binding information. Alternatively, the IP binding database can be statically populated with the `ip source-binding` command.

To enable IP source lockdown resource extended on the device (supports dynamically sharing hardware resources of IP source lockdown with other features):

```

switch(config)# ip source-lockdown resource-extended
Do you want to continue (y/n)? y

```

To enable IPv4/IPv6 source lockdown for all VLANs on the selected interface (port):

```
switch(config)# interface 1/1/1
switch(config-if)# ipv4 source-lockdown
switch(config-if)# ipv6 source-lockdown
```

```
switch# show ipv4 source-lockdown
INTERFACE LOCKDOWN HW-STATUS
-----
1/1/1   Yes       Yes
```

```
switch# show ipv6 source-lockdown
INTERFACE LOCKDOWN HW-STATUS
-----
1/1/1   Yes       Yes
```

To add static IPv4/IPv6 client source binding information to the switch IPv4/IPv6 binding database:

```
Ipv4 source-binding <VLAN-ID> <IPv4-ADDR> <MAC-ADDR> <IFNAME>
ipv6 source-binding <VLAN-ID> <IPv6-ADDR> <MAC-ADDR> <IFNAME>
```

```
switch# show ipv4 source-binding
PORT    VLAN    MAC-ADDRESS    HW-STATUS    FROM    IPv4-ADDRESS
-----
1/1/1    2       aa:bb:cc:dd:ee:ff    Yes         static   1.2.3.4
1/1/2    12      aa:ab:cc:dd:ee:ff    Yes         static   10.20.30.40
```

```
switch# show ipv6 source-binding
PORT    VLAN    MAC-ADDRESS    HW-STATUS    FROM    IPv6-ADDRESS
-----
1/1/1    1234    00:50:56:96:e4:cf    Yes/No       static   3000::1
```

Securing Routing Protocols

The following sections describe the workflows for securing OSPF and BGP routing protocols.

OSPF Passive Interfaces

Unlike BGP, most routing protocols tend to discover neighbors via the sending and receiving Hello packets. Since these neighbor relationships build dynamically, the administrator should control which neighbor relationships can be formed and administrators should ensure that the potential neighbors are known and trusted.

To limit where OSPF can learn neighbors, AOS-CX supports the passive OSPF interfaces. A passive OSPF interface has its IP subnets announced, but it does not establish neighbor relationships with other OSPF devices on the interface.

You must make all OSPF enabled interfaces passive. Setting the OSPF enabled interfaces to from default to passive is done in the OSPF router instance context.

```
switch(config)# router ospf 1
switch(config-ospf-1)# passive-interface default
```

The passive interface is then removed from each interface where OSPF neighbor relationships are allowed. Since this is an interface-level configuration change, it can be done from the interface context:

```
switch(config)# interface 1/1/1
switch(config-if)# no ip ospf passive
```

OSPF Neighbor Authentication

All OSPF exchanges are authenticated. However, the default authentication used by network vendors is "null," meaning empty or zero. OSPF also supports use of a simple plaintext password and cryptographic authentication. AOS-CX supports several OSPFv2 authentication methods, including SHA cryptographic hashes up to 512 bits, to authenticate messages between OSPF neighbors.

When configuring authentication between OSPF neighbors, the authentication method and authentication key must be the same on the interfaces connected on the both devices.

To configure SHA-512 authentication, change the default authentication method from null to hmac-sha-512 from the interface context:switch

```
(config-if)# ip ospf authentication hmac-sha-512
```

Then configure a SHA key to be used for the connection, the key can be entered as plaintext or as a hashed ciphertext string:

```
switch(config-if)# ip ospf sha-key 1 plaintext ospfshakestring
```

Alternatively, the AOS-CX keychain feature may be used to specify a system-level cryptographic authentication key which can be used by multiple OSPF interfaces:

```
switch(config)# keychain ospf-keychain
switch(config-keychain)# key 1
switch(config-keychain-key)# cryptographic-algorithm hmac-sha-512
switch(config-keychain-key)# key-string plaintext ospfshakestring
switch(config-keychain-key)# interface 1/1/49
switch(config-if)# ip ospf authentication keychain
switch(config-if)# ip ospf keychain ospf-keychain
```

OSPFv3 Area Authentication and Encryption with IPsec

OSPFv3 neighbors may use interface-level authentication. An alternative method might be used to provide encryption, or authentication, or both for an entire OSPFv3 area using the IPsec protocol, which automatically applies the configured methods to all member interfaces. There are two IPsec encapsulation types supported on AOS-CX to secure OSPFv3 areas:

- IPv6 authentication header (AH), which adds an IPv6 authentication header to OSPFv3 packets.
- Encrypted Security Payload (ESP), which provides both authentication and encryption for OSPFv3 packets.

