

AOS-CX 10.17 Monitoring Guide

4100i, 6000, 6100 Switch Series



**Hewlett Packard
Enterprise**

Published: November 2025

Version: 1

Copyright Information

© Copyright 2025 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.

Acknowledgment

Intel®, Itanium®, Optane™, Pentium®, Xeon®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the U.S. and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

About this document	7
Applicable products	7
Latest version available online	7
Command syntax notation conventions	7
About the examples	8
Identifying switch ports and interfaces	8
Monitoring hardware through visual observation	10
Confirming normal operation of the switch by reading LEDs	10
Detecting if the switch is not ready for a failover event	11
Finding faulted components using the switch LEDs	11
Limitations, Conflicts or Exclusions	12
Alarms	14
Alarm commands	14
alarm	14
alarm input	16
alarm snooze	18
show alarm	18
Boot commands	21
boot set-default	21
boot system	21
show boot-history	23
L1-100Mbps downshift	26
Limitations with speed downshift	26
L1-100Mbps downshift commands	26
downshift enable	26
show interface	27
show interface statistics	33
show interface downshift-enable	36
show running-config interface	37
Mirroring	39
Mirror statistics	40
Classifier policies and mirroring sessions	40
Mirroring commands	40
clear mirror	40
comment	41
destination interface	42
destination tunnel	43
diagnostic	45
disable	46
enable	46
mirror session	47
mirror endpoint	48
show mirror	49

show mirror endpoint	52
source interface	53
Monitoring a device using SNMP	56
Power-over-Ethernet	57
PoE commands	58
lldp dot3 poe	58
lldp med poe	59
power-over-ethernet	59
power-over-ethernet allocate-by	60
power-over-ethernet always-on	61
power-over-ethernet assigned-class	62
power-over-ethernet pre-std-detect	63
power-over-ethernet priority	64
power-over-ethernet quick-poe	65
power-over-ethernet threshold	66
power-over-ethernet trap	67
show lldp local	67
show lldp neighbor	68
show power-over-ethernet	69
Aruba AirWave	72
SNMP support and AirWave	72
SNMP on the switch	72
Supported features with AirWave and the AOS-CX switch	73
Configuring the AOS-CX switch to be monitored by AirWave	73
AirWave commands	74
logging	74
snmp-server community	76
snmp-server host	77
snmp-server vrf	79
snmpv3 context	79
snmpv3 user	80
Support and Other Resources	83
Accessing HPE Aruba Networking Support	83
Accessing Updates	84
Warranty Information	84
Regulatory Information	84
Documentation Feedback	84

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

Applicable products

This document applies to the following products:

- HPE Aruba Networking 4100i Switch Series (JL817A, JL818A)
- HPE Aruba Networking 6000 Switch Series (R8N85A, R8N86A, R8N87A, R8N88A, R8N89A, R9Y03A, S4R20A, S4R21A, S4R22A, S4R23A, S4R24A, S4R25A, S4R26A, S4R27A, S4R28, S4R29A)
- HPE Aruba Networking 6100 Switch Series (JL675A, JL676A, JL677A, JL678A, JL679A)

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](#).

Command syntax notation conventions

Convention	Usage
<code>example-text</code>	Identifies commands and their options and operands, code examples, filenames, pathnames, and output displayed in a command window. Items that appear like the example text in the previous column are to be entered exactly as shown and are required unless enclosed in brackets ([]).
example-text	In code and screen examples, indicates text entered by a user.
Any of the following: <ul style="list-style-type: none">▪ <code><example-text></code>▪ <code><example-text></code>▪ <i>example-text</i>▪ <i>example-text</i>	Identifies a placeholder—such as a parameter or a variable—that you must substitute with an actual value in a command or in code: <ul style="list-style-type: none">▪ For output formats where italic text cannot be displayed, variables are enclosed in angle brackets (< >). Substitute the text—including the enclosing angle brackets—with an actual value.▪ For output formats where italic text can be displayed, variables might or might not be enclosed in angle brackets. Substitute the text including the enclosing angle brackets, if any, with an actual value.
	Vertical bar. A logical OR that separates multiple items from which you can choose only one. Any spaces that are on either side of the vertical bar are included for

Convention	Usage
	readability and are not a required part of the command syntax.
{ }	Braces. Indicates that at least one of the enclosed items is required.
[]	Brackets. Indicates that the enclosed item or items are optional.
... or ...	Ellipsis: <ul style="list-style-type: none"> ▪ In code and screen examples, a vertical or horizontal ellipsis indicates an omission of information. ▪ In syntax using brackets and braces, an ellipsis indicates items that can be repeated. When an item followed by ellipses is enclosed in brackets, zero or more items can be specified.

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 HPE Aruba Networking 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 HPE Aruba Networking 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.

Confirming normal operation of the switch by reading LEDs

This task describes using the switch LEDs to confirm that the switch is operating normally.



For complete information on LED behaviors for your AOS-CX switch, refer to the **Installation and Getting Started Guide** for that switch series, available for download from the [Aruba Switch Documentation](#) section of the [Aruba Hardware Documentation and Translations Portal](#).

Procedure

1. Quick check: Verify that the chassis has power and there are no fault conditions.
On the front of the switch, verify that the states of the following LEDs are On Green:
 - Power
 - Health
2. Verify that the Health LEDs of all installed line modules are On Green.
3. Verify that the Health LEDs of all installed management modules are On Green.
4. Verify that the network ports are operating normally.
 - a. On the active management module, check the Status Front section. Verify that each LED that indicates a line module is in one of the following states:
 - On Green (normal operation)
 - Off (no line module installed)
 - b. On each line module, verify that each port LED is in one of the following states:
 - On Green, Half-Bright Green, or Flickering Green (normal operation)
 - Off (no cable connected or port off by default in config)
5. Verify that the power supplies are operating normally.
 - a. On the active management module, check the Status Front section. Verify that each LED that indicates a power supply is in one of the following states:
 - On Green (normal operation)
 - Off (no power supply installed)
 - b. On each power supply, verify that LEDs are in the following states:
 - Power LED: On Green
 - Fault LED: Off
6. Verify that the rear components are operating normally by checking the Status Rear section of the active management module:
 - a. Verify that the LEDs for the fabric modules are in one of the following states:
 - On Green (normal operation)
 - Off (component not installed)
 - b. Verify that the LEDs for the fan trays and fans are On Green.

7. Verify that the standby management module is ready to take over as the active management module. On the standby management module, verify the states of the following LEDs:
 - Health LED is On Green.
 - Management state standby (Stby) LED is On Green.

Detecting if the switch is not ready for a failover event

This task describes using the switch LEDs to detect if the switch is not ready for the loss of a fabric module or for a failover from the active management module to the standby management module.



Although you can detect power supply failures by viewing the LEDs, you must use software commands to determine if the power supply redundancy is sufficient to power the chassis if a power supply fails. For complete information on LED behaviors for your AOS-CX switch, refer to the **Installation and Getting Started Guide** for that switch series, available for download from the [Aruba Switch Documentation](#) section of the [Aruba Hardware Documentation and Translations Portal](#).

Procedure

1. Detect if the standby management module is shut down.
If the standby management module is shut down, the LED states are as follows:
 - The standby management module health LED is Off.
 - The standby management state active (Actv) LED is Off.
 - The standby management state standby (Stby) LED is Off.
 - On the active management module in the Status Front Management Modules section, the LED for the standby management module is Off. For example, if the active management module is Management Module LED 5, Management Modules LED 6 is Off.
2. Detect if the standby management module is in a transient state. If the standby management module is booting, updating, or in another transient state, the LED states are as follows:
 - The standby management module health LED is Slow Flash Green when the service operating system is running or during an operating system update.
 - The standby management module Booting LED is Slow Flash Green when the AOS-CX operating system is booting.
 - The standby management state active (Actv) LED is Off.
 - The standby management state standby (Stby) LED is Off.
 - On the active management module in the Status Front Management Modules section, the LED for the standby management module is Slow Flash Green.
3. Detect if a fabric module is shut down or not present. If a fabric module is shut down or not present, the LED states are as follows:
 - On the active management module, in the Status Rear section, the LED for the fabric module is Off.
 - On the rear display module, the LED for the fabric module is Off.
 - On the fabric module, the health LED is Off. However, the fabric module is behind fan 1 and is not directly visible.

Finding faulted components using the switch LEDs

This task describes using the switch LEDs to find components that are in a fault condition.



All green LEDs—except for chassis power LEDs and the Usr1 LED—are off when the LED mode is set to Light Faults (The Usr1 LED of the LED Mode section of the active management module is On Green and the default behavior for the Usr1 LED is being used.). For complete information on LED behaviors for your AOS-CX switch, refer to the **Installation and Getting Started Guide** for that switch series, available for download from the [Aruba Switch Documentation](#) section of the [Aruba Hardware Documentation and Translations Portal](#).

Procedure

1. Find the switch that has the fault condition, which is indicated by a chassis health LED in the state of Slow Flash Orange.

The chassis health LED is located on the front of the switch and on the rear panel of the switch.

2. If you are at the back of the switch, on the rear panel, look for LEDs that are in the Slow Flash Orange state:

The Status Rear area has LEDs for power supplies, fabric modules, fan trays, and fans. The number on the LED represents the unit number of the component.

If the only LED in a state of Slow Flash Orange is the Chassis health LED, go to the front of the switch.

3. At the front of the switch, on the active management module, look for LEDs that are in the Slow Flash Orange state:
 - The Status Front area has LEDs for power supplies, line and fabric modules, and management modules. The number on the LED indicates the slot number of the component.
 - The Status Rear area has LEDs for fabric modules and fan trays, with a single LED for all the fans in the fan tray. The number on the LED represents the slot or bay number of the component.

4. Use the number indicated by the LED that is flashing to locate the slot that contains the faulted component.

The fabric modules are located behind the fan trays, and the fabric module number corresponds to the fan tray number.

5. At the front of the switch, on line modules, look for LEDs that are in the Slow Flash Orange state: Module LEDs and Port LEDs indicate faults if their states are Slow Flash Orange.

Limitations, Conflicts or Exclusions

- Platform should support Forwarding Info infra to fetch the nexthop and interface details.
- Traffic Insight feature needs to be running on all devices in the path to match on the Policy/ACL drop matches. If Traffic Insight is not supported on the devices, then the Flow drop will not be detected.
- If Traffic Insight feature is not supported or enabled along with Application recognition, Flow Tracking, IPFix, TI with AppFlow monitor on the first-hop device, then the flow trace query should have the complete 5 tuple information. Partial query having only the Client IP and destination (AppName, DomainName) will result in lookup failure.
- Client-Insight ARPtoGW needs to be enabled on the first Access switch if acting as an L2 device to make sure that the MAC-IP association of the GW are available.
- Client-IP-Tracker needs to be enabled on the first Access switch if acting as an L2 device to make sure that the MAC-IP association of the Client/Source are available.
- If the Clients have static IPs, the **gateway-ip** needs to be provided in the flow trace query.

- The feature sends out proprietary packets using UDP Source L4 Port set to 55000 and Destination L4 Port set to 55001, which are private ports as per IANA.
- In case of Flow Trace capability unaware devices present on the path, the pathtrace post that device could result in an incorrect path.
- If the last-hop device IP is not provided along with the 5-tuple query, then packet would get leaked to the destination device if the last hop switch is not Flow Trace capable.
- The MAC and Routing tables may be constantly changing. The egress port received is valid for that snapshot in time. It may not remain the same in the future.
- This feature is not applicable for broadcast, multicast or unknown unicast packets.
- IPv6 support is not provided.
- Below drops are not captured in current release:
 - GBP Policy drops
 - QoS drops
 - Per-Hop latency
 - Path Latency
 - GBP relay issues
 - PBR drops
- Support for UBT, IPSEC or GRE tunnels is not present.
- Periodic probes for the 5-tuple query is not supported.
- Number of non-flowtrace capable devices in the path (applicable for L3 only).
- Automatic Reverse Flow Trace is not supported.
- If Application provided is not active in TI for the provided source IP, then the flow trace query will fail.
- If destination details are not provided, then the flow trace query will fail.
- Uses IP exception CoPP class which can be updated while executing the **flowtrace** query.
- If the ARP table in case of L3 or the CIPT/CIARPGW table in case of L2 does not have the MAC-IP binding on the first-hop device, then the query is not initiated and failed.
- LLDP neighbor information is used to figure out if the adjacent device is a Flow Trace capable or not. So all the devices in the network need to have LLDP enabled to make sure that the device classification is correct.
- If the TLV packet payload increases above the MTU of the egress interface, the flow trace will be stopped on that device and a partial trace information will be sent back to the initiating device. Each node will add multiple TLV information depending on the attributes getting collected.
- It is advisable to have the MTU in the network to be jumbo capable to make sure the flow trace is complete especially when there are more than 7-8 nodes between the source and destination device.

Alarms allow you to configure physical alarms for alarm input ports and specific systems events.



Alarms are only available on the 4100i switch.

Alarm commands

alarm

```
alarm
  power-supply|temperature
    action <LOG-AND-TRAP>|<RELAY>
  input IN1|IN2
    action <LOG-AND-TRAP>|<RELAY>
    name <STRING>
    trigger <CLOSED>| <OPEN>
  no...
  snooze
    <time in minutes>
  repeat
```

Description

Configures input alarm and actions, or global events to be forwarded to the output alarm port, or temporarily disables (snoozes) active event forwarding to the alarm relay for a specified time interval.

