AOS-CX 10.09 Monitoring Guide

8400 Switch Series

a Hewlett Packard Enterprise company

Published: September 2022
Edition: 4

Open Source Code

This product includes code licensed under the GNU General Public License, the GNU Lesser General Public License, and/or certain other open source licenses. A complete machine-readable copy of the source code corresponding to such code 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, send a check or money order in the amount of US $10.00 to:

Hewlett Packard Enterprise Company
6280 America Center Drive
San Jose, CA 95002
USA

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.

Chapter 1
About this document

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

Applicable products

This document applies to the following products:

- Aruba 8400 Switch Series (JL375A, JL376A)

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
`example-text`	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: ■ *<example-text>* ■ <example-text> ■ *example-text* ■ *example-text*	Identifies a placeholder—such as a parameter or a variable—that you must substitute with an actual value in a command or in code: ■ 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 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.

Convention	Usage
... or . . .	Ellipsis: ■ 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 8400 Switch Series

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

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

Identifying modular switch components

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

The software notation for describing member, slot, and port information depends on the switch hardware.

The physical interfaces on the Aruba 8400 Switch Series use the format:

member/*slot*/*port*

member

Specifies the chassis number. In this release of the software, the value of *member* is always `1`.

slot

Specifies physical location in the switch chassis.

port

Specifies the physical port on the module.

The slot numbers are unique to each type of component—in contrast to being unique within a chassis.

Line Modules and Management Modules

Line modules are on the front of the switch in slots 1/1 through 1/4 and 1/7 through 1/10.

The number of ports depend on the line module. Line module ports are labeled in software as port or interface, depending on the context.

For example, `interface 1/1/1` is the logical interface associated with the physical interface member 1, slot 1, port 1.

Management modules are on the front of the switch in slots 1/5 and 1/6.

Figure 1 *Aruba 8400 Switch Series line module and management module slots*

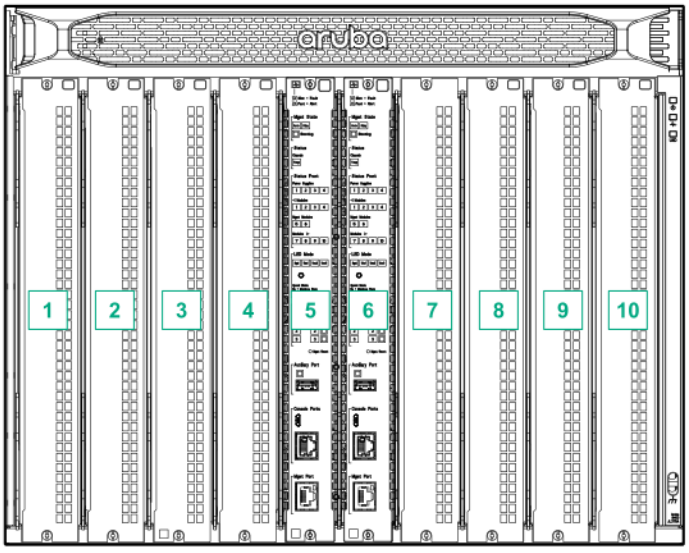

Power supplies

Power supplies are on the front of the switch behind the bezel above the line modules and management modules. Power supplies are labeled in software as Member/PSU: 1/1 through 1/4.

Fan trays

Fan trays are on the rear of the switch and are labeled in software as Member/Tray: 1/1 through 1/3.

Fans

Fans are on the rear of the switch in fan trays and are labeled in software as Member/Tray/Fan:

- 1/1/1 through 1/1/6
- 1/2/1 through 1/2/6
- 1/3/1 through 1/3/6

Fabric modules

Fabric modules are on the rear of the switch, behind the fan trays, in slots 1/1 through 1/3.

Rear display module

The rear display module is on the rear of the switch and is not labeled with a member or slot number.

Confirming normal operation of the switch by reading LEDs

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

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.