IPsec authentication and encryption are configured from the OSPFv3 router process context. Both authentication and encryption require a specified Security Policy Index (SPI), which is an integer value between 256 and 4,294,967,295; this value is used on each OSPFv3 router in the secured area to match a configured IPsec authentication and/or encryption policy. Each OSPFv3 IPsec policy on a switch must use a different SPI value, and the SPI value (as well as authentication, or encryption keys, or both) must match across all OSPFv3 neighbor interfaces using that policy within the secured area.

To configure AH authentication for OSPFv3 area 1, specify the SPI, authentication method (md5 or sha1), key type (plaintext, hex-string, or ciphertext) and the key string itself. If a key type and string are not specified, the user is prompted to enter a plaintext key interactively:

```
switch(config-ospfv3-1)# area 1 authentication ipsec spi 1024 sha1
Enter the IPsec authentication key: *****
Re-Enter the IPsec authentication key: *****
```

To configure ESP encryption for area 1, specify the SPI, authentication method, authentication key type and string, encryption type (3des, aes, des, or null), key type, and encryption key string. If the encryption type and key string are not specified, you are prompted to enter a plaintext key interactively. If the authentication key type and string are not specified, you are prompted to enter both a plaintext authentication key as well as the desired encryption type and plaintext key.

```
switch(config-ospfv3-1)# area 1 encryption ipsec spi 1024 sha1
Enter the IPsec authentication key: *****
Re-Enter the IPsec authentication key: *****
Enter the IPsec encryption type (3des/aes/des/null)? aes
Enter the IPsec encryption key: *****
Re-Enter the IPsec encryption key: *****
```

Depending on the selected encryption type, a plaintext or hexadecimal encryption key must be set to a specific length as mentioned below:

- 3DES:
 - Hexadecimal: 48 digits
 - Plaintext: 24 characters
- DES:
 - Hexadecimal: 16 digits
 - Plaintext: 8 characters
- AES:
 - Hexadecimal: 32, 48, or 64 digits
 - Plaintext: 16, 24, or 32 characters

For AES encryption, the specified key lengths correspond to AES128, AES192, or AES256, respectively; the type of the key that will be used is automatically determined by the length of the entered encryption key. AOS-CX recommends using AES over DES or 3DES as it is stronger.

```
switch(config)# show run ospf
switch(config)# show run ospfv3
```

BGP

The IETF Best Current Practices for BGP Security (BCP194) focuses on the following three items:

- Utilizing the control-plane ACL functionality to limit BGP communication to configured BGP peers.
- Securing BGP sessions between peers by using authentication.
- Use TTL Security Mechanisms to prevent spoofing attacks from third parties.

Control Plane ACL for BGP Peering Sessions

Devices running BGP listen for connections on TCP port 179. When establishing a BGP peer session, one device actively establishes a relationship with the other peer by sending the first TCP SYN packet. This device is at the outgoing side of the connection. The other peer, hearing the TCP SYN, responds with a SYN or ACK at the incoming connection. As each peer can assume either role, ACL entries need to be configured for BGP in both directions.

Building on the same Control Plane ACL example as before, the below entries permit traffic from **10.20.0.10** so that it can establish a BGP peering session with the device. Either side could play the outgoing or incoming role in the connection, so the ACL requires two entries per peer:

```
switch(config)# access-list ip CONTROLPLANE
switch(config-acl-ip)# 800 comment LOCKDOWN BGP SESSIONS
switch(config-acl-ip)# 805 permit tcp 10.20.0.10 gt 1023 any eq 179
switch(config-acl-ip)# 810 permit tcp 10.20.0.10 eq 179 any gt 1023
```