Parameter	Description
temperature	Selects the alarm for ambient temperatures reaching the threshold limit of 70°C.
power-supply	Selects the alarm events from the power supply.
action	Specifies the action to be taken when the monitored event occurs.
<LOG-AND-TRAP>	Generates an event log and SNMP trap.
<RELAY>	Relays an event to alarm output port.
input IN1 IN2	Specifies the input alarm port.
action	Specifies the action to be taken when the monitored event occurs.

Parameter	Description
<LOG-AND-TRAP>	Generates an event log and SNMP trap.
<RELAY>	Relays an event to alarm output port.
name <STRING>	Specifies the external device connected.
trigger	Triggers an alarm based on a normally open or closed circuit.
<CLOSED>	Generates an alarm event when the circuit is closed.
<OPEN>	Generates an alarm event when the circuit is open.
no ...	Negates any configured parameter or removes the specified configuration.
snooze	Snooze the relay action for specified time interval.
<time in minutes>	Specifies the snooze time in minutes. The range is 0-1440 minutes..
repeat	Repeats the previous snooze time.

Examples

Disabling a temperature event to remove the configuration for all actions associated with the event:

```
switch(config)# no alarm temperature
```

Configuring an alarm for a temperature event for the log-and-trap action:

```
switch(config)# alarm temperature action log-and-trap
```

Removing the configuration for the temperature event for a log-and-trap action:

```
switch(config)# no alarm temperature action log-and-trap
```

Configuring an alarm for a power-supply event for the relay action:

```
switch(config)# alarm power-supply action relay
```

Removing the configuration for the power-supply event for the relay action:

```
switch(config)# no alarm power-supply action relay
```

Configuring an alarm on input port 1 named Door-Sensor:

```
switch(config)# alarm input IN1 name Door-Sensor
```

Removing the configuring for an alarm on input port 1:

```
switch(config)# no alarm input IN1
```

Configuring an alarm on input port 1 with log-and-trap action:

```
switch(config)# alarm input IN1 action log-and-trap
```

Removing the configuration for an alarm on input port 1 with log-and-trap action:

```
switch(config)# no alarm input IN1 action log-and-trap
```

Configuring an alarm on input port 1 to trigger an alarm when the door sensor circuit is closed:

```
switch(config)# alarm input IN1 trigger closed
```

Configuring an alarm relay action for 10 minutes:

```
switch(config)# alarm snooze 10
```

Configuring an alarm relay action to be repeated with previously configured snooze time:

```
switch(config)# alarm snooze repeat
```

Command History

Release	Modification
10.08	Feature introduced.

Command Information

Platforms	Command context	Authority
4100i	config	Administrators or local user group members with execution rights for this command.

alarm input

```
alarm input {IN1 | IN2} [name <STRING>] [action <LOG-AND-TRAP> | <RELAY>] [trigger <CLOSED | OPEN>]
no alarm input {IN1 | IN2} [name <STRING>] [action <LOG-AND-TRAP> | <RELAY>] [trigger <CLOSED | OPEN>]
```

Description

Configures input alarm and actions. The **no** form of this command removes the specified configuration.

Parameter	Description
{ IN1 IN2 }	Specifies the input alarm port..
name	Specifies the external device connected.
<STRING>	Descriptive string.
action	Specifies the action to be taken when the monitored event occurs.
<LOG-AND-TRAP>	Generates an event log and SNMP trap.
<RELAY>	Relays an event to alarm output port.
trigger	Triggers an alarm based on a normally open or closed circuit.
<CLOSED>	Generates an alarm event when the circuit is closed.
<OPEN>	Generates an alarm event when the circuit is open.

Examples

Configuring an alarm on input port 1 named Door-Sensor:

```
switch(config)# alarm input IN1 name Door-Sensor
```

Removing the configuring for an alarm on input port 1:

```
switch(config)# no alarm input IN1
```

Configuring an alarm on input port 1 with log-and-trap action:

```
switch(config)# alarm input IN1 action log-and-trap
```

Removing the configuration for an alarm on input port 1 with log-and-trap action:

```
switch(config)# no alarm input IN1 action log-and-trap
```

Configuring an alarm on input port 1 to trigger an alarm when the door sensor circuit is closed:

```
switch(config)# alarm input IN1 trigger closed
```

Command History

Release	Modification
10.08	Feature introduced.

Command Information

Platforms	Command context	Authority
4100i	config	Administrators or local user group members with execution rights for this command.

alarm snooze

```
alarm snooze [time in minutes] [repeat]
```

Description

Configures any active event forwarded to alarm relay to be snoozed.

Parameter	Description
time in minutes	Specifies the value for time in minutes. The range is 0-1440 minutes..
repeat	Repeats the previous snooze time.

Examples

Configuring an alarm relay action for 10 minutes:

```
switch(config)# alarm snooze 10
```

Configuring an alarm relay action to be repeated with previously configured snooze time:

```
switch(config)# alarm snooze repeat
```

Command History

Release	Modification
10.08	Feature introduced.

Command Information

Platforms	Command context	Authority
4100i	config	Administrators or local user group members with execution rights for this command.

show alarm

```
show alarm
  input [IN1 | IN2]
  power supply
  temperature
  timer
```

Description

Shows all of the details of global status monitoring alarm events.

Parameter	Description
{IN1 IN2}	Specifies the input alarm port..
power supply	Selects the alarm events from the power supply.
timer	Shows the status of an alarm snooze timer's status and duration.
temperature	Selects the alarm for ambient temperatures reaching threshold limits.

Examples

Showing details for all global alarm events on the switch:

```
switch# show alarm
Alarm Snooze Timer Status: active
Duration remaining: 1 min 32 sec
-----
Global Alarm: Temperature
-----
Alarm Event      Status      log-and-trap  Relay
-----
Temperature      inactive    true           false
-----
Global Alarm: power-supply
-----
Alarm Event      Status      log-and-trap  Relay
-----
power-supply     active      false         true
```

Showing details for the temperature alarm:

```
switch# show alarm temperature
Alarm Snooze Timer Status: Active
Duration remaining: 5 min 36 sec
-----
Global Alarm: Temperature
-----
Alarm Event      Status      log-and-trap  Relay
-----
Temperature      inactive    true           false
```

Showing the status of an alarm snooze timer when it is inactive :

```
switch# show alarm timer
Alarm Snooze Timer Status: inactive
```

Showing the status of an active 3-minute alarm snooze timer:

```
switch# show alarm timer
Alarm Snooze Timer Status: active
Duration remaining: 2 min 55 sec
```

Showing details for all alarm input ports on the switch:

```
switch# show alarm input
Alarm Snooze Timer Status: inactive
-----
Input Alarm IN1, Name: Door-Sensor
-----
Alarm Port      Status      log-and-trap  Relay      Trigger
-----
IN1             inactive   true          false     closed
-----
Input Alarm IN2, Name: N/A
-----
Alarm Port      Status      log-and-trap  Relay      Trigger
-----
IN2             active     false         true      open
```

Showing details for alarm input ports IN1:

```
switch# show alarm input IN1
Alarm Snooze Timer Status: inactive
-----
Input Alarm IN1, Name: Door-Sensor
-----
Alarm Port      Status      log-and-trap  Relay      Trigger
-----
IN1             inactive   true          false     closed
```

Command History

Release	Modification
10.08	Feature introduced.

Command Information

Platforms	Command context	Authority
4100i	Manager (#)	Administrators or local user group members with execution rights for this command.

boot set-default

```
boot set-default {primary | secondary}
```

Description

Sets the default operating system image to use when the system is booted. Changes to this setting may impact other features, such as the job scheduler feature configured via the **reload at** or **reload after** commands.

Parameter	Description
primary	Selects the primary network operating system image.
secondary	Selects the secondary network operating system image.

Example

Selecting the primary image as the default boot image:

```
switch# boot set-default primary
Default boot image set to primary.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Manager (#)	Administrators or local user group members with execution rights for this command.

boot system

```
boot system [primary | secondary | serviceos]
```

Description

Reboots all modules on the switch. By default, the configured default operating system image is used. Optional parameters enable you to specify which system image to use for the reboot operation and for future reboot operations.

Parameter	Description
<code>primary</code>	Selects the primary operating system image for this reboot and sets the configured default operating system image to primary for future reboots.
<code>secondary</code>	Selects the secondary operating system image for this reboot and sets the configured default operating system image to secondary for future reboots.
<code>serviceos</code>	Selects the service operating system for this reboot. Does not change the configured default operating system image. The service operating system acts as a standalone bootloader and recovery OS for switches running the AOS-CX operating system and is used in rare cases when troubleshooting a switch.

Usage

This command reboots the entire system. If you do not select one of the optional parameters, the system reboots from the configured default boot image.

You can use the **show images** command to show information about the primary and secondary system images.

Choosing one of the optional parameters affects the setting for the default boot image:

- If you select the **primary** or **secondary** optional parameter, that image becomes the configured default boot image for future system reboots. The command fails if the switch is not able to set the operating system image to the image you selected.

You can use the **boot set-default** command to change the configured default operating system image.

- If you select **serviceos** as the optional parameter, the configured default boot image remains the same, and the system reboots all management modules with the service operating system.

If the configuration of the switch has changed since the last reboot, when you execute the **boot system** command you are prompted to save the configuration and you are prompted to confirm the reboot operation.

Saving the configuration is not required. However, if you attempt to save the configuration and there is an error during the save operation, the **boot system** command is aborted.

Examples

Rebooting the system from the configured default operating system image:

```
switch# boot system
Do you want to save the current configuration (y/n)? y
The running configuration was saved to the startup configuration.

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.
The system is going down for reboot.
```

Rebooting the system from the secondary operating system image, setting the secondary operating system image as the configured default boot image:

```
switch# boot system secondary
Default boot image set to secondary.

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.
```

Canceling a system reboot:

```
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)? n
Reboot aborted.
switch#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Manager (#)	Administrators or local user group members with execution rights for this command.

show boot-history

```
show boot-history [all|{vsf member <1-10>}]
```

Description

Shows boot history information. When no parameters are specified, shows the most recent information about the current boot operation, and the three previous boot operations for the switch. When the **all** parameter is specified, the output of this command shows the boot information for the active management module.



To view boot-history on a standby, the command must be sent on the conductor console.

Parameter	Description
all	Optional. Shows boot information for the active management module.
vsf member <1-10>	Optional. Display boot history for the specified VSF member

Usage

This command displays the boot-index, boot-ID, and up time in seconds for the current boot. If there is a previous boot, it displays boot-index, boot-ID, reboot time (based on the time zone configured in the system) and reboot reasons. Previous boot information is displayed in reverse chronological order.

The output of this command includes the following information:

Parameter	Description
Index	The position of the boot in the history file. Range: 0 to 3.
Boot ID	A unique ID for the boot . A system-generated 128-bit string.
Current Boot, up for <time>	For the current boot, the show boot-history command shows the number of seconds the module has been running on the current software.
<Timestamp>: boot reason	For previous boot operations, the show boot-history command shows the time at which the operation occurred and the reason for the boot. The reason for the boot is one of the following values: <ul style="list-style-type: none"> ▪ <DAEMON-NAME> crash: The daemon identified by <DAEMON-NAME> caused the module to boot. ▪ Kernel crash: The operating system software associated with the module caused the module to boot. ▪ Uncontrolled reboot: The reason for the reboot is not known. ▪ Reboot requested through database: The reboot occurred because of a request made through the CLI or other API.

Examples

Showing the boot history of the active management module:

```
switch# show boot-history
Management module
=====

Index : 2
Boot ID : c34a2c2499004a02bbeeff4992e1fdbd
Current Boot, up for 1 days 13 hrs 13 mins 27 secs

Index : 1
```

```

Boot ID : bfba9bc486304e57904ac717a0ccbdcd
02 Sep 23 02:55:33 : CPU request reset with 0x20201, Version: FL.10.14.0000-1619-
ga9ec1805bd442~dirty
02 Sep 23 02:55:33 : Switch boot count is 2

Index : 0
Boot ID : a88a71b7ca9a4574af7e3b811ddfdc7e
02 Sep 23 02:49:26 : Reboot requested by user, Version: FL.10.14.0000-1619-
ga9ec1805bd442~dirty
02 Sep 23 02:50:02 : Switch boot count is 1

Index : 3
Boot ID : f00ba10c8c44457f83fee303d014a89a
25 Aug 23 10:27:42 : Power on reset with 0x1, Version: FL.10.14.0000-1465-
g9df95249d06b0~dirty
25 Aug 23 10:28:18 : Switch boot count is 3
25 Aug 23 10:29:02 : Primary overtemperature fault detected with 0x2 in PSU 1/1

```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Manager (#)	Administrators or local user group members with execution rights for this command.

The speed downshift feature allows the user to link-up at sub-optimal speeds when failing to link-up at the highest advertised speed. There are fixed number of link attempts made to establish link at highest advertised speed and when all of them fail and attempt is made to link-up at a lower possible speed.

This feature requires underlying PHY to have support for the same and hence capability is only added to select set of ports. If a link cannot be established at the highest common denominator within a set number of link attempts, the PHY advertises the next highest speed using auto-negotiation.

Limitations with speed downshift

- Link up may be delayed as certain number of retries are done to establish the link at highest advertise speeds by both link partners before downshifting.
- Link may be established at sub-optimal speed.