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.).

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.

Chassis LEDs

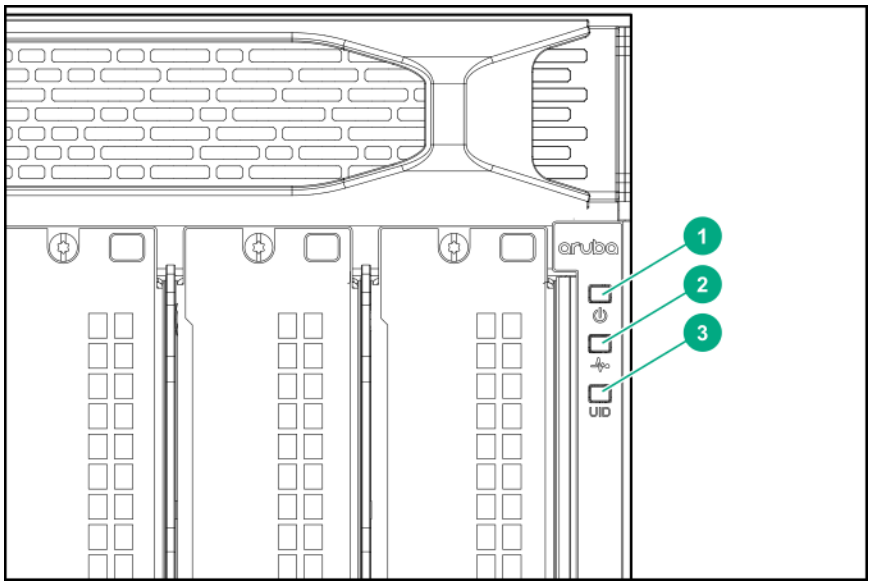

LED	
1	Chassis power LED
2	Chassis health LED
3	Unit identification (UID)

Chassis LED behavior

Chassis power LED

LED state	Meaning
Off	No power to chassis
On Green	Powered on

Chassis health LED

LED state	Meaning
Off	No power to chassis
On Green	Normal operation
Slow Flash Green	Booting
Slow Flash Orange	Fault condition somewhere in the system

Chassis UID (unique identifier) LED

LED state	Meaning
Off	Not activated
On Blue	Location aid
Slow Flash Blue	Location aid

Fabric module and fan module LEDs

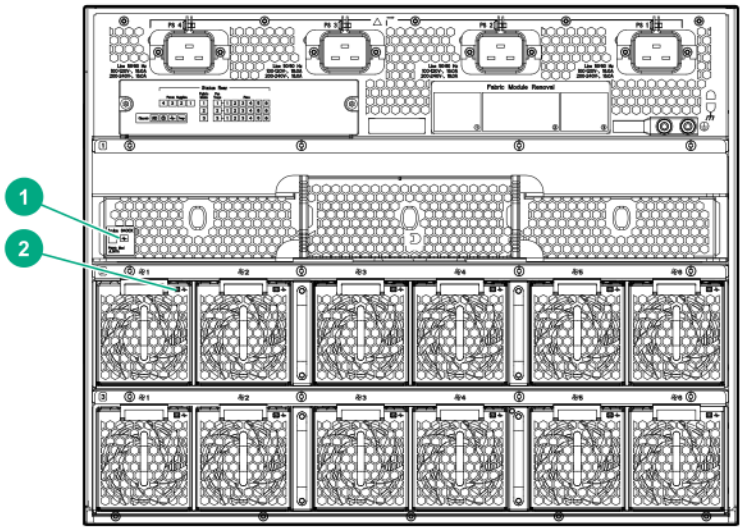

	LED
1	Fabric module LEDs Fabric module 1 shown with fan tray 1 and fan modules removed. Fabric modules are located behind the fan trays. Fabric module LEDs are visible if the entire fan tray is removed or if Fan 1 is removed from the fan tray.
2	Fan Module LEDs

Fabric module LED behavior

The fabric module health LED is not directly visible when Fan 1 is in place. The purpose of the health ID for fabric modules is to allow you to ensure that the module is shut down and ready for removal.