After allowing traffic from all configured peers, block all other devices from establishing a BGP peering session by denying all other traffic to or from TCP port 179.

```
switch(config-acl-ip)# 890 deny tcp any gt 1023 any eq 179
switch(config-acl-ip)# 895 deny tcp any eq 179 any gt 1023
```

Authenticate BGP Peers Using MD5

The TCP sessions between the two peers can be secured by adding MD5 protection to the TCP session header. The MD5 digest acts like a password between peers. This configuration is done within the BGP configuration context, and both peers need to configure the same password.

```
switch(config-bgp)# neighbor 10.20.0.10 password plaintext meatballs4me!
```

BGP TTL Security

Assuming most routing neighbors are typically directly connected, a simple method to block remote spoofing from remote devices is to check the TTL of the packets sent from the peer and dropped packets whose TTL is less than the expected amount. Following example uses the BGP peer specified above. Assuming the maximum TTL value is 255, the packets sent from the peer are compared against the hop-count, entered below as a value of 1.

```
switch(config-bgp)# neighbor 10.20.0.10 ttl-security-hops 1
```

With a maximum TTL value of 255 and a configured hop count value of 1, the packets with a TTL below 254 will be dropped.

```
switch # show run bgp
```

Multicast Security

The following sections describe security and hardening workflows for multicast traffic.

SSDP

The Simple Service Discovery Protocol (SSDP) is an application layer protocol and one of the key protocols that implement Universal Plug and Play (UPnP). SSDP enables network devices to discover and advertise network services by sending multicast discovery and advertisement messages to multicast IPv4 group address **239.255.255.250:1900** or multicast IPv6 group address **FF0x::C**. With UPnP, each device generates a unique multicast flow (Source IP, SSDP Group IP). In a multicast network with many end user devices, this can consume a large amount of multicast hardware and software resources as each device creates a unique (S, G) flow and the resources are limited. In networks where there is a need to control, drop, or minimize SSDP traffic, summarized static multicast routes can be configured to save network resources and to avoid denial of services.

The following example shows a typical static multicast route: [Incoming interface, Source, Group] > [Set of downstream interfaces]:

```
switch(config)# ip multicast-static-route vlan10 any 239.250.255.250 1/1/2
```

```
switch# show ip multicast-static-route all-vrfs
VRF : default
Group Address : 239.250.255.250
Source Address : any
Route type : Static
Incoming interface : 1/1/2
Outgoing Interface List :
Interface State
-----
vlan10 forwarding
```

If the SSDP service is not enabled in the network, best practices is to disable SSDP either through VLAN ACLs or through a policy, as shown in the following examples:

```
switch(config)# access-list ip drop_ssdp
switch(config-acl-ip)#10 deny udp any 239.255.255.250 eq 1900
switch(config)# vlan 10
switch(config-vlan-10)# apply access-list ip drop_ssdp in
switch(config)# interface 1/1/1
switch(config-if)# no shutdown
switch(config-if)# no routing
switch(config-if)# vlan access 10
switch(config)# interface vlan 10
switch(config-if-vlan)# ip address 192.168.1.2/24
switch(config-if-vlan)# ip igmp enable
switch(config-if-vlan)# ip pim-sparse enable
switch(config)# router pim
switch(config-pim)# enable
```

```
switch(config)# class ip drop_class
switch(config-class-ip)# 10 match any any 239.255.255.250
switch(config)# policy drop_ssdp
switch(config-policy)# 10 class ip drop_class action drop
```

```
switch(config)# vlan 10
switch(config-vlan-10)# apply policy drop_ssdp in
switch(config)# interface 1/1/1
switch(config-if)# no shutdown
switch(config-if)# no routing
switch(config-if)# vlan access 10
switch(config)# interface vlan 10
switch(config-if-vlan)# ip address 192.168.1.2/24
switch(config-if-vlan)# ip igmp enable
switch(config-if-vlan)# ip pim-sparse enable
switch(config)# router pim
switch(config-pim)# enable
```

To view the configuration change, issue the command **show run pim**.

Hardening IGMP and MLD Snooping

IGMP snooping runs on a Layer 2 device as a multicast constraining mechanism to improve multicast forwarding efficiency. It creates Layer 2 multicast forwarding entries from IGMP packets that are exchanged between the hosts and the router. If IGMP snooping is not enabled, the snooping switch floods multicast packets to all hosts in a VLAN. IGMP L2 snooping switch provides the benefit of conserving bandwidth on those segments of the network where no node has expressed interest in receiving packets addressed to the group address. When IGMP snooping is enabled, the L2 snooping switch forwards multicast packets of known multicast groups to only the receivers.