L1-100Mbps downshift commands

downshift enable

```
downshift-enable  
no downshift-enable
```

Description

Enables/disables automatic speed downshift on an interface that supports downshift, generally 1GBASE-T ports. When enabled, downshift allows an interface to link at a lower advertised speed when unable to establish a stable link at the maximum speed. Downshifting only applies to physical interfaces that are not members of a LAG and is only available when auto-negotiation is enabled. When only one speed is advertised, downshift will not be triggered.

Examples

```
switch(config-if)# interface 1/1/1  
switch(config-if)# downshift-enable
```

```
Warning: this is a non-standard mode for use only when standards-based  
auto-negotiation is not able to establish a stable link. Enabling this  
may cause the port to link at a lower than expected speed and should  
not be used on ports that are members of a LAG. Support calls may require  
this feature to be disabled
```

```
Continue (y/n)?
```

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

When automatic downshift is enabled:

```
switch(config-if)# show running-config interface
interface 1/1/1
    downshift-enable
```

Disabling automatic speed downshift:

```
switch(config-if)# interface 1/1/1
switch(config-if)# no downshift-enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

show interface

```
show interface [<IFNNAME>|<IFRANGE>] [brief | physical]
show interface [<IFNNAME>|<IFRANGE>] [extended [non-zero] | [human-readable]]
show interface [<IFNNAME>] monitor [human-readable]
show interface [lag | vlan ] [<ID>] [brief]
show interface lag [<LAG-ID>] [extended [non-zero] | [human-readable]]
show interface lag [<LAG-ID>] monitor [human-readable]
```

Description

Shows active configurations and operational status information for interfaces.

Parameter	Description
<IFNAME>	Specifies a interface name.
<IFRANGE>	Specifies the port identifier range.
brief	Shows brief info in tabular format.
physical	Shows the physical connection info in tabular format.
extended	Shows additional statistics, including the tx filtered and rx filtered counters. <ul style="list-style-type: none">Rx filter packets are protocol packets received when the protocol is disabled on the switch and there is only one port in the VLAN. Protocols include OSPF, PIM, RIP, LACP, and LLDP.An example of a Tx filtered packet would be a multicast packet

Parameter	Description
	being filtered from going out of the ingress port.
human-readable	Shows statistics rounded to the nearest power of 1000, for example, 1K, 345M, 2G. This is available only in the CLI interface output.
non-zero	Shows only non zero statistics.
LAG	Shows LAG interface information.
monitor	Continuously monitor interface statistics.
VLAN	Shows VLAN interface information.
<LAG-ID>	Specifies the LAG number. Range: 1-256
<VLAN-ID>	Specifies the VLAN ID. Range: 1-4094

Examples

Showing information when interface 1/1/1 is configured:

```

MDI mode: MDIX
VLAN Mode: native-untagged
Native VLAN: 1
Allowed VLAN List: all
Rate collection interval: 300 seconds
Rates
-----
Mbits / sec          RX          TX          Total (RX+TX)
KPkts / sec         0.00        0.00        0.00
Unicast              0.00        0.00        0.00
Multicast            0.00        0.00        0.00
Broadcast            0.00        0.00        0.00
Utilization %       0.00        0.00        0.00
Statistics
-----
                RX          TX          Total
Packets         0          0          0
Unicast         0          0          0
Multicast       0          0          0
Broadcast       0          0          0
Bytes           0          0          0
Jumbos          0          0          0
Dropped         0          0          0
Filtered        0          0          0
Pause Frames    0          0          0
Errors          0          0          0
CRC/FCS         0          n/a        0
Collision      n/a         0          0
Runts           0          n/a        0
Giants          0          n/a        0

```

Showing information when the interface is currently linked at a downshifted speed:

```
switch(config-if)# show interface 1/1/1
Interface 1/1/1 is up
...
Auto-negotiation is on with downshift active
```

Showing information when the interface is currently linked with energy-efficient-ethernet negotiated:

```
switch(config-if)# show interface 1/1/1
Interface 1/1/1 is up
...
Energy-Efficient Ethernet is enabled and active
```

Showing information when the interface is configured with EEE and the EEE has auto-negotiated:

```
switch(config-if)# show interface 1/1/1 physical
-----
EEE          PoE Power      Link  Admin      Speed      Flow-Control
Port        Type           Status Admin      Status |   Status |
Status |   Config (Watts)  State Information Description  Config
-----
1/1/1      1GbT          up    up          1G         auto      off    off
on         on            --    10M/100M/1G  --
```

Showing the monitor information:



In monitor mode, the CLI refreshes data automatically until it is exited by entering **q**. Pressing **?** opens the help menu to display which options are available in this context.

```
Interface 1/1/1 is up
Rate                RX                TX                Total (RX+TX)
-----
Mbits / sec        30196.43          30196.43          60392.85
MPkts / sec        58977.39          58977.40          117954.79
Unicast            0.00              0.00              0.00
Multicast          58977.39          58977.40          117954.79
Broadcast          0.00              0.00              0.00
Utilization %      75.49             75.49             150.98
Statistic          RX                TX                Total (RX+TX)
-----
Packets            4756527649        4756527865        9513055514
Unicast            0                  0                  0
Multicast          4756527649        4756527865        9513055514
Broadcast          2                  0                  2
Bytes              304417778668      304417795428      608835574096
Jumbos             0                  0                  0
Dropped            0                  19028847730       19028847730
Pause Frames       0                  0                  0
Errors             0                  0                  0
CRC/FCS            0                  n/a                0
help: ?, quit: q
```

```

Help for Interface Monitor
h  Toggle human-readable mode
c  Clear interface statistics
Does not apply to rates
Arrows, PgUp, PgDn, Home, End
Navigate interface statistics
Delay: 2
help: ?, quit: q

```

Showing the output for interface 1/1/1 in human-readable format:



In human-readable format, the < 1 symbol for **Utilization** indicates that the amount of packets is between zero and one. This is true in cases where the number of bytes increases but the number of packets and the **Utilization** value is not displayed even in the normal output, where the human-readable parameter is not included in the command.

```

switch(config-if)# show interface 1/1/1 human-readable
Interface 1/1/1 is up
Rate                               RX                               TX                               Total (RX+TX)
-----
Bits / sec                         3M                               3M                               6M
Pkts / sec                         316                             316                             633
Unicast                            319                             319                             638
Multicast                          0                               0                               0
Broadcast                          0                               0                               0
Utilization %                      < 1                             < 1                             < 1
Statistic                          RX                               TX                               Total
-----
Packets                            577K                            577K                            1M
Unicast                            577K                            577K                            1M
Multicast                          0                               51                              51
Broadcast                          0                               15                              15
Bytes                              744M                            745M                            1G
Jumbos                            0                               0                               0
Dropped                            0                               0                               0
Filtered                            0                               0                               0
Pause Frames                       0                               0                               0
Errors                              0                               0                               0
CRC/FCS                            0                               n/a                             0
Collision                          n/a                             0                               0
Runts                              0                               n/a                             0
Giants                              0                               n/a                             0

```

Showing information about extended counters:



The output of the `show interface extended` command varies depending on the switch model and configuration.

```

switch(config-if)# show interface 1/1/17 extended
-----
Interface 1/1/17
-----
Statistics                               Value
-----
Dot1d Tp Port In Frames                  547

```

```

Dot1d Tp Port Out Frames          608
Dot3 In Pause Frames              0
Dot3 Out Pause Frames             0
Ethernet Stats Broadcast Packets  19
Ethernet Stats Bytes              40162
Ethernet Stats Packets            342
...

```

```

-----
Error-Statistics                   Value
-----
Dot1d Base Port MTU Exceeded Discards 0
Dot3 Control In Unknown Opcodes      0
Dot3 Stats Alignment Errors          0
Dot3 Stats FCS Errors                 0
Dot3 Stats Frame Too Longs           0
Dot3 Stats Internal Mac Transmit Errors 0
Ethernet RX Oversize Packets         0
...

```

Showing interface link-status:

```
switch# show interface link-status
```

```

-----
Port                Type                Physical Link      Last
Link State          Transitions        Change
-----
1/1/1               1G-BT              down     0         --
1/1/2               1G-BT              up       1         1 minute ago (Fri Mar 09
12:36:56 UTC 2018)
1/1/3               1G-BT              up       1         1 minute ago (Fri Mar 09
12:36:56 UTC 2018)
1/1/4               --                 down     0         --
1/1/5               --                 down     0         --

```

Showing interface loopback 1 link-status:

```

-----
Port                Type                Physical Link      Last
Link State          Transitions        Change
-----
loopback1          --                 up       --        --

```

Showing interface 1/1/2-1/1/3 link-status:

```

-----
Port                Type                Physical Link      Last
Link State          Transitions        Change
-----
1/1/2               1G-BT              up       1         1 minute ago (Fri Mar 09
12:36:56 UTC 2018)
1/1/3               1G-BT              up       1         1 minute ago (Fri Mar 09
12:36:56 UTC 2018)

```

Showing interface link-status:

```
switch# show interface link-status
```

Port	Type	Physical Link State	Link Transitions	Link Flaps Ignored	Last Change
1/1/1	1G-BT	down	0	0	--
1/1/2	1G-BT	up	1	0	1 minute ago (Fri Mar 09 12:36:56 UTC 2018)
1/1/3	1G-BT	up	1	0	1 minute ago (Fri Mar 09 12:36:56 UTC 2018)
1/1/4	--	down	0	0	--
1/1/5	--	down	0	0	--

Showing state information when interface is blocked:

```
8360(config-if)# show interface 1/1/1
```

```
Interface 1/1/1 is up (Blocked)
Admin state is up
State information: Blocked by UDLD
Link state: up for 1 minute (since Mon Jun 10 09:25:27 UTC 2024)
Link transitions: 1
Description:
Persona:
Hardware: Ethernet, MAC Address: 00:fd:45:67:85:91
MTU 1500
Type 10G-LR / 10G SFP+ LR
Full-duplex
qos trust none
Speed 10000 Mb/s
Auto-negotiation is off
Flow-control: off
Error-control: off
VLAN Mode: access
Access VLAN: 1
Rate collection interval: 300 seconds
```

Rate	RX	TX	Total (RX+TX)
Mbits / sec	0.00	0.00	0.00
KPkts / sec	0.00	0.00	0.00
Unicast	0.00	0.00	0.00
Multicast	0.00	0.00	0.00
Broadcast	0.00	0.00	0.00
Utilization %	0.00	0.00	0.00
Statistic	RX	TX	Total
Packets	15	15	30
Unicast	12	12	24
Multicast	3	3	6
Broadcast	0	0	0
Bytes	1350	1350	2700
Jumbos	0	0	0
Dropped	0	0	0
Pause Frames	0	0	0
Errors	0	0	0
CRC/FCS	0	n/a	0
Collision	n/a	0	0
Runts	0	n/a	0
Giants	0	n/a	0

Command History

Release	Modification
10.15	Added state information when port goes into down state.
10.11	Added <code>monitor</code> parameter.
10.10	Added <code>human-readable</code> parameter.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show interface statistics

```

show interface [<IFNAME>|<IFRANGE>] statistics [non-zero] [human-readable]
show interface [<IFNAME>|<IFRANGE>] statistics monitor [non-zero] [human-readable]
show interface [<IFNAME>|<IFRANGE>] error-statistics [non-zero] [human-readable]
show interface [<IFNAME>|<IFRANGE>] error-statistics monitor [non-zero] [human-readable]
show interface lag [<LAG-ID>] statistics [non-zero] [human-readable]
show interface lag [<LAG-ID>] statistics monitor [non-zero] [human-readable]
show interface lag [<LAG-ID>] error-statistics [non-zero] [human-readable]
show interface lag [<LAG-ID>] error-statistics monitor [non-zero] [human-readable]

```

Description

Shows statistics for switch interfaces such as packets transmitted and received, bytes transmitted and received, broadcast and multicast packets.

Parameter	Description
<IFNAME>	Specifies a interface name.
<IFRANGE>	Specifies the port identifier range.
LAG	Shows LAG interface information.
<LAG-ID>	Specifies the LAG number. Range: 1-256
monitor	Continuously monitor interface statistics.
human-readable	Shows statistics rounded to the nearest power of 1000, for example, 1K, 345M, 2G.
non-zero	Shows only non zero statistics.

Examples

Showing statistics of all interfaces:

```
show interface statistics
```

Interface	RX Bytes	RX Packets	RX Drops	TX Bytes	TX Packets	TX Drops	RX Broadcast	RX Multicast	TX Broadcast	TX Multicast	RX Pause	TX Pause
1/1/1	2727136	1975	0	17796	195	0	82	1788	96	54	0	0
1/1/10	0	0	0	0	0	0	0	0	0	0	0	0
1/1/11	0	0	0	0	0	0	0	0	0	0	0	0
1/1/12	0	0	0	0	0	0	0	0	0	0	0	0
...												
1/1/30 - lag1	0	0	0	11271	92	0	0	0	0	51	0	0
1/1/31 - lag2	2360	25	50	2732119	2040	0	0	0	178	1839	0	0
1/1/32 - lag2	0	0	0	11373	93	0	0	0	0	51	0	0
vlan1	0	0	0	0	0	0	0	0	0	0	0	0

Showing statistics of all interfaces with only non-zero statistics:

```
show interface statistics non-zero
```

Interface	RX Bytes	RX Packets	RX Drops	TX Bytes	TX Packets	TX Drops	RX Broadcast	RX Multicast	TX Broadcast	TX Multicast	RX Pause	TX Pause
1/1/1	2727136	1975	0	17796	195	0	82	1788	96	54	0	0
1/1/30 - lag1	0	0	0	11271	92	0	0	0	0	51	0	0
1/1/31 - lag2	2360	25	50	2732119	2040	0	0	0	178	1839	0	0
1/1/32 - lag2	0	0	0	11373	93	0	0	0	0	51	0	0

Showing statistics of all interfaces in the human-readable format:

```
show interface statistics human-readable
```

Interface	RX Bytes	RX Pkts	RX Drops	TX Bytes	TX Pkts	TX Drops	RX Bcast	RX Mcast	TX Bcast	TX Mcast	RX Pause	TX Pause
1/1/1	744M	577K	0	745M	578K	0	0	0	73	287	0	0
1/1/2	474M	367K	0	475M	369K	0	0	0	73	288	0	0
1/1/3	0	0	0	0	0	0	0	0	0	0	0	0

Showing statistics of a single interfaces:

```
show interface 1/1/2 statistics
```

Interface	RX Bytes	RX Packets	RX Drops	TX Bytes	TX Packets	TX Drops	RX Broadcast	RX Multicast	TX Broadcast	TX Multicast	RX Pause	TX Pause
1/1/2	2725080	1931	0	25877	253	0	21	1788	65	55	0	0

Showing statistics of all members of a LAG interface:

```
show interface lag1 statistics
```

Interface	RX Bytes	RX Packets	RX Drops	TX Bytes	TX Packets	TX Drops	RX Broadcast	RX Multicast	TX Broadcast	TX Multicast	RX Pause	TX Pause
1/1/3 - lag1	2424	26	0	2734082	2062	0	0	0	191	1848	0	0
1/1/30 - lag1	0	0	0	12383	100	0	0	0	0	59	0	0
lag1	2424	26	0	2746465	2162	0	0	0	191	1907	0	0

Showing error statistics of all interfaces:

```
show interface error-statistics
```

Interface	RX Errors	TX Errors	Giants	Runts	CRC/FCS	Collisions
1/1/1	190	20	100647	0	0	0
1/1/10	0	0	100	290	7165	949
1/1/11	0	0	0	0	0	0
1/1/12	0	0	0	0	0	0
...						
1/1/30 - lag1	1500	500	45800	0	0	0
1/1/31 - lag2	0	0	11	27	0	0
1/1/32 - lag2	0	0	0	0	6	18

Showing monitor statistics:



The rows and columns of show interface monitor statistics depends on the length of width of the client terminal. The CLI can be navigated using the arrow keys as well as the PageUp, PageDown, Home, and End keys.

```
show interface statistics monitor
```

```
-----  
Interface                RX Bytes      RX Packets >>  
-----  
1/1/1                    3440525421984 53758209526  
1/1/2                    3440526607008 53758228042  
1/1/3                    3440527785312 53758246453  
1/1/30                   3440559671264 53758744653  
1/1/31                   3440560851680 53758763098  
1/1/32                   3440562028704 53758781489
```

```
help: ?, quit: q
```

Help for Interface Monitor

```
f  Toggle full statistics  
h  Toggle human-readable mode  
n  Toggle non-zero mode  
r  Toggle rate display
```

```
c  Clear interface statistics  
    Does not apply to rates
```

```
Arrows, PgUp, PgDn, Home, End  
    Navigate interface statistics
```

Delay:2

```
help: ?, quit: q
```

Showing monitor error statistics in human-readable format:

```
show interface 1/1/1-1/1/3,1/1/30-1/1/32 error-statistics monitor human-readable
```

```
-----  
Interface                RX Errors    TX Errors    RX Giants    RX Runts    CRC/FCS    Collisions  
-----  
1/1/1                    0            0            0            0            0            0  
1/1/2                    0            0            0            0            0            0  
1/1/3                    0            0            0            0            0            0  
1/1/30                   0            0            0            0            0            0  
1/1/31                   0            0            0            0            0            0  
1/1/32                   0            0            0            0            0            0
```

Human-readable

```
help: ?, quit: q
```

Help for Interface Monitor

```
h  Toggle human-readable mode  
n  Toggle non-zero mode
```

```
c  Clear interface statistics  
    Does not apply to rates
```

```
Arrows, PgUp, PgDn, Home, End  
    Navigate interface statistics
```

Delay:2

```
help: ?, quit: q
```

Command History

Release	Modification
10.11	Added <code>monitor</code> parameter.
10.10	Added <code>human-readable</code> parameter.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show interface downshift-enable

```
show interface [<IFNAME>|<IFRANGE>] downshift-enable
```

Description

Displays speed downshift information, including the interface speed status and configuration.

Parameter	Description
<IFNAME>	Specifies a interface name.
<IFRANGE>	Specifies the port identifier range.

Examples

Showing automatic downshift information:

```
switch(config-if)# show interface downshift-enable
-----
Port          Downshift          Speed
             Enabled | Active   Status   | Config
-----
1/1/1         yes      yes    100M-FDx  auto
1/1/2         yes      no     1G       auto
1/1/3         yes      no    100M-FDx 100M-FDx
1/1/4         no       no     --       auto
```

Showing automatic downshift information on per interface:

```
switch(config-if)# show interface 1/1/2 downshift-enable
-----
Port          Downshift          Speed
             Enabled | Active   Status   | Config
-----
1/1/2         yes      no     1G       auto
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show running-config interface

```
show running-config interface [<IFNNAME>|<IFRANGE>]
show running-config interface [lag | loopback | tunnel | vlan ] [<ID>]
```

Description

Displays active configurations of various switch interfaces.

Parameter	Description
<IFNAME>	Specifies a interface name.
<IFRANGE>	Specifies the port identifier range.
LAG	Specifies LAG interface information
LOOPBACK	Specifies loopback interface information.
TUNNEL	Specifies tunnel interface information.
VLAN	Specifies VLAN interface information.
<LAG-ID>	Specifies the LAG number. Range: 1-256.
<LOOPBACK-ID>	Specifies the LOOPBACK number. Range: 0-255.
<TUNNEL-ID>	Specifies the tunnel ID. Range: 1-255.
<VLAN-ID>	Specifies the VLAN ID. Range: 1-4094.
VXLAN	Specifies the VXLAN interface information.
<VXLAN-ID>	Specifies the VXLAN interface identifier. Default: 1.

Examples

Showing 1/1/2 interface configuration:

```
switch(config-if)# show running-config interface 1/1/2

interface 1/1/2
no shutdown
```

```
description DC-23
exit
```

Showing loopback interfaces configured:

```
switch(config-if)# show running-config interface loopback

interface loopback 1
  description lb interface 1
  exit
interface loopback 2
  description lb interface 2
  exit
```

Showing loopback interfaces not configured:

```
switch(config-if)# show running-config interface loopback

No loopback interfaces configured.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

Mirroring allows you to replicate all traffic arriving and/or leaving the selected system interfaces. This data can be used for collection or analysis.

The traffic replicated using mirroring can be sent to a separate interface on the same switch as the traffic source for analysis or inspection. Such a collection of interfaces and settings is called a mirror session.

A mirror session can be configured with many traffic sources but only a single output, or destination. In the initial configuration, the mirror session is disabled. You have enable the feature to start the replication.



Care must be taken in choosing the number and rates of sources to avoid over-saturating a session destination. A mirror session with multiple 10G sources can overwhelm a single 10G destination and important data may be lost.

This version of AOS-CX support the following mirror capabilities:

- Support for a VLAN as source for a mirror session
- 6100, 6000 and 4100i Switch series support of 2 simultaneously enabled Mirroring Sessions with LAG sources, regardless of direction
- Ability for a given Mirror source in one session to act as source in another Mirror Session
- Support for a Layer 2/bridged Link Aggregation Group (LAG) as Session destination
- Support for a Layer 3/Route-only Link Aggregation Group (LAG) as Session destination

The following interface types are not supported as mirror source or destination:



- Management
 - Persona
 - Loopback
 - VxLAN
 - Subinterfaces
-

Hit Count Behavior for Policy Mirroring Actions

When using classifier policies with mirroring actions, it is important to note that packets mirrored via these policy actions will have their hit counts recorded within the Policy hit counts. These packets will not be recorded as output packets in the Mirror Session statistics, except in the case of Mirror-to-CPU sessions. This distinction means that traffic metrics for packets mirrored via Policy Actions should be obtained from Policy hit counts, as these packets will not appear in the Mirror Session output statistics.

Mirroring and CoPP

When configuring a Mirror Session with the destination set to the CPU, it is important to understand how mirrored traffic interacts with Control Plane Policing (CoPP).

Control plane packets that are already destined for the switch (for example, routing protocol packets and management traffic) will be processed by their respective CoPP classes. These packets will still be mirrored to the CPU as part of the Mirror-to-CPU session, but they will not be counted under the Mirror-to-CPU CoPP class. The Mirror-to-CPU CoPP class is specifically designed to manage dataplane traffic that is mirrored to the CPU. Any control plane traffic mirrored to the CPU will bypass this class and be handled by the CoPP class corresponding to its original purpose (e.g., OSPF, BGP, or ARP). Administrators should account for this distinction when analyzing traffic mirrored to the CPU.

For example, if a mirror session is configured to mirror all traffic from a source interface to the CPU, and that interface receives OSPF packets, the OSPF packets will be mirrored to the CPU. These packets will be processed by the OSPF CoPP class, not the Mirror-to-CPU CoPP class.

Mirror statistics

Mirror statistics are reset for a Mirror-to-CPU session when an interface is added or removed from a LAG that is a source interface in the Mirror session and during a failover.

Classifier policies and mirroring sessions

Network traffic can be mirrored to a destination interface in two ways:

- Using a mirroring session alone.
- Using Classifier Policies with mirror actions in conjunction with a mirroring session.

Basic mirroring sessions provide coarse control over the type of traffic mirrored from a source: all received, all transmitted, or both. However, a traffic class within a Classifier Policy applied to a source can provide much finer grained control of mirrored traffic. For example, a policy can match on many different aspects of the Ethernet or IPv4 or IPv6 header information in each frame or packet received or transmitted on an interface.

The steps to configure a policy and class with a mirror action are the following:

1. Configuring a mirroring session with a destination interface.
2. Enabling the mirroring session.
3. Configuring the Classifier Policy, specifying the mirroring session ID in the mirror action.

If the packets being mirrored are received from a VLAN that is not allowed on the mirror destination, the mirrored packets would be dropped at the mirror destination interface. When the mirrored packets are dropped at the destination, the mirror output packet and byte count will increment, however the packets will not be received at the mirror destination.

The mirror destination port among the active mirror sessions must be unique. That is, if an interface is configured as a source or destination in an active mirror session, the same port cannot be used as a destination in another active mirror session.

Mirroring commands

clear mirror

```
clear mirror [all | <SESSION-ID>]
```

Description

Clears the mirror statistics for all configured mirror sessions or a specified session

Parameter	Description
all	Specifies all configured sessions.
<SESSION-ID>	Specifies a numeric identifier for the session. Range: 1 to 4

Examples

Clearing mirror statistics for all configured mirror sessions:

```
switch# clear mirror all
```

Clearing mirror statistics for mirror session 1:

```
switch# clear mirror 1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Manager (#)	Administrators or local user group members with execution rights for this command.

comment

```
comment <COMMENT>
no comment
```

Description

Specifies a comment for the mirroring session.

When used in mirror endpoint command context, specifies a comment for the mirror endpoint.

The **no** form of this command removes the comment.

Parameter	Description
<COMMENT>	A comment string of up to 64 characters composed of letters, numbers, underscores, dashes, spaces, and periods.

Usage

Comments are optional and can be added or removed at any time without affecting the state of the mirroring session.

Adding a comment to a session that already has a comment replaces the existing comment.