LED state	Meaning
Off	Powered down, ready for removal
On Green	Normal operation
Slow Flash Green	Booting or initializing
Slow Flash Orange	Fault condition
Fast Flash Orange	Hardware (ASIC) mismatch

Fan module LED behavior

LED state	Meaning
Off	Shut down or not operating
On Green	Normal Operation
Slow Flash Orange	Fault condition

Line module LEDs

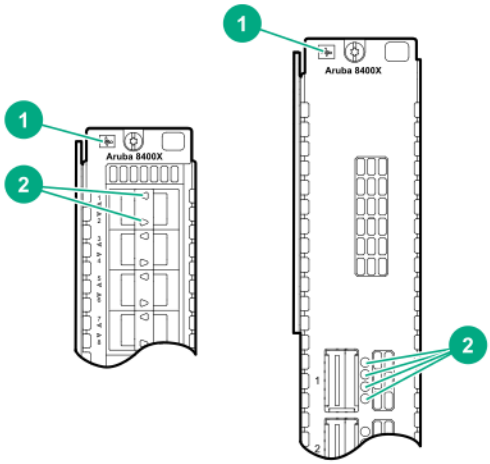

LED	
1	Line module health LED
2	Line module port LEDs

Line module LED behavior

LED state	Meaning
Off	Powered down, ready for removal
On Green	Ready or normal operation
Slow Flash Green	Booting or initializing
Slow Flash Orange	Fault condition
Fast Flash Orange	Hardware (ASIC) mismatch

Line module port LEDs when LED mode is Link/Activity

LED behavior for port LEDs is set by LED Mode button on management module.

LED State	Meaning
Off	Port is disabled, not connected or not receiving a link indication from the connected device
Half-Bright Green	Port is enabled and receiving a link indication from the connected device
Switching between Half-Bright and Full-Bright Green	Port is actively transferring data
On Green	Port is at high utilization
Slow Flash Orange	Port is in a fault condition

Line module port LEDs when LED mode is Spd

LED State	Meaning
Off	Port is disabled, not connected or not receiving a link indication from the connected device
On Green	Port is operating at maximum port rate
Slow Flash Green	Port is operating below maximum port rate
Slow Flash Orange	Fault condition

Line module port LEDs when LED mode is Usr1 (light faults)

This mode is primarily to assist users who cannot discern between green and orange LEDs. It is also helpful to users who want to locate an occurring fault quickly.

LED State	Meaning
Off	All Green LEDs (except the chassis power and Usr1 LEDs) are Off in this LED mode
Slow Flash Orange	Fault condition

Management module LEDs, ports, and buttons

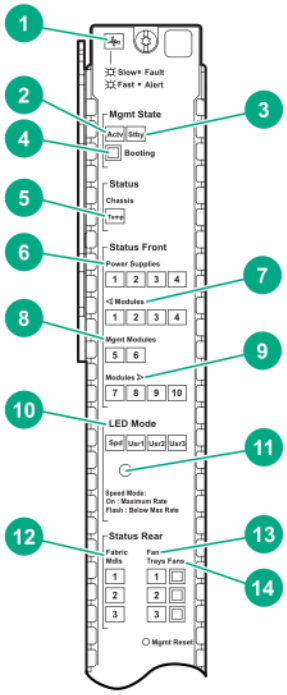

	LED
1	Management module health LED
2	Active state (Actv)
3	Standby state (Stby)
4	Booting
5	Chassis temperature status (Temp)
6	Power supply status (1 2 3 4)
7	Line modules status (1 2 3 4) Module slots 1 - 4
8	Management module status (5 6) Module slots 5 and 6
9	Line modules status (7 8 9 10) Module slots 7 - 10
10	Line module port LED State (Spd) (Usr1) (Usr2) (Usr3) selected by LED Mode button.
11	LED Mode button Changes display of the line module port LEDs from the default Link/Activity behavior to cycle through the speed (Spd) and user (Usr) options
12	Fabric module status (1 2 3)
13	Fan tray status (1 2 3)
14	Fans (fan modules 1 - 6 in indicated fan tray)