Multicast Listener Discovery (MLD) snooping optimizes multicast traffic across the network to prevent traffic from flooding ports on a VLAN. For example, one of the features of MLD snooping lets you configure the network so that traffic is forwarded only to ports that initiate an MLD request for multicast. Another feature of MLD lets you enable a setting so that packets that do not match the configured version are dropped. Ports can be blocked from traffic.

The device cannot provide multicast services to legal users when it has many invalid multicast entries that are created based on IGMP or MLD reports from malicious users. To control the multicast groups that hosts can join, configure a multicast group policy on the Layer 2 device that is enabled with IGMP snooping or MLD snooping. When a host sends an IGMP or MLD report to request a multicast program, the Layer 2 device uses the multicast group policy to filter the report. The Layer 2 device adds the port of the host to the outgoing port list only if the report is permitted by the multicast group policy.

```
switch(config)# ip igmp snooping apply access list <ACL-NAME>
```

```
switch(config)# ipv6 mld snooping apply access list <ACL-NAME>
```

- Existing classifier commands are used to configure the ACL.
- If an IGMPv3 packet with multiple group addresses is received, the switch only processes the permitted group addresses from the ACL rule set. The packet is forwarded to the querier and PIM router even though one of the groups present in the packet is blocked by ACL. This avoids a delay in learning the permitted groups. Since the access switch configured with ACL blocks the traffic for the groups which are denied, the forwarding of join messages has no impact. If all the groups in the packet are denied by the ACL rule, the packet is not forwarded to the querier and PIM router. Existing join messages will time out.
- In a deployment with IGMPv2, if there is no match or if there is a deny rule match, the packet is dropped.

Hardening PIM and PIMv6

In a network where IP multicast traffic is transmitted for multimedia applications, this traffic is blocked at routed interface (VLAN) boundaries unless a multicast routing protocol is running. Protocol Independent Multicast (PIM) is a family of routing protocols that form multicast trees to forward traffic from multicast sources to subnets that have used a protocol such as IGMP to request the traffic.

PIM Accept-Register

PIM **Accept-register** is a security feature that allows control over which sources and groups can register with the Rendezvous Point (RP). The command is configured in addition to the RP configuration and includes an access list option that controls sources and groups based on the following access list configuration:

```
switch(config)# access-list ip pim_reg_acl
switch(config-acl-ip)# 10 permit any 20.1.1.1 225.1.1.2
switch(config-acl-ip)# 20 deny any 30.1.1.1 225.1.1.3
switch(config)# router pim
switch(config-pim)# accept-register access-list pim_reg_acl
```

```
switch(config)# access-list ipv6 pim_regv6_acl
switch(config-acl-ipv6)# 10 permit any 20:::1 ff1e:::1
switch(config-acl-ipv6)# 20 deny any 30:::1 ff1e:::3
switch(config)# router pim6
switch(config-pim6)# accept-register access-list pim_regv6_acl
```

When a register ACL is associated with a PIM Router, the PIM protocol will store the source and destination address details along with the action (permit or deny). If there are any existing flows, the user will need to disable and enable PIM on the interface to apply the ACL. Upon receiving the register messages, a lookup is performed to check if the S and G in the packet is in the permitted list. If there is no match or if there is a deny rule match, a register stop message is immediately sent and the packet is dropped and no further action is taken. Permitted packets will go through the normal flow.

PIM Accept-RP

PIM **Accept-RP** prevents unwanted rendezvous points in the PIM sparse mode domain. By default, an RP will accept all multicast groups in the 224.0.0.0/4 range (the entire class D range), but if required, you can configure the router to allow only PIM join/prune messages toward the wanted groups.

```
switch(config)# access-list ip pim_rp_grp_acl
switch(config-acl-ip)# 10 permit any any 225.1.1.2/255.255.255.0
switch(config-acl-ip)# 20 permit any any 239.1.1.2/255.255.255.0
switch(config-acl-ip)# router pim
switch(config-pim)# accept-rp 30.1.1.1 access-list pim_rp_grp_acl
```

```
switch(config-pim)# access-list ip pim_rpv6_grp_acl
switch(config-acl-ipv6)# 10 permit any any ff2e::2/64
switch(config-acl-ipv6)# 20 permit any any ff1e::1/64
switch(config-acl-ipv6)# router pim6
switch(config-pim6)# accept-rp 30:::1 access-list pim_rpv6_grp_acl
```

PIM SSM

Protocol Independent Multicast - Source-Specific Multicast (PIM-SSM) uses a subset of PIM sparse mode and IGMP version 3 (IGMPv3)/MLD version 2 (MLDv2) to allow multicast receivers to receive traffic from specified source addresses.