Examples

Adding a comment to a mirror session:

```
switch(config-mirror-3) # comment This Mirror will be removed during next  
maintenance window
```

Removing the comment from mirror session 3:

```
switch(config-mirror-3) # no comment
```

Adding a comment to a mirror endpoint:

```
switch(config-mirror-endpoint-test) # comment Monitor endpoint traffic
```

Replacing the existing comment for mirror endpoint:

```
switch(config-mirror-endpoint-test) # comment Monitor statistics on each endpoint  
interfaces
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config-mirror- <i><SESSION-ID></i> config-mirror-endpoint	Administrators or local user group members with execution rights for this command.

destination interface

```
destination interface {<INTERFACE-ID>|<LAG-NAME>}  
no destination interface {<INTERFACE-ID>|<LAG-NAME>}
```

Description

Configures the specified interface as the destination of the mirrored traffic.

The **no** form of this command immediately disables the mirroring session and removes the specified destination interface from the configuration.

Parameter	Description
<INTERFACE-ID>	Specifies a interface. Format: member/slot/port.
<LAG-NAME>	Specifies a LAG (link aggregation group) identifier.

Usage

Configuring a different destination interface in an enabled mirroring session causes all mirrored traffic to use the new destination interface. This action might cause a temporary suspension of mirrored source traffic during the reconfiguration. Mirroring traffic to a LAG destination with a device running LACP attached to the destination will cause the LAG link to flap if LACP packets are part of the mirrored traffic source.

Examples

Configuring a mirroring session and adding an interface as a destination:

```
switch(config)# mirror session 1
switch(config-mirror-1)# destination interface 1/1/1
```

Replacing the existing destination with different interface:

```
switch(config-mirror-1)# destination interface 1/1/12
```

Removing a destination:

```
switch(config-mirror-1)# no destination interface 1/1/12
```

Switch	Destination interface limit per mirror session (4 possible sessions)
4100i	1
6000	1
6100	1

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config-mirror-<SESSION-ID>	Administrators or local user group members with execution rights for this command.

destination tunnel

Description

Specifies the tunnel where all mirrored traffic for the session is transmitted. Only one tunnel destination is allowed per session.

You may configure multiple mirror sessions with the same source/destination IP address pair, however, only one of those sessions sharing the same source/destination IP address pair can be enabled at a given time.

ERSPAN is not supported leaving the switch by the OOB port. If VRF management is configured for an ERSPAN session, the session will be in "mirror_err_tunnel_oob_port_not_supported" operation status.

ERSPAN is not supported leaving the switch encapsulated within another tunnel (e.g. GRE IPv4). When the path to the destination IP address will leave via a tunnel, the session will be in "tunnel_route_resolution_not_populated" operation status.



The interface/LAG used to transmit ERSPAN packets should not be a source in the same mirror session.

The **no** form of this command will cease the use of the tunnel and disable the session.

Parameter	Description
<TUNNEL-IPv4-ADDR>	Specifies the tunnel address in IPv4 format (x.x.x.x), where x is a decimal number from 0 to 255.
<SOURCE-IPv4-ADDR>	Specifies the source address in IPv4 format (x.x.x.x), where x is a decimal number from 0 to 255.
<DSCP-VALUE>	Specifies the DSCP value to be carried within the DS field of ERSPAN packet header. Range: 0 to 63. Default: 0.
<VRF-NAME>	Specifies a VRF name. Default: default.

Examples

Creating a Mirror Session and adding tunnel destination, source, dscp, and VRF:

```
switch# config
switch(config)# mirror session 1
switch(config-mirror-1)# destination tunnel 1.1.1.1 source 2.2.2.2 dscp 10 vrf default
```

Replacing the existing tunnel destination:

```
switch(config-mirror-1)# destination tunnel 11.12.13.14 source 2.2.2.2 dscp 10 vrf default
```

Replacing the existing destination with a different DSCP value:

```
switch(config-mirror-1)# destination tunnel 11.12.13.14 source 2.2.2.2 dscp 2 vrf default
```

Removing the destination:

```
switch(config-mirror-1)# no destination tunnel
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-mirror-<SESSION-ID>	Administrators or local user group members with execution rights for this command.

diagnostic

diagnostic

```
diag utilities tshark [file]
diag utilities tshark [delete-file]
```

Description

Captures packets from a mirror-to-cpu session, and save the most recent 32MB to pcap file which can then be copied and analyzed. When capturing a mirror-to-cpu session to a file, packets will not be dumped to the console.



The `diagnostic` command must be entered prior to the `diag utilities tshark` command.

Use the **delete-file** form of this command to delete the most recent capture file.

Since **file** and **delete-file** are optional, the behavior of the base command **diag utilities tshark** does **not** save anything to a file, and instead dumps the tshark session to the console until **CTRL + c** is entered.

Parameter	Description
file	Saves captured packets to a temporary file.
delete-file	Deletes the most recent captured file.

Example

Performing diagnostic:

```
switch# diagnostic

switch# diagnostic utilities tshark file
Inspecting traffic mirrored to the CPU until Ctrl-C is entered
^CEnding traffic inspection.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Manager (#)	Administrators or local user group members with execution rights for this command.

disable

disable

Description

Disables the mirroring session specified by the current command context.

Usage

By default, mirroring sessions are disabled.

When a mirroring session is disabled, the **show mirror** command for that session ID shows an **Admin Status** of **disable** and an **Operation Status** of **disabled**.

Example

Disabling a mirroring session:

```
switch(config)# mirror session 3
switch(config-mirror-3)# disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config-mirror- <i><SESSION-ID></i>	Administrators or local user group members with execution rights for this command.

enable

enable

Description

Enables the mirroring session for the current command context.

Usage

By default, mirroring sessions are disabled.

When a mirroring session is enabled, the **show mirror** command for that session ID shows an **Admin Status** of **enable** and an **Operation Status** of **enabled**.

If sFlow is enabled on an interface and a mirroring session specifies the same interface as the source of received traffic (the source is configured with a direction of **rx** or **both**):

- The attempt to enable the mirroring session fails and an error is returned.



When adding, removing, or changing the configuration of a source interface in an enabled mirroring session, packets from other mirror sources using the same destination interface might be interrupted.

Example

Configuring and enabling a mirroring session:

```
switch(config)# mirror session 3
switch(config-mirror-3)# source interface 1/1/2 rx
switch(config-mirror-3)# destination interface 1/1/3
switch(config-mirror-3)# comment Monitor router port ingress-only traffic
switch(config-mirror-3)# enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config-mirror- <i><SESSION-ID></i>	Administrators or local user group members with execution rights for this command.

mirror session

```
mirror session <SESSION-ID>
no mirror session <SESSION-ID>
```

Description

Creates a mirroring session configuration context or enters an existing mirroring session configuration context.

From this context, you can enter commands to configure and enable or disable the mirroring session.

The **no** form of this command removes an existing mirroring session from the configuration.

Parameter	Description
<i><SESSION-ID></i>	Specifies the session identifier. Range: 1 to 4

Examples

```
switch(config)# mirror session 1
switch(config-mirror-1)#
```

```
switch(config)# mirror session 3
switch(config-mirror-3)#

switch(config)# no mirror session 1
switch(config)#
```



When configuring mirroring via the command-line interface, not all configuration errors will result in an immediate error message. After making configuration changes, always check the operation status of your mirror sessions using the **show mirror** command.

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config	Administrators or local user group members with execution rights for this command.

mirror endpoint

```
mirror endpoint <NAME>
no mirror endpoint <NAME>
```

Description

Creates the specified mirror endpoint or enters its context if it already exists. The specifics of a mirror endpoint are created or altered while in the mirror endpoint context and the mirror endpoint is enabled or disabled from this context. It may be possible to support different encapsulations by different ASICs. For example, UDP for PVOS compatibility. Termination of GRE encapsulation is also supported. The **no** form of this command removes an existing mirror endpoint. An enabled mirror endpoint is automatically disabled first before removal.

Parameter	Description
<NAME>	Specifies mirror endpoint name.

Examples

Creating a mirror endpoint named test :

```
switch(config)# mirror endpoint test
```

Deleting mirror endpoint named test:

```
switch(config)# no mirror endpoint test
```

Configuring a mirror endpoint named test :

```
6100(config)# mirror endpoint test
6100(config-mirror-endpoint-test)#
6100(config-mirror-endpoint-test)# destination
    interface Specify interfaces to send traffic
6100(config-mirror-endpoint-test)# destination interface
    IFNAMELIST An interface, a range or a comma seperated list of interfaces
6100(config-mirror-endpoint-test)# destination interface 1/1/3
    <cr>
6100(config-mirror-endpoint-test)# destination interface 1/1/3
6100(config-mirror-endpoint-test)#
6100(config-mirror-endpoint-test)# source 1.1.1.1 destination 1.1.1.2 id 1 vrf
default
6100(config-mirror-endpoint-test)#
```



Only physical ports can be configured as interface for mirror-endpoint destination. LAG port is not supported as interface for mirror-endpoint destination.



The maximum allowed number of destination interfaces for both mirror-session and mirror-endpoint is 1.

Command History

Release	Modification
10.13.1000	Added support for 4100i, 6000, and 6100 switches.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config	Administrators or local user group members with execution rights for this command.

show mirror

```
show mirror [<SESSION-ID>]
```

Description

Shows information about mirroring sessions. If **<SESSION-ID>** is not specified, then the command shows a summary of all configured mirroring sessions. If **<SESSION-ID>** is specified, then the command shows detailed information about the specified mirroring session.

Parameter	Description
<SESSION-ID>	Specifies the session identifier. Range: 1 to 4

Usage

Information in the **Admin Status** column of the command output indicates the configured status. The admin status is one of the following values:

- **enable**: The mirroring session is enabled.
- **disable**: The mirroring session has been configured but not yet enabled, or has been disabled.

Information in the **Operation Status** column indicates the status of the mirroring session. Operation status is one of the following values:

- **dest_doesnt_exist**: The configured destination interface is not found in the system. The mirroring session cannot be enabled.
- **destination_shutdown**: **The mirroring session is enabled, but the destination interface is shut down. No traffic** can be monitored.
- **disabled**: The mirroring session is disabled and is not in an error condition.
- **enabled**: The mirroring session is enabled.
- **external/driver_error**: An internal ASIC hardware error occurred.
- **hit_active_sessions_capacity**: The mirroring session could not be enabled because the maximum number of supported mirroring sessions are already enabled.
- **internal_error**: An invalid parameter was passed to the ASIC software layer.
- **no_dest_configured**: The mirroring session does not have a destination interface configured.
- **no_name_configured**: A software error occurred. The mirroring session does not have a session ID in its configuration.
- **null_mirror**: A software error occurred. The session object reference is invalid.
- **out_of_memory**: The system is out of memory, reboot recommended.
- **tunnel_route_resolution_not_populated**: If the destination tunnel IP address is not reachable.
- **unknown_error**: An unexpected error occurred.

Examples

Showing summary information about all configured mirroring sessions:

```
switch# show mirror
ID  Admin Status  Operation Status
---  -
1   enable       enabled
2   disable      disabled
3   disable      disabled
4   enable       internal_error
```

Showing detailed information about a single mirroring session:

```
switch# show mirror 3
Mirror Session: 3
Admin Status: disable
```

```
Operation Status: disabled
Comment: Monitor router port ingress-only traffic
Source: interface 1/1/2 rx
Destination: interface 1/1/3
switch#
```

Show the details of mirror session 1 with an empty LAG as destination:

```
switch: show mirror 1
Admin Status: enable
Operation Status: dest_doesnt_exist
Source: interface 1/1/1 rx
Source: interface tx none
Destination: interface lag1
Output Packets: 0
Output Bytes: 0
```

Show the details of mirror session 1 with a LAG and a LAG member of same LAG as sources. In this scenario, traffic received on all LAG member interfaces is mirrored, is it is not necessary to define 1/1/1 as a source since it is already part of LAG1; this does not impact the mirror output packet or byte counters.

```
switch: show mirror 1
Admin Status: enable
Operation Status: enabled
Source: interface 1/1/1 rx
Source: interface lag1 rx
Destination: interface 1/1/2
Output Packets: 0
Output Bytes: 0
```

Showing the details of mirror **session 1** for a mirror configuration enabled with invalid interface which is not part of a lag:

```
switch# show mirror 1
Mirror Session: 1
Admin Status: enable
Operation Status: disabled
Source: interface 1/1/1 (unknown lag) both
Destination: cpu
Output Packets: 0
Output Bytes: 0
```

Certain configurations can cause this condition while the mirror source configuration remains unchanged:

1. Removing of a port from a LAG (Link Aggregation Group).
2. Port **split** or **unsplit** operations that result in LAG membership removal.

To recover from this issue, unconfigure the affected interface (previously part of a LAG), then restore the LAG membership configuration.

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show mirror endpoint

show mirror endpoint [<NAME>]

Description

Shows a list of all configured mirror endpoints, their Admin Status and their Operation Status. The optional parameter will display the details of the specified mirror endpoint if it exists.

Parameter	Description
<NAME>	Specifies name of the mirror endpoint instance to be displayed.

Examples

Showing a summary of all configured mirror endpoints on the switch:

```
switch# show mirror endpoint
Name      Admin Status  Operation Status
-----  -
test      enable        enabled
monitor   disable       disabled
```

Showing the details of enabled mirror endpoint test:

```
switch# show mirror endpoint test
Mirror Endpoint: audit
Admin Status: enable
Operation Status: enabled
Comment: Mirror Endpoint Audit
Type: gre
Tunnel: source 1.1.1.1 destination 1.1.1.2 id 1 vrf default
Interface: 1/1/3
Output Packets: 123456789
Output Bytes: 0
```



"Output Packets" in "show mirror endpoint [name]" is only supported for statistics.
 "Output Bytes" in "show mirror endpoint [name]" is not supported due to ASIC limitation.

Command History

Release	Modification
10.13.1000	Added support for 4100i, 6000, and 6100 switches.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

source interface

```
source interface {<PORT-NUM> | <LAG-NAME>} [<DIRECTION>]
no source interface {<PORT-NUM> | <LAG-NAME>} [<DIRECTION>]
```

Description

Configures the specified interface (either an Ethernet port or a LAG) as a source of traffic to be mirrored. The **no** form of this command ceases mirroring traffic from the specified source interface and removes the source interface from the mirroring session configuration.

Parameter	Description
<PORT-NUM>	Specifies a physical port on the switch. Use the format member/slot/port (for example, 1/3/1).
<LAG-NAME>	Specifies the identifier for the LAG (link aggregation group).
<DIRECTION>	Selects the direction of traffic to be mirrored from this source interface. There is no default for this parameter. Valid values are the following: <ul style="list-style-type: none">both: Mirror both transmitted and received packets.rx: Mirror only received packets.tx: Mirror only transmitted packets.

Usage

There is a limit of source interfaces in each direction of a given mirror session:

Switch	Source interface limit per mirror session (4 possible sessions)
4100i	28

Switch	Source interface limit per mirror session (4 possible sessions)
6000	52
6100	52

However, there is a practical limit to the amount of traffic that a mirror destination can transmit. For example, mirroring session with multiple 10G sources can overwhelm a single 10G destination.



When adding, removing, or changing the configuration of a source port in an enabled mirroring session, packets from other mirror sources using the same destination port might be interrupted.

Examples

Configuring a mirrored traffic source interface:

```
switch(config-mirror-1)# source interface
LAG-NAME      Enter a LAG name. For example, lag10
PORT-NUM      Enter a port number
```

Creating a mirroring session and configuring a source interface to mirror both transmitted and received packets:

```
switch(config)# mirror session 1
switch(config-mirror-1)# source interface 1/1/1 both
```

Creating a second mirroring session and configuring two source interfaces. One port mirroring only transmitted packets and the other mirroring both transmitted and received packets:

```
switch(config)# mirror session 2
switch(config-mirror-2)# source interface 1/1/3 tx
switch(config-mirror-2)# source interface 1/2/1 both
```

Removing the first source interface:

```
switch(config-mirror-2)# no source interface 1/2/3
```

Configuring a source interface to mirror received packets only:

```
switch(config-mirror-3)# source interface 1/1/2 rx
```

Configuring a source interface to mirror both transmitted and received packets:

```
switch(config-mirror-1)# source interface 1/1/1 both
```

Configuring a LAG as source interface to mirror both transmitted and received packets:

```
switch(config-mirror-4) # source interface lag1 both
```

Stopping the mirroring of received packets from a configured source interface:

```
switch(config-mirror-4) # no source interface lag1 rx
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config-mirror-<SESSION-ID>	Administrators or local user group members with execution rights for this command.