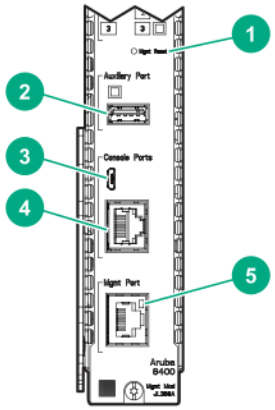

	LED
1	Mgmt Reset button Recessed button used to reset the selected management module
2	Auxiliary port
3	USB Micro-B Console Port
4	Serial Console Port (RJ-45)
5	Mgmt Port (OOBM Port) Activity/Link LED

Management module LED behavior

Management module health LED

LED state	Meaning
Off	Shut down or ready for removal
On Green	Booting the AOS-CX operating system, ready, or normal operation
Slow Flash Green	Booting the BIOS or service operating system, or updating firmware
Slow Flash Orange	Management module fault
Fast Flash Orange	Software (NOS) mismatch

Chassis temperature status (Temp) LED

LED state	Meaning
On Green	Temperature at or below rating
Slow Flash Orange	Environment air temperature or individual component temperature above rating

Management state LEDs

Standby management modules do not show component status LEDs.

Table 1: *Actv (active state) and Stby (standby state) LEDs*

LED State	Active state (Actv)	Standby state (Stby)
On Green (only one on at a time)	This module is the active management module and it is ready.	This module is the standby management module and it is ready.
Both LEDs Off	This management module is not ready (booting or shut down).	

Table 2: *Booting*

LED state	Meaning
Off	Operating system has not started booting Management module is ready
Slow Flash Green	AOS-CX operating system is booting

Status Front LEDs

Table 3: *Power Supplies: 1-4*

LED state	Meaning
Off	Not present
On Green	Operating normally
Slow Flash Orange	Fault condition

Table 4: *Line Modules: 1-4 and 7-10*

LED state	Meaning
Off	Powered down, ready for removal
On Green	Ready or normal operation
Slow Flash Green	Booting or initializing
Slow Flash Orange	Fault condition
Fast Flash Orange	Hardware (ASIC) mismatch

Table 5: *Management Modules: 5-6*

LED state	Meaning
Off	Not present or module is shut down and ready for removal
On Green	Operating normally

Status Rear LEDs

Table 6: *Fabric Modules: 1-3*

LED state	Meaning
Off	Powered down, ready for removal
On Green	Normal operation
Slow Flash Green	Booting or initializing
Slow Flash Orange	Fault condition
Fast Flash Orange	Hardware (ASIC) mismatch

Table 7: *Fan Trays: 1-3*

LED state	Meaning
On Green	Operating normally
Slow Flash Orange	Fault condition or not present

Table 8: *Fans: 1-6 in each fan tray*

LED state	Meaning
On Green	Operating normally
Slow Flash Orange	One or more fans in the given fan tray is in a fault condition or is not present

LED Mode

The LED Mode button changes the behavior of port LEDs from the default mode to cycle through various other settings. The selected LED mode reverts to the default mode after 10 minutes. The LED modes are the following:

Default mode

In the default mode:

- All LEDs in the LED Mode section are off.
- Port LED behavior indicates link and activity without rate indications.
- If a line module port is in a fault condition, its LED changes from Flashing Green (indicating activity) to Slow Flash Orange (fault condition).

Speed mode

In the speed mode:

- The Spd (Speed) LED is On Green.
- Port LED behavior indicates link speed. A Port LED is On Green when the link is at maximum speed and Flashing Green when the link is below maximum speed.