The default PIM SSM group range is IPv4 - **232.0.0.0/8** and IPv6 - **FF3x::/32**. A range access list allows the PIM router to modify the default SSM range

- In the ACL used to specify the PIM-SSM range, ACEs should contain only multicast group addresses in the destination IP field, else the ACE will be ignored.
- Modifying the PIM-SSM range can lead to momentary traffic loss until PIM rebuilds the states.
- It is recommended to keep the SSM range the same across the network.

```
switch(config)# access-list ip pim_ssm_grp_range_acl
switch(config-acl-ip)# 10 permit any any 225.1.1.2/255.255.255.0
switch(config-acl-ip)# 20 permit any any 239.1.1.2/255.255.255.0
switch(config)# router pim
switch(config-pim)# pim-ssm range-access-list pim_ssm_grp_range_acl
```

```
switch(config)# access-list ipv6 pim_ssm_v6grp_range_acl
switch(config-acl-ipv6)# 10 permit any any ff2e::2/64
switch(config-acl-ipv6)# 20 permit any any ff1e::1/64
switch(config)# router pim6
switch(config-pim6)# pim-ssm range-access-list pim_ssm_v6grp_range_acl
```

To view the configuration changes, issue the commands **show run pim** or **show run pim6**.

Securing MSDP

Multicast Source Discovery Protocol (MSDP) is a mechanism to connect multiple Protocol Independent Multicast sparse mode (PIM-SM) domains. MSDP allows multicast sources for a group to be known to all rendezvous points (RPs) in different domains. An RP runs MSDP over TCP to discover multicast sources in other domains. The main advantage of MSDP is that it reduces the complexity of interconnecting multiple PIM-SM domains by allowing PIM-SM domains to use an interdomain source tree (rather than a common shared tree).

To enhance MSDP security, enable MD5 authentication for both MSDP peers to establish a TCP connection. If the MD5 authentication fails, the TCP connection cannot be established.

```
switch(config)# router msdp
switch(config-msdp)# ip msdp peer 10.1.1.1
switch(config-msdp-peer)#
switch(config-msdp-peer)# password
Enter the MD5 password: *****
Re-Enter the MD5 password: *****
```

MSDP uses SA (Source Active) messages that contain S,G (Source Group) information for RPs (Rendezvous Points) in PIM sparse domains.

By default, the MSDP enabled router forwards all the SA messages, and the peer router processes all the received messages. This command allows the user to configure an ACL on the MSDP peer to filter SA messages. User can prevent the incoming and outgoing SA messages on MSDP router by creating incoming and outgoing filter lists using an ACL.

```
switch(config-msdp-peer)# sa-filter in access-list msdp_sa_filter1
switch(config-msdp-peer)# sa-filter out access-list msdp_sa_filter2
```

To view the configuration changes, issue the command **show run msdp**.

NAE Scripts

The Network Analytics Engine is a first-of-its-kind built-in framework for network assurance and remediation. Combining the full automation and deep visibility capabilities of the AOS-CX operating system, this unique framework enables monitoring, collecting network data, evaluating conditions, and taking corrective actions through simple scripting agents. This engine is integrated with the AOS-CX system configuration and time series databases, enabling to examine historical trends and predict future problems due to scale, security, and performance bottlenecks. With that information, administrators can create software modules that automatically detect such issues and take appropriate action.

The following list describes the HPE Aruba Networking certified NAE scripts hosted on the [Aruba Solution Exchange](#). These scripts can help in hardening the switch control plane.