Monitoring a device using SNMP

Configuring SNMP: Refer to the *SNMP/MIB Guide* for information on how to add SNMP so a device can be monitored from a network management system (NMS).

Configuring an SNMP trap receiver: Refer to the *SNMP/MIB Guide* and specific information about the `show snmp trap` command to enable SNMP traps.

- The Power-over-Ethernet (PoE) subsystem manages power supplied to devices using standard Ethernet data cables. A Power Sourcing Equipment (PSE) supplies DC power as well as Ethernet connectivity to a Powered Device (PD) using a standard Ethernet cable. The maximum current depends on the PD Requested Class.
- A PoE subsystem contains two parts : a PSE and PD. A Power Sourcing Equipment (PSE) is a device that provides power through a standard Ethernet cable. A PoE capable switch functions as PSE. All Aruba PoE switches are considered as PSEs. A PD is a device powered by a PSE. Examples of PD are VoIP phones, Wireless APs, and IP cameras.
- When a PD or any network cable is connected to a PSE port, the PSE applies a detection voltage and measures the resistance value of the PD. If resistance is within IEEE 802.3 standard values (23 - 26k ohm), the connected device is treated as PD and classification begins. For legacy devices to be detected, you must enable prestandard detection on the switch.
- PDs are divided into different types and classes based on PD power requirements. The power supplied by the PSE is higher than the power PD draws to accommodate for the line losses that can result with the use of the standard maximum length cable(100m).
 - Type 1: PSE can supply maximum of 15.4W, and PD can draw a maximum of 13W.
 - Type 2: PSE can supply maximum of 30W, and PD can draw a maximum of 25.5W.
 - Type 3: PSE can supply maximum of 60W, and PD can draw a maximum of 51W.
 - Type 4: PSE can supply maximum of 90W, and PD can draw a maximum of 71W.
- Classes of PD:
 - Class 0: Type1 PD, it can draw a maximum of 13W.
 - Class 1: Type1 PD, it can draw a maximum of 3.84W.
 - Class 2: Type1 PD, it can draw a maximum of 6.49W.
 - Class 3: Type1 PD, it can draw a maximum of 13W.
 - Class 4: Type2 PD, it can draw a maximum of 25.5W.
 - Class 5: Type3 PD, it can draw a maximum of 40W.
 - Class 6: Type3 PD, it can draw a maximum of 51W.
 - Class 7: Type4 PD, it can draw a maximum of 62W.
 - Class 8: Type4 PD, it can draw a maximum of 71.3W.
- IEEE 802.3bt introduced 4-Pair PoE as a means of supplying higher power to PDs that need more than the current 25.5W supplied by IEEE 802.3at. To increase the available power without damaging the Ethernet cable, the standard introduced the ability to use all four pairs within the Ethernet cable instead of the two pairs used by previous standards (802.3at, 802.3af).
- Supported protocols:
 - Compatibility with IEEE 802.3af, 802.3at, 802.3bt and prestandard.
 - Long first class event supported on Type 3-4 PSE.
 - Support for Single Signature (SS) Type 0-6 and Dual Signature (DS) Type 0-4 PDs.
 - Multi-Event classification permits mutual ID of SS Class 0-6 and DS Class 0-4.

- Support LLDP Data Link Layer (DLL) Type 1-2 extension 12-octet TLV and Type 3-4 extension 29-octet TLV.
- Default PSE assigned class delivers the maximum PSE capable power at initial power up based on PD requested class.
- Always-on PoE is a feature that provides the ability for a switch to continue to provide power across user initiated reboots through software. Always-on PoE is enabled by default and no additional configuration is needed.



PDs only remain powered, no data transfer or PoE power negotiation can occur until the switch has completely booted up and in normal operation. PD faults occurring prior to full switch boot up will result in PoE power removal and restart the detection process only after switch returns to normal operation.

PoE commands

All PoE configuration commands except **quick-poe**, **threshold configuration** and **always-on poe configuration** are entered at the **config-if** context. The PoE threshold command is used at the system level whereas the **always-on poe** and **power-over-ethernet quick-poe** commands are set at the slot level. These commands can only be configured in the global configuration context.

lldp dot3 poe

```
lldp dot3 poe
no lldp dot3 poe
```

Description

Enables 802.3 TLV list in LLDP to advertise for Power over Ethernet Data Link Layer Classification. LLDP dot3 TLV is by default enabled for PoE.

The **no** form of this command disables 802.3 TLV list in LLDP.

Examples

Enabling 802.3 TLV list in LLDP:

```
switch(config)# interface 1/1/1
switch(config-if)# lldp dot3 poe
```

Disabling 802.3 TLV list in LLDP:

```
switch(config-if)# no lldp dot3 poe
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

lldp med poe

```
lldp med poe [priority-override]
no lldp med poe [priority-override]
```

Description

Enables MED TLV list in LLDP to advertise for Power over Ethernet Data Link Layer Classification. Also enables the lldp-MED TLV priority to override user configured port priority for Power over Ethernet. When both dot3 and MED are enabled, dot 3 will take precedence. MED TLV is by default enabled for PoE. Priority over-ride is by default disabled.

The **no** form of this command disables MED TLV list in LLDP.

Parameter	Description
[priority-override]	System defined name of the interface.

Examples

Enabling and disabling LLDP MED PoE:

```
switch(config)# interface 1/1/1
switch(config-if)# lldp med poe
switch(config-if)# no lldp med poe
```

Enabling and disabling LLDP MED PoE priority override:

```
switch(config-if)# lldp med poe priority-override
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

power-over-ethernet

```
power-over-ethernet
no power-over-ethernet
```

Description

Enables per-interface power distribution. Per-port power is enabled by default with priority low. PoE cannot be disabled for individual ports when Quick PoE is enabled for the entire switch or line module. The **no** form of this command disables per-interface power distribution.

Examples

Enabling per-interface power distribution:

```
switch(config)# interface 1/1/1  
switch(config-if)# power-over-ethernet
```

Disabling per-interface power distribution:

```
switch(config-if)# no power-over-ethernet
```

Showing Quick PoE enabled:

```
switch(config)# power-over-ethernet quick-poe 1/1  
switch(config)# no power-over-ethernet  
Interface PoE cannot be disabled when Quick PoE is enabled.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

power-over-ethernet allocate-by

```
power-over-ethernet allocate-by {usage | class}  
no power-over-ethernet allocate-by {usage | class}
```

Description

Configures the power allocation method. Power allocation method is initially based on usage. PSE Allocated power value will change to LLDP negotiated power if and when LLDP exchange takes place between PSE and PD. When there is no LLDP negotiation, PSE Allocated Power Value will be the actual instantaneous power draw and reserve power based on actual consumption. In allocate-by class, power allocation is based on PD requested class and PSE allocated power value will be the LLDP negotiated power when LLDP exchange takes place between PSE and PD. When there is no LLDP negotiation, PSE Allocate Power will be based on PD class. Reserve power is based on PD Class. By default, power allocation is by usage.

The power allocation method can be changed on an interface through port-access (User roles or RADIUS). An allocation method when configured through port-access will replace the user configured method.

The **no** form of this command resets the action to default.

Parameter	Description
usage	Configures the usage-based allocation method.
class	Configures the class-based allocation method.

Usage

If you enable **pd-class-override** for an interface, the **allocate-by** configuration of that interface will be automatically changed to **class**. However, if you change the allocation method to **usage** when **pd-class-override** is still enabled, you will receive an error message stating that "The power allocation method cannot be changed when pd-class-override is enabled."

To remove **pd-class-override**, you can use the **no power-over-ethernet pd-class-override** command. It is important to note that **pd-class-override** requires the allocation method to be set to **class** and is enforced when configured through CLI. However, if you override the allocation method to **usage** via port-access, **pd-class-override** will not be in effect. Therefore, it is recommended that you do not override the allocation method to **usage** through port-access on interfaces configured with **pd-class-override**.

Examples

Configuring the power allocation method:

```
switch(config)# interface 1/1/1
switch(config-if)# power-over-ethernet allocate-by usage
switch(config-if)# power-over-ethernet allocate-by class
```

Resetting power allocation method:

```
switch(config-if)# no power-over-ethernet allocate-by class
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

power-over-ethernet always-on

```
power-over-ethernet always-on <MODULE-ID>
no power-over-ethernet always-on <MODULE-ID>
```



This command is not available on the 6000 or 6100 Switch Series.



This command is available only on the 4100i Switch Series.

Description

Always-on PoE is a feature that provides the ability to the switch to continue to provide power across a soft reboot. It is applicable only to the interfaces which were connected and delivering before the soft reboot. Also, power will not be delivered if power to the switch is interrupted. This command enables or disables the always-on PoE feature at the switch or the slot level. By default, always-on PoE is enabled at the switch or the slot level.

The **no** form of this command disables power distribution on soft reboot.

Parameter	Description
<MODULE-ID>	Module number to apply always-on PoE configuration.

Examples

Enabling per-interface power distribution:

```
switch(config)# power-over-ethernet always-on 1/1
```

Disabling per-interface power distribution:

```
switch(config)# no power-over-ethernet always-on 1/1
```

Command History

Release	Modification
10.08	Command introduced

Command Information

Platforms	Command context	Authority
4100i	config	Administrators or local user group members with execution rights for this command.

power-over-ethernet assigned-class

```
power-over-ethernet assigned-class {3 | 4 | 6}
no power-over-ethernet assigned-class
```

Description

Limit PoE power based on the assigned class. When an user assigns a maximum class to an interface, the PSE will limit the maximum power delivered to the PD up to a total power draw not exceeding the PSE assigned-class power. Power demotion occurs when a PD requested class is higher than the PSE assigned class, permitting the PD to receive power and operate in a reduced power mode. PoE ports cannot set an assigned class when Quick PoE is enabled on the sybsystem. The default assigned class is 4 for 2-pair capable PSE and 6 for 4-pair capable PSE.

The **no** form of this command resets the action to default.

Examples

Setting PoE assigned class:

```
switch(config)# interface 1/1/1  
switch(config-if)# power-over-ethernet assigned-class 4
```

Resetting PoE assigned class to default:

```
switch(config-if)# no power-over-ethernet assigned-class 4
```

Showing Quick PoE enabled:

```
switch(config)# power-over-ethernet quick-poe 1/1  
switch(config)# interface 1/1/1  
switch(config)# power-over-ethernet assigned-class 4  
Interface assigned class cannot be configured when Quick PoE is enabled.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

power-over-ethernet pre-std-detect

```
power-over-ethernet pre-std-detect  
no power-over-ethernet pre-std-detect
```

Description

Before IEEE 802.3 released the first Power over Ethernet standard (802.3af), vendors had shipped PoE capable switches and PD's. HPE Aruba Networking switches are backward compatible and will support both IEEE standard and pre-standard 802.3af Power over Ethernet PD's concurrently. This CLI allows the user to enable or disable pre-802.3af-standard device detection and powering on the specific port. When pre-std-detect is enabled, power will be delivered on PairA only. Default is disabled.

The **no** form of this command resets the action to default.

Examples

Enabling pre-standard device detection:

```
switch(config)# interface 1/1/1  
switch(config-if)# power-over-ethernet pre-std-detect
```

Disabling pre-standard device detection:

```
switch(config-if)# no power-over-ethernet pre-std-detect
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

power-over-ethernet priority

```
power-over-ethernet priority {critical | high | low}  
no power-over-ethernet priority {critical | high | low}
```

Description

Sets PoE priority for an interface. Specifying critical, high, or low indicates the priority of the interface in the event of power over-subscription. Within the same priority level, higher power-priority line-module ports have higher precedence. With same PoE priority and same line-module priority, lower numbered line-module ports have higher precedence. Per-interface PoE priority is low by default.

The **no** form of this command resets the priority to default PoE priority "low".