User 1 Mode

The behavior for the Usr1 (User 1) mode is to Light Faults. Use this mode to assist users who cannot discern between Green and Orange LEDs or who want to locate an occurring fault quickly.

In Light Faults mode:

- The Usr1 (User 1) LED is On Green.
- All other Green LEDs except the chassis power LEDs are Off.
- LEDs that indicate a fault condition are Slow Flash Orange or Fast Flash Orange.

Auxiliary Port LED

Off (Not used).

Management Port LEDs

LED state	Meaning
Off	Port is disabled, not connected or not receiving a link indication from the connected device
Half-Bright Green	Port is enabled and receiving a link indication from the connected device
Switching between Half-Bright and Full-Bright Green	Port is actively transferring data
On Green	Port is at high utilization

Power supply LEDs

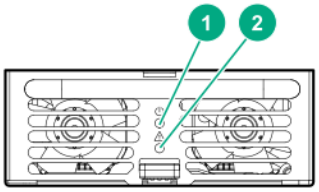

	LED
1	Power LED (green)
2	Fault LED (orange)

Power supply LED behavior

Power LED state	Fault LED state	Power supply condition
On Green	Off	Normal operation
Off	Off	No AC power to power supply

Power LED state	Fault LED state	Power supply condition
Flashing Green	Off	AC power is present, but the power supply is not supplying power to the chassis
On Green	Flashing Orange	Power limit exceeded
Off	On Orange	Temperature limit exceeded or abnormal output detected

Rear panel LEDs

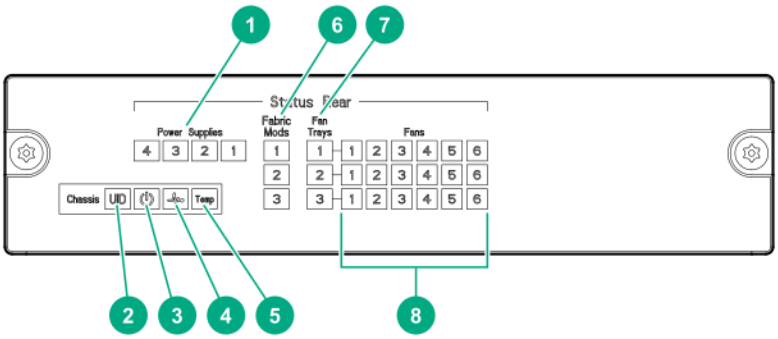

	LED
1	Power supply status (1 2 3)
2	Unit identification (UID)
3	Chassis power LED
4	Chassis health LED
5	Chassis temperature status (Temp)
6	Fabric modules
7	Fan tray status (1 2 3)
8	Fan status each fan tray (1 2 3 4 5 6)
Not shown*	Fan health (Health LED) (*For a given fan module, this LED appears to the right of the module release lever.)
Not shown**	Fabric Module Health (Health Icon) (**To view this LED for a given fabric module (1 through 3), remove the first (leftmost) fan module from the fan tray slot in which the fabric module is installed.)

Rear panel LED behavior

Chassis Power LED

LED state	Meaning
Off	No power to chassis
On Green	Powered on

Chassis Health LED

LED state	Meaning
On Green	Normal operation
Slow Flash Green	Booting
Slow Flash Orange	Fault condition somewhere in the system

Chassis Unique Identifier (UID) LED

LED state	Meaning
Off	Not activated
On Blue	Location aid
Slow Flash Blue	Location aid

Chassis temperature status (Temp) LED

LED state	Meaning
On Green	Temperature at or below rating
Slow Flash Orange	Environment air temperature or individual component temperature above rating

Status Rear LEDs

Table 1: *Fabric Modules 1-3*

LED state	Meaning
Off	Powered down, ready for removal
On Green	Normal operation
Slow Flash Green	Booting or initializing
Slow Flash Orange	Fault condition
Fast Flash Orange	Hardware (ASIC) mismatch

Table 2: *Fan Trays: 1-3*

LED state	Meaning
On Green	Operating normally
Slow Flash Orange	Fault condition or not present

Table 3: *Fans: 1-6 in each fan tray*

LED state	Meaning
On Green	Operating normally
Slow Flash Orange	Fan is in a fault condition or is not present

LED states

Off

No discernable light emitted.