Table 1: HPE Aruba Networking certified NAE scripts

Module	Objective	Script
Control Plane Policing	The purpose of this script is to detect anomalous traffic that is getting dropped by Control Plane Policing.	copp.5.1
Spanning Tree	This script monitors the STP BPDU counters, Topology Change Notifications (TCN) and raise an alert if the rate of TCN exceeds a certain threshold value for a specified time. It also monitors the number of STP packets dropped at CoPP.	tp_bpdu_tcn_rate_monitor.2.0
ARP	The purpose of this script is to help in monitoring the number of ARP requests coming to the switch CPU.	arp_request_monitor.2.0

HPE is a leader in the ICT (Information and Communication Technology) industry for supply chain cybersecurity. HPE recognizes the importance of secure software and hardware development, enabling the availability of parts from trusted sources, building products with advanced security features, and accessing data that is protected within secure environments.

As cybersecurity threats evolve, HPE continues to identify and mitigate cybersecurity risks within the supply chain and provide secure products so you can concentrate on your business goals.

Supply chain attacks caused by malicious actors infiltrating systems through partners or technology vendors with access to data and resources are on the rise. Mitigating cybersecurity risks and preventing attacks in the supply chain is essential to provide secure products and services. Following are the list of security best practices that can be followed from AOS-CX switching perspective.

- Boot the switch in [Enhanced Secure Mode](#).
- Disable the [USB Auxiliary port](#) when not in use.
- Disable all physical interfaces and the OOBM port using the following commands:

```
switch(config)# interface <physical interface range>
switch(config)# disable
switch(config)# exit

switch(config)# interface mgmt
switch(config)# shutdown
switch(config)# exit
```

- Disable all management protocols (https-server, SSH, SNMP) and force the console into the device configuration to disable the management protocols on all the enabled VRFs using the following commands:

```
switch(config)# no ssh server vrf <vrf-name>
switch(config)# no https-server vrf <vrf-name>
switch(config)# no snmp-server vrf <vrf-name>
```

- Enable [password complexity](#) with a strict set of requirements.
- Enable the [ServiceOS password prompt](#).
- Disable the Central client using the following commands:

```
switch(config)# aruba-central
switch(config-aruba-central)# disable
switch(config-aruba-central)# exit
```

- Enable only [NDcPP approved SSH algorithms](#).

Accessing HPE Aruba Networking Support

HPE Aruba Networking Support Services	https://www.arubanetworks.com/support-services/
AOS-CX Switch Software Documentation Portal	https://www.arubanetworks.com/techdocs/AOS-CX/help_portal/Content/home.htm
HPE Aruba Networking Support Portal	https://networkingsupport.hpe.com/home
North America telephone	1-800-943-4526 (US & Canada Toll-Free Number) +1-408-754-1200 (Primary - Toll Number) +1-650-385-6582 (Backup - Toll Number - Use only when all other numbers are not working)
International telephone	https://www.arubanetworks.com/support-services/contact-support/

Be sure to collect the following information before contacting Support:

- 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

Other useful sites

Other websites that can be used to find information:

Airheads social forums and Knowledge Base	https://community.arubanetworks.com/
HPE Aruba Networking Hardware Documentation and Translations Portal	https://www.arubanetworks.com/techdocs/hardware/DocumentationPortal/Content/home.htm

HPE Aruba Networking software	https://networkingsupport.hpe.com/downloads
Software licensing and Feature Packs	https://lms.arubanetworks.com/
End-of-Life information	https://www.arubanetworks.com/support-services/end-of-life/
HPE Aruba Networking Developer Hub	https://developer.arubanetworks.com/

Accessing Updates

You can access updates from the HPE Aruba Networking Support Portal at <https://networkingsupport.hpe.com>.

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 subscribe to eNewsletters and alerts:

<https://networkingsupport.hpe.com/notifications/subscriptions> (requires an active HPE Aruba Networking Support Portal account to manage subscriptions). Security notices are viewable without an HPE Aruba Networking Support Portal account.

Warranty Information

To view warranty information for your product, go to <https://www.arubanetworks.com/support-services/product-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 <https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

Additional regulatory information

HPE Aruba Networking is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements, environmental data (company programs, product recycling, energy efficiency), and safety information and compliance data, (RoHS and WEEE). For more information, see <https://www.arubanetworks.com/company/about-us/environmental-citizenship/>.

Documentation Feedback

HPE Aruba Networking 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-switching@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.