Examples

Configuring PoE priority:

```
switch(config)# interface 1/1/1  
switch(config-if)# power-over-ethernet priority critical  
switch(config-if)# power-over-ethernet priority high
```

Resetting the PoE priority to default:

```
switch(config-if)# no power-over-ethernet priority high
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

power-over-ethernet quick-poe

power-over-ethernet quick-poe <MODULE-ID>
no power-over-ethernet quick-poe <MODULE-ID>



This command is not available on the 6000 or 6100 Switch Series.



This command is available only on the 4100i Switch Series.

Description

Quick PoE is a feature that provides the ability for the switch to provide power to the connected powered device as soon as switch goes through cold reboot. When quick PoE is enabled on the subsystem PoE port disablement and PD demotion is not allowed. also quick PoE enablement is not allowed if any of the port is disabled on the subsystem. User should not over-subscribe the PoE power when quick PoE is enabled. Quick PoE saved configuration will work irrespective of the configuration change at reboot.

Enables quick PoE feature on the switch or the subsystem level. By default, quick-PoE is disabled for the subsystem.

The **no** form of this command disables quick PoE.

Parameter	Description
<MODULE-ID>	Specifies module number for quick PoE configuration .

Examples

Enabling and disabling quick PoE:

```
switch(config)# power-over-ethernet quick-poe 1/2
switch(config)# no power-over-ethernet quick-poe 1/2
```

```
switch(config-if)# power-over-ethernet quick-poe 1/1
PoE must be enabled on all interfaces before enabling Quick PoE
```

```
switch(config-if)# power-over-ethernet quick-poe 1/3  
All interfaces must use the default assigned class before enabling Quick PoE
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i	config-if	Administrators or local user group members with execution rights for this command.

power-over-ethernet threshold

```
power-over-ethernet threshold <PERCENTAGE>  
no power-over-ethernet threshold <PERCENTAGE>
```

Description

Sets the threshold at which the system will send an excess power consumption notification trap. Default value is 80 percentage.

The **no** form of this command resets the action to default.

Parameter	Description
<PERCENTAGE>	Excess power consumption trap threshold. Range 1-99.

Examples

Setting the power-over-ethernet threshold:

```
switch(config)# power-over-ethernet threshold 75
```

Resetting the power-over-ethernet threshold to default:

```
switch(config-if)# no power-over-ethernet threshold 75
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config	Administrators or local user group members with execution rights for this command.

power-over-ethernet trap

```
power-over-ethernet trap
no power-over-ethernet trap
```

Description

This command enables/disables the SNMP trap generation for PoE related events at system level. PoE trap generation is enabled by default.

The **no** form of this command resets the priority to default PoE priority "low".

Examples

Enabling SNMP trap generation for PoE:

```
switch(config)# power-over-ethernet trap
```

Disabling SNMP trap generation for PoE:

```
switch(config-if)# no power-over-ethernet trap
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	config-if	Administrators or local user group members with execution rights for this command.

show lldp local

```
show lldp local-device [<INTERFACE-ID>]
```

Description

Displays information advertised by the switch if the LLDP feature is enabled by user.

Parameter	Description
<INTERFACE-ID>	Specifies an interface. Format: member/slot/port

Examples

Showing LLDP local device:

```
switch# show lldp local-device 1/1/10
Local Port Data
=====

Port-ID           : 1/1/10
Port-Desc         : "1/1/10"
Port VLAN ID     : 0

PoE Plus Information

PoE Device Type   : Type 2 PSE
Power Source      : Primary
Power Priority     : low
PSE Allocated Power: 25.0 W
PD Requested Power : 25.0 W
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show lldp neighbor

```
show lldp neighbor [<INTERFACE-ID>]
```

Description

Displays detailed information about a particular neighbor connected to a particular interface.

Parameter	Description
<INTERFACE-ID>	Specifies an interface. Format: member/slot/port

Examples

Showing LLDP neighbor information when there is only one neighbor:

```
switch# show lldp neighbor-info 1/1/10

Port           : 1/1/10
Neighbor Entries : 1
Neighbor Entries Deleted : 0
Neighbor Entries Dropped : 0
```

```

Neighbor Entries Aged-Out      : 0
Neighbor Chassis-Name         : 84:d4:7e:ce:5d:68
Neighbor Chassis-Description  : ArubaOS (MODEL: 325), Version HPE_ANW IAP
Neighbor Chassis-ID           : 84:d4:7e:ce:5d:68
Neighbor Management-Address    : 169.254.41.250
Chassis Capabilities Available : Bridge, WLAN
Chassis Capabilities Enabled   :
Neighbor Port-ID              : 84:d4:7e:ce:5d:68
Neighbor Port-Desc            : eth0
TTL                            : 120
Neighbor Port VLAN ID         :
Neighbor PoEplus information   : DOT3
Neighbor Device Type          : TYPE2 PD
Neighbor Power Priority        : Unkown
Neighbor Power Source         : Primary
Neighbor Power Requested      : 25.0 W
Neighbor Power Allocated      : 0.0 W
Neighbor Power Supported      : No
Neighbor Power Enabled        : No
Neighbor Power Class          : 5
Neighbor Power Paircontrol    : No
Neighbor Power Pairs          : SIGNAL

```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

show power-over-ethernet

show power-over-ethernet [member <MEMBER-ID>] [brief]

Description

Displays the status information of the full system.

Parameter	Description
<MEMBER-ID>	Displays the detailed status of given member.
brief	Display the brief status of all ports or the given port.

Examples

Showing sample output for show power-over-ethernet :

```

switch# show power-over-ethernet

System Power Status

PoE Power Status           : No redundancy
Operational Power Status   : No redundancy
Total Available Power      : 360 W
Total Configured Power     : 0 W
Total Failover Pwr Avl    : 0 W
Total Redundancy Power     : 0 W
Total Power Drawn         : 0 W
Total Power Reserved       : 0 W
Total Remaining Power      : 360 W
Trap Threshold             : 80 %
Trap Enabled               : Yes
Always-on PoE Enabled      : 1/1
Quick PoE Enabled         : None

```

```

Internal Power

PS      Total Power      Status
-----
1/1     0                 Ok
1/2     0                 Absent

```

Showing sample output for power-over-ethernet brief per-port:

```

switch# show power-over-ethernet 1/1/1 brief

Status and Configuration Information for port 1/1/1

Member 1Power Status
  Available: 370 W Reserved: 55.60 W Remaining: 314.40 W
  Always-on PoE Enabled: 1/1
PoE      Pwr Power      Pre-std Alloc PSE Pwr PD Pwr PoE Port      PD      Cls Type
Port     En  Priority Detect   Act  Rsrvd Draw  Status Sign
-----
1/1/1   Yes Low      Off     Class  0.0 W  0.0 W Denied  None   4    2

```

Showing sample output for power-over-ethernet brief for interface range:

```

switch# show power-over-ethernet 1/1/1-1/1/2 brief

Status and Configuration Information for port 1/1/1-1/1/2

Power Status
  Available: 360 W Reserved: 0.00 W Remaining: 360.00 W
  Always-on PoE Enabled: 1/1
  Quick PoE Enabled: None
PoE      Pwr Power      Pre-std Alloc PSE Pwr PD Pwr PoE Port      PD      Cls Type
Port     En  Priority Detect   Act  Rsrvd Draw  Status Sign
-----
1/1/1   Yes Low      Off     Usage  0.0 W  0.0 W Searching N/A   N/A N/A
1/1/2   Yes Low      Off     Usage  0.6 W  0.0 W Searching N/A   N/A N/A

```

Showing sample output for power-over-ethernet for a missing line card:

```
switch# show power-over-ethernet 1/3 brief
```

```
Module 1/3 is not physically present.
```

Showing sample output for power-over-ethernet port when physical interface is not present:

```
switch# show power-over-ethernet 2/1/1
```

```
Interface 2/1/1 is not present.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
4100i 6000 6100	Operator (>) or Manager (#)	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

You can manage and monitor the AOS-CX switch through Aruba AirWave. The following benefits and functions include:

- Configuration (partial configuration)
- Device topology
- Immediate and historical trend reports
- Monitoring of the device and user connected to the network.
- Network discovery
- Syslogs and trap receiver

For information about which versions of Aruba AirWave support AOS-CX, see the *AOS-CX Release Notes*.

SNMP support and AirWave

For AirWave to discover and monitor the switch, you must:

- Enable the SNMP services on the switch.
- Configure the SNMP agent to use the SNMP version supported by the management station.

SNMP on the switch

The switch provides SNMP services through the management channel and the data interfaces. Functionality, such as device discovery from NMS, syslog and trap forwarding, can be any channel configured by you.

Although the SNMP server can be enabled on both VRFs (`mgmt` and `default`), only one instance of SNMP can be running. The highest priority is on the `default` VRF.

For example, assume that SNMP is first enabled on the `mgmt` VRF (`snmp-server vrf mgmt`). Then, SNMP is enabled on the `default` VRF (`snmp-server vrf default`) without disabling SNMP on the `mgmt` (using an equivalent `no` form of the command). The `show running-config` command displays both `snmp-server vrf` commands; however, the SNMP instance is running only on the `default` VRF (highest priority).

```
switch# config
switch(config)# snmp-server vrf mgmt
switch(config)# snmp-server vrf default
switch(config)# show running-config
Current configuration:
!
!Version AOS-CX Virtual.10.01.
led locator on
!
!
!
snmp-server vrf default
```

```
snmp-server vrf mgmt
!  
...
```

Supported features with AirWave and the AOS-CX switch

AirWave supports the following features with the AOS-CX switch:

Device management	Device discovery using SNMPv2C and SNMPv3
	Device dashboards
Monitoring management	Device health attributes (device status/reachability)
	Interface and VLAN management
	Initiates an SSH connection from Aruba AirWave to AOS-CX so that the device outputs from the AOS-CX CLI can be displayed in the Aruba AirWave user interface.
	Firmware versions
	Displays neighbor devices connected to AOS-CX switches
Device topology	
Configuration management	Partial configuration
Alarm management	Alarm triggers (device and interface up/down, new device discoveries, custom event triggers)
	Syslogs and traps
Report management	Device inventory, interface utilization, and device reachability reports
	Summary report of device model, firmware, and boot loader version

Configuring the AOS-CX switch to be monitored by AirWave

Prerequisites

Aruba AirWave is active on the network.

Procedure

1. Enable SNMP on the switch by entering the `snmp-server vrf` command.

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

2. Configure the SNMPv2C community to public by entering the `snmp-server community public` command. In this instance, `public` is a read-only community string.

```
switch(config)# snmp-server community public
```

3. The community-string is used by SNMPv1 and SNMPv2C for unencrypted authentication. SNMPv3 lets you encrypt the authentication mechanism. To enable SNMPv3, enter the `snmpv3 user` and `snmpv3 context` commands.

```
switch(config)# snmpv3 user Admin auth sha auth-pass ciphertext
AQBapZHf2d20GYr/xcGUzYzm0zjNf/4VKHtSqbNImqtfYbJYCgAAALkGFJVcSp3nZ3o=
priv des priv-pass ciphertext
AQBapb0H2poBQKXPoVsC9L9qzZyfJQnzR7hmTr7LGsOsI7K3CgAAAKP98Rq2jfTrFwQ=

switch(config)# snmpv3 context Admin
```

For discovering devices in AirWave through the SNMPv3 community, the SNMPv3 context name is not mandatory. Devices can still be discovered in Aruba AirWave without the SNMPv3 context name.

4. Enter the `logging` command for enabling syslog forwarding to a remote syslog server, such as AirWave:

```
switch(config)# logging 10.0.10.2 severity debug
```

5. SNMP traps enable an agent to notify the management station of significant events by way of an unsolicited SNMP message. Enable SNMP traps by entering the `snmp-server host` command:

```
switch(config)# snmp-server host 10.10.10.10 trap version v2c vrf default
```

SNMP traps cannot be forwarded from AOS-CX 10.00 switches that have the VRF configured as `mgmt`. Later versions of AOS-CX support SNMP trap forwarding even when the VRF is configured as `default` or `mgmt`.

6. For information on how to add a device for monitoring in the Aruba AirWave user interface, see the documentation for Aruba AirWave.

AirWave commands

logging

```
logging {<IPV4-ADDR> | <IPV6-ADDR> | <FQDN | HOSTNAME>} [ {udp [<PORT-NUM>] }|{tcp
[<PORT-NUM>] | {tls [<PORT-NUM> [auth-mode {certificate|subject-name}] [legacy-tls-
renegotiation]]} [severity <LEVEL>] [vrf <VRF-NAME>] [include-auditable-events]
[filter <FILTER-NAME>] [ rate-limit-burst <BURST> [rate-limit-interval <INTERVAL>] ]
```

```
no logging {<IPV4-ADDR> | <IPV6-ADDR> | <HOSTNAME>}
```

Description

Enables syslog forwarding to a remote syslog server.

The `no` form of this command disables syslog forwarding to a remote syslog server.