On

No discernable intensity variation, full-bright.

Activity flicker

The port LED switches to full-bright to indicate that traffic is passing through the port. If the port LED is steady full-bright, the port utilization is continuous.

Half-bright

No discernable intensity variation, intensity is 50% of full-bright.

Slow Flash

0.8s on / 0.8s off

Fast Flash

0.4s on / 0.4s off

boot fabric-module

`boot fabric-module <SLOT-ID>`

Description

Reboots the specified fabric module.

Parameter	Description
<SLOT-ID>	Specifies the member and slot of the module in the format `member/slot`. For example, to specify the module in member 1 slot 3, enter `1/3`.

Usage

The `boot fabric-module` command reboots the specified fabric module. Traffic performance is affected while the module is down.

If the specified module is the only fabric module in an up state, rebooting that module stops traffic switching between line modules and the line modules power down. The line modules power up when one fabric module returns to an up state.

This command is valid for fabric modules only.

Examples

Rebooting the fabric module in slot **1/3** when auto-confirm is not enabled:

```
switch# boot fabric-module 1/3
This command will reboot the specified fabric module.  Traffic performance may
be affected while the module is down.  Rebooting the last fabric module will
stop traffic switching between line modules.
Do you want to continue (y/n)? y

switch#
```

Rebooting the fabric module in slot **1/1** when auto-confirm is enabled:

```
switch# boot fabric-module 1/3
This command will reboot the specified fabric module.  Traffic performance may
be affected while the module is down.  Rebooting the last fabric module will
stop traffic switching between line modules.
Do you want to continue (y/n) y (auto-confirm)

switch#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	Manager (#)	Administrators or local user group members with execution rights for this command.

boot line-module

`boot line-module <SLOT-ID>`

Description

Reboots the specified line module.

Parameter	Description
<SLOT-ID>	Specifies the member and slot of the module in the format `member/slot`. For example, to specify the module in member 1 slot 3, enter `1/3`.

Usage

Reboots the specified line module. Any traffic for the switch passing through the affected module (SSH, TELNET, and SNMP) is interrupted. It can take up to 2 minutes to reboot the module. During that time, you can monitor progress by viewing the event log.

This command is valid for line modules only.

Examples

Reloading the module in slot 1/1:

```
switch# boot line-module 1/1
This command will reboot the specified line module and interfaces on this
module will not send or receive packets while the module is down. Any
traffic passing through the line module will be interrupted. Management
sessions connected through the line module will be affected. It might take
up to 2 minutes to complete rebooting the module. During that time, you can
monitor progress by viewing the event log.
Do you want to continue (y/n)? y
switch#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	Manager (#)	Administrators or local user group members with execution rights for this command.

boot management-module

`boot management-module {active | standby | <SLOT-ID>}`

Description

Reboots the specified management module. Choose the management module to reboot by role (active or standby) or by slot number.

Parameter	Description
active	Selects the active management module.
standby	Selects the standby management module.
<SLOT-ID>	Specifies the member and slot of the management module in the format `member/slot`. For example, to specify the module in member 1 slot 5, enter `1/5`.

Usage

This command reboots a single management module in a chassis. Choose the management module to reboot by role (active or standby) or by slot number.

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

If you reboot the active management module and the standby management module is available, the active management module reboots and the standby management module becomes the active management module.

If you reboot the active management module and the standby management module is not available, you are warned, you are prompted to save the configuration, and you are prompted to confirm the operation.

If you reboot the standby management module, the standby management module reboots and remains the standby management module.

If you attempt to reboot a management module that is not available, the `boot` command is aborted.

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` command is aborted.

Hewlett Packarrd Enterprise recommends that you use the `boot management-module` command instead of pressing the module reset button to reboot a management module because if you are rebooting the only available management module, the `boot management-module` command enables you to save the configuration, cancel the reboot, or both.

Examples

Rebooting the active management module when the standby management module is available:

```
switch# boot management-module active
The management-module in slot 1/5 is going down for reboot now.
```

Rebooting the active management module when the standby management module is not available:

```
switch# boot management-module 1/5
The management module in slot 1/5 is currently active and no
standby management module was found.
This will reboot the entire switch.

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.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	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.

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
`primary`	Selects the primary operating system image for this reboot and sets the configured default operating system image to `primary` for future reboots.
`secondary`	Selects the secondary operating system image for this reboot and sets the configured default operating system image to `secondary` for future reboots.
`serviceos`	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.
```

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]
```

Description

Shows boot information. When no parameters are specified, shows the most recent information about the boot operation, and the three previous boot operations for the active management module. When the `all` parameter is specified, shows the boot information for the active management module and all available line modules. To view boot-history on the standby, the command must be sent on the standby console.

Parameter	Description
`all`	Shows boot information for the active management module and all available line modules.

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.

`Index`
The position of the boot in the history file. Range: 0 to3.

`Boot ID`
A unique ID for the boot . A system-generated 128-bit string.

`Current Boot, up for <SECONDS> seconds`
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:

`<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.

`Reboot requested through database`
The reboot occurred because of a request made through the CLI or other API.

`Uncontrolled reboot`
The reason for the reboot is not known.

Examples

Showing the boot history of the active management module:

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

Index : 3
```

```
Boot ID : f1bf071bdd04492bbf8439c6e479d612
Current Boot, up for 22 hrs 12 mins 22 secs

Index : 2
Boot ID : edfa2d6598d24e989668306c4a56a06d
07 Aug 18 16:28:01 : Reboot requested through database

Index : 1
Boot ID : 0bda8d0361df4a7e8e3acdc1dba5caad
07 Aug 18 14:08:46 : Reboot requested through database

Index : 0
Boot ID : 23da2b0e26d048d7b3f4b6721b69c110
07 Aug 18 13:00:46 : Reboot requested through database
switch#
```

Showing the boot history of the active management module and all line modules:

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

Index : 3
Boot ID : f1bf071bdd04492bbf8439c6e479d612
Current Boot, up for 22 hrs 12 mins 22 secs

Index : 2
Boot ID : edfa2d6598d24e989668306c4a56a06d
07 Aug 18 16:28:01 : Reboot requested through database

Index : 1
Boot ID : 0bda8d0361df4a7e8e3acdc1dba5caad
07 Aug 18 14:08:46 : Reboot requested through database

Index : 0
Boot ID : 23da2b0e26d048d7b3f4b6721b69c110
07 Aug 18 13:00:46 : Reboot requested through database