Parameter	Description
{<IPV4-ADDR> <IPV6-ADDR> <HOSTNAME>}	Selects the IPv4 address, IPv6 address, or host name of the remote syslog server. Required.
[udp [<PORT-NUM>] tcp [<PORT-NUM>]]	Specifies the UDP port or TCP port of the remote syslog server to receive the forwarded syslog messages.
udp [<PORT-NUM>]	Range: 1 to 65535. Default: 514
tcp [<PORT-NUM>]	Range: 1 to 65535. Default: 1470
tls [<PORT-NUM>]	Range: 1 to 65535. Default: 6514
include-auditable-events	Specifies that auditable messages are also logged to the remote syslog server.
severity <LEVEL>	Specifies the severity of the syslog messages: <ul style="list-style-type: none"> ▪ alert: Forwards syslog messages with the severity of alert (6) and emergency (7). ▪ crit: Forwards syslog messages with the severity of critical (5) and above. ▪ debug: Forwards syslog messages with the severity of debug (0) and above. ▪ emerg: Forwards syslog messages with the severity of emergency (7) only. ▪ err: Forwards syslog messages with the severity of err (4) and above ▪ info: Forwards syslog messages with the severity of info (1) and above. Default. ▪ notice: Forwards syslog messages with the severity of notice (2) and above. ▪ warning: Forwards syslog messages with the severity of warning (3) and above.
auth-mode	Specifies the TLS authentication mode used to validate the certificate. <ul style="list-style-type: none"> ▪ certificate: Validates the peer using trust anchor certificate based authentication. Default. ▪ subject-name: Validates the peer using trust anchor certificates as well as subject-name based authentication.
legacy-tls-renegotiation	Enables the TLS connection with a remote syslog server supporting legacy renegotiation.
vrf <VRF-NAME>	Specifies the VRF used to connect to the syslog server. Optional. Default: default

Examples

Enabling the syslog forwarding to remote syslog server 10.0.10.2:

```
switch(config)# logging 10.0.10.2
```

Enabling the syslog forwarding of messages with a severity of `err` (4) and above to TCP port 4242 on remote syslog server 10.0.10.9 with VRF `lab_vrf`:

```
switch(config)# logging 10.0.10.9 tcp 4242 severity err vrf lab_vrf
```

Disabling syslog forwarding to a remote syslog server:

```
switch(config)# no logging
```

Enabling syslog forwarding over TLS to a remote syslog server using subject-name authentication mode:

```
switch(config)# logging example.com tls auth-mode subject name
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	<code>config</code>	Administrators or local user group members with execution rights for this command.

snmp-server community

```
snmp-server community <STRING>  
no snmp-server community <STRING>
```

Description

Adds an SNMPv1/SNMPv2c community string. A community string is a password that controls read access to the SNMP agent. A network management program must supply this name when attempting to get SNMP information from the switch. A maximum of 10 community strings are supported. Once you create your own community string, the default community string (`public`) is deleted.

The `no` form of this command removes the specified SNMPv1/SNMPv2c community string. When no community string exists, a default community string with the value `public` is automatically defined.

Parameter	Description
<STRING>	Specifies the SNMPv1/SNMPv2c community string. Range: 1 to 32 printable ASCII characters, excluding space and question mark.

Examples

Setting the SNMPv1/SNMPv2c community string to **private**:

```
switch(config)# snmp-server community private
```

Removing SNMPv1/SNMPv2c community string **private**:

```
switch(config)# no snmp-server community private
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config	Administrators or local user group members with execution rights for this command.

snmp-server host

```
snmp-server host <IPv4-ADDR> trap version <VERSION> [community <STRING>]
[port <UDP-PORT>] [vrf <VRF-NAME>]
no snmp-server host <IPv4-ADDR> trap version <VERSION> [community <STRING>]
[port <UDP-PORT>] [vrf <VRF-NAME>]
snmp-server host <IPv4-ADDR> inform version v2c [community <STRING>]
[port <UDP-PORT>] [vrf <VRF-NAME>]
no snmp-server host <IPv4-ADDR> inform version v2c [community <STRING>]
[port <UDP-PORT>] [vrf <VRF-NAME>]
snmp-server host <IPv4-ADDR> [trap version v3 | inform version v3] user <NAME>
[port <UDP-PORT>] [vrf <VRF-NAME>]
no snmp-server host <IPv4-ADDR> [trap version v3 | inform version v3] user <NAME>
[port <UDP-PORT>] [vrf <VRF-NAME>]
```

Description

Configures a trap/informs receiver to which the SNMP agent can send SNMP v1/v2c/v3 traps or v2c informs. A maximum of 30 SNMP traps/informs receivers can be configured.

The `no` form of this command removes the specified trap/inform receiver.



Configuring `snmpv3` informs is not supported.

Parameter	Description
<IPv4-ADDR>	Specifies the IP address of a trap receiver in IPv4 format (x.x.x.x), where x is a decimal number from 0 to 255. You can remove leading zeros. For example, the address 192.169.005.100 becomes 192.168.5.100.
trap version <VERSION>	Specifies the trap notification type for SNMPv1 or v2c. Available options are: v1 or v2c.
inform version v2c	Specifies the inform notification type for SNMPv2c.
trap version v3	Specifies the trap notification type for SNMPv3.

Parameter	Description
<code>user <NAME></code>	Specifies the SNMPv3 user name to be used in the SNMP trap notifications.
<code>community <STRING></code>	Specifies the name of the community string to use when sending trap notifications. Range: 1 - 32 printable ASCII characters, excluding space and question mark. Default: <code>public</code> .
<code><UDP-PORT></code>	Specifies the UDP port on which notifications are sent. Range: 1 - 65535. Default: 162.
<code>vrf <VRF-NAME></code>	Specifies the name of the VRF on which to send the notifications.

Examples

```

switch(config)# snmp-server host 10.10.10.10 trap version v1
switch(config)# no snmp-server host 10.10.10.10 trap version v1
switch(config)# snmp-server host 10.10.10.10 trap version v2c community public
switch(config)# no snmp-server host 10.10.10.10 trap version v2c community public
switch(config)# snmp-server host 10.10.10.10 trap version v2c community public
port 5000
switch(config)# no snmp-server host 10.10.10.10 trap version v2c community public
port 5000
switch(config)# snmp-server host 10.10.10.10 trap version v2c community public
port 5000 vrf default
switch(config)# no snmp-server host 10.10.10.10 trap version v2c community public
port 5000 vrf default
switch(config)# snmp-server host 10.10.10.10 inform version v2c community public
switch(config)# no snmp-server host 10.10.10.10 inform version v2c community
public
switch(config)# snmp-server host 10.10.10.10 inform version v2c community public
port 5000
switch(config)# no snmp-server host 10.10.10.10 inform version v2c community
public port 5000
switch(config)# snmp-server host 10.10.10.10 inform version v2c community public
port 5000 vrf default
switch(config)# no snmp-server host 10.10.10.10 inform version v2c community
public port 5000 vrf default

switch(config)# snmp-server host 10.10.10.10 trap version v3 user Admin
switch(config)# no snmp-server host 10.10.10.10 trap version v3 user Admin
switch(config)# snmp-server host 10.10.10.10 trap version v3 user Admin port 2000
switch(config)# no snmp-server host 10.10.10.10 trap version v3 user Admin port
2000

```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	<code>config</code>	Administrators or local user group members with execution rights for this command.

snmp-server vrf

```
snmp-server vrf <VRF-NAME>  
no snmp-server vrf <VRF-NAME>
```

Description

Configures the VRF on which the SNMP agent listens for incoming requests. By default, the SNMP agent does not listen on any VRF.

The `no` form of this command stops the SNMP agent from listening for incoming requests on the specified VRF.

Parameter	Description
<VRF-NAME>	Specifies the VRF on which the SNMP agent listens for incoming requests. The SNMP agent can listen on either the <code>mgmt</code> or <code>default</code> VRF. If configured for both, the SNMP agent listens on <code>default</code> , which has a higher priority.

Example

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

```
switch(config)# no snmp-server vrf default
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	<code>config</code>	Administrators or local user group members with execution rights for this command.

snmpv3 context

```
snmpv3 context <NAME> vrf <VRF-NAME> [community <STRING>]  
no snmpv3 context <NAME> [vrf <VRF-NAME>]
```

Description

Creates an SNMPv3 context on the specified VRF.

The `no` form of this command removes the specified SNMP context.

Parameter	Description
<NAME>	Specifies the name of the context. Range: 1 to 32 printable ASCII

Parameter	Description
	characters, excluding space and question mark (?).
<code>vrf <VRF-NAME></code>	Specifies the VRF associated with the context. Default: default.
<code>community <STRING></code>	Specifies the SNMP community string associated with the context. Range: 1 to 32 printable ASCII characters, excluding space and question mark. Default: public.

Examples

Creating an SNMPv3 context named **newContext**:

```
switch(config)# snmpv3 context newContext
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	<code>config</code>	Administrators or local user group members with execution rights for this command.

snmpv3 user

```
snmpv3 user <NAME> [auth <AUTH-PROTOCOL> auth-pass {plaintext | ciphertext}
<AUTH-PWORD> [priv <PRIV-PROTOCOL> priv-pass {plaintext | ciphertext} <PRIV-PWORD>] ]
no snmpv3 user <NAME> [auth <AUTH-PROTOCOL> auth-pass
<AUTH-PWORD> [priv <PRIV-PROTOCOL> priv-pass <PRIV-PWORD>] ]
```

Description

Creates an SNMPv3 user and adds it to an SNMPv3 context.

The `no` form of this command removes the specified SNMPv3 user.

Parameter	Description
<code><NAME></code>	Specifies the SNMPv3 username. Range 1 - 32 printable ASCII characters, excluding space and question mark.
<code>auth <AUTH-PROTOCOL></code>	Specifies the authentication protocol used to validate user logins. Available options are: <code>md5</code> or <code>sha</code> .
<code>auth-pass {plaintext ciphertext} <AUTH-PWORD></code>	Specifies the SNMPv3 user password. Range for <code>plaintext</code> is 8 - 32 printable ASCII characters, excluding space and question

Parameter	Description
	mark. Range for <code>ciphertext</code> is 1 - 120 printable ASCII characters. This option is only used when copying user configuration settings between switches. It enables you to duplicate a user's configuration on another switch without having to know their password.
<code>priv <PRIV-PROTOCOL></code>	Specifies the SNMPv3 security protocol (encryption method). Available options are: <code>aes</code> or <code>des</code> .
<code>priv-pass {plaintext ciphertext} <PRIV-PWORD></code>	Specifies the SNMPv3 user privacy passphrase. Range for <code>plaintext</code> is 8 - 32 printable ASCII characters, excluding space and question mark. Range for <code>ciphertext</code> is 1 - 120 printable ASCII characters. This option is only used when copying user configuration settings between switches. It enables you to duplicate a user's configuration on another switch without having to know their password.

Examples

Defining an SNMPv3 user named **Admin** using **sha** authentication with the plaintext password **mypassword** and using **des** security with the plaintext password **myprivpass**:

```
switch(config)# snmpv3 user Admin auth sha auth-pass plaintext mypassword priv des
priv-pass plaintext myprivpass
```

Removing an SNMPv3 user named `Admin`:

```
switch(config)# no snmpv3 user Admin
```

Defining an SNMPv3 user named **Admin** using **sha** authentication with the plaintext password **mypassword** and using **des** security with the plaintext password **myprivpass**:

```
switch(config)# snmpv3 user Admin auth sha auth-pass plaintext mypassword priv des
priv-pass plaintext myprivpass
```

Copying an SNMP user from switch 1 to switch 2.

On switch 1, configure a user called **Admin**, then issue the `show running-config` command to display switch configuration settings. The `snmpv3 user` command uses the `ciphertext` option to protect the users's passwords.

```
switch1(config)# snmpv3 user Admin auth sha auth-pass plaintext mypassword
priv des priv-pass plaintext myprivpass
switch1(config)# exit
switch1# show running-config
Current configuration:
```

```

!
!Version AOS-CX TL.10.00.0003-8017-gdeb0606~dirty
!
!
!
snmpv3 user Admin auth sha auth-pass ciphertext
AQBapZHf2d20GYr/xcGUzYzm0zjNf/4VKHtSqbNImqtfYbJYCgAAALkGFJVcSp3nZ3o=
priv des priv-pass ciphertext
AQBapb0H2poBQKXPoVsC9L9qzZyfJQnzR7hmTr7LGsOsI7K3CgAAAKP98Rq2jfTrFwQ=
ssh server vrf mgmt
!
!
!
!
interface mgmt
    no shutdown
    ip dhcp
vlan 1

```

On switch 2, execute the snmpv3 user command that was displayed by `show running-config` on switch 1. This creates the user on switch 2 with the same configuration settings.

```

switch1(config)# snmpv3 user Admin auth sha auth-pass
ciphertextAQBapZHf2d20GYr/xcGUzYzm0zjNf/4VKHtSqbNImqtfYbJYCgAAALkGFJVcSp3nZ3o=priv
des priv-pass ciphertext
AQBapb0H2poBQKXPoVsC9L9qzZyfJQnzR7hmTr7LGsOsI7K3CgAAAKP98Rq2jfTrFwQ=

```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	config	Administrators or local user group members with execution rights for this command.

Accessing HPE Aruba Networking Support

HPE Aruba Networking Support Services	https://www.hpe.com/us/en/networking/hpe-aruba-networking-support-services.html
AOS-CX Switch Software Documentation Portal	https://arubanetworking.hpe.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) +1-650-750-0350 (Backup—Toll Number)
International telephone	https://www.hpe.com/psnow/doc/a50011948enw

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:

HPE Aruba Networking Developer Hub	https://developer.arubanetworks.com/hpe-aruba-networking-aoscx/docs/about
Airheads social forums and Knowledge Base	https://community.arubanetworks.com/
AOS-CX Software Technical Update channel on YouTube.	Videos on new features introduced in this release: https://www.youtube.com/playlist?list=PLsYGHuNuBZcbWPEjjHuVMqP-Q_UL3CskS

HPE Aruba Networking Hardware Documentation and Translations Portal	https://arubanetworking.hpe.com/techdocs/hardware/DocumentationPortal/Content/home.htm
HPE Aruba Networking software	https://networkingsupport.hpe.com/downloads
Software licensing and Feature Packs	https://licensemanagement.hpe.com/
End-of-Life information	https://networkingsupport.hpe.com/end-of-life

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.