Line module 1/1
=================
Index : 3
10 Aug 17 12:45:46 : dune_agent crashed
...
```

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.

Chapter 6
Switch system and hardware commands

Switch system and hardware commands are general commands used to configure fundamental settings on the switch.

Refer to the Fundamentals Guide to view the switch system and hardware commands.

The switch has limited capacity to store data, collected by switch features and protocols. You can provide virtually unlimited storage capacity by adding user-supplied external storage volumes. Supported volume types and storage protocols include: NFSv3, NFSv4, and SCP (sshfs).

One application of external storage is the saving and restoring of DHCP lease files over SCP or NFS network attached storage systems. SCP file system protocol uses a user mode process to emulate a network file system. The key advantage is packet level encryption and simple configuration. The key disadvantage is slow performance.

You can set up external storage volume credentials and then enable it. A storage management process acts on your requests by enabling the storage volume using the requested storage protocol. You can disable the external storage volume or set it up but leave it disable.

The feature maintains storage volume state. The states are: *disabled* (down), *connecting* (establishing connection), *operational* (up), and *unaccessible* (unavailable).

If a storage volume is unavailable, the system attempts to reconnect periodically. Multiple volumes could connect concurrently. If one connection times out the others can connect immediately.

The system supports server connection through data and management ports.

Data port support requires server IP address on a default VRF.

Once a storage volume is enabled, applications can use the volume to store retrieve and delete files and directories.

External storage commands

address (external storage)

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

Description

Specifies the NAS IP address or hostname.

The `no` form of this command deletes an IP address or hostname.

Parameter	Description
<IPV4-ADDR>	Specifies the NAS server IPv4 address, Global.
<IPV6-ADDR>	Specifies the IPv6 address of the NAS server.
<HOSTNAME>	Specifies the hostname of the NAS server. String.

Examples

Creating the logfiles storage volume with IP address 10.1.1.1:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# address 10.1.1.1
```

Deleting an external storage volume named logfiles:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# no address 10.1.1.1
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-external-storage-*<VOLUME-NAME>*	Administrators or local user group members with execution rights for this command.

directory

```
directory <DIRECTORY-NAME>
no directory <DIRECTORY-NAME>
```

Description

Selects an existing directory on the external storage volume.

The no form of this command clears a directory of an external storage volume.

Parameter	Description
<DIRECTORY-NAME>	Specifies the external storage directory for mapping the volume.

Examples

Creating a volume named logfiles that is mapped under /home on the server:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# directory /home
```

Clearing the directory /home:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# no directory /home
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config-external-storage-<VOLUME-NAME>`	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

disable external-storage logfiles

```
disable
no disable
```

Description

Disables the external storage volume.

The `no` form of this command enables the external storage volume. This is identical to the `enable` command.

Examples

Disabling a volume named logfiles:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config-external-storage-<VOLUME-NAME>`	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

enable (external-storage logfiles)

```
enable
no enable
```

Description

Enables the external storage volume.

The `no` form of this command disables the external storage volume. This is identical to the `disable` command.

Examples

Creating and then enabling a volume named logfiles:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# enable
```

Disables the external storage volume:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# disable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config-external-storage-<VOLUME-NAME>`	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

external-storage

`external-storage <VOLUME-NAME>`
`no external-storage <VOLUME-NAME>`

Description

Creates or updates an external storage volume.

The `no` form of this command deletes an external storage volume.

Examples

Creating the logfiles storage volume:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)#
```

Deleting the logfiles storage volume:

```
switch(config)# no external-storage logfiles
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config`	Administrators or local user group members with execution rights for this command.

password (external-storage)

```
password [{plaintext | ciphertext} <PASSWORD>]
no password {plaintext | ciphertext} <PASSWORD>
```

Description

Sets the password for network attached storage server login.

The `no` form of this command clears the password for network attached storage server login.

Parameter	Description
`{ciphertext | plaintext}`	Selects the password format.
<PASSWORD>	Specifies the password. **NOTE:** When the password is not provided on the command line, plaintext password prompting occurs upon pressing Enter. The entered password characters are masked with asterisks.

Examples

Creating a volume named logfiles with password Xj#9:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# password plaintext Xj#9
```

Creating a volume named bak1 with a prompted plaintext password:

```
switch(config)# external-storage bak1
switch(config-external-storage-bak1)# password
Enter the NAS server password: **********
Re-Enter the NAS server password: **********
```

Clearing the password for volume logfiles:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# no password plaintext Xj#9
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config-external-storage-<VOLUME-NAME>`	Administrators or local user group members with execution rights for this command.

show external-storage

`show external-storage [<VOLUME-NAME>]`

Description

Shows external storage configuration and state for all volumes or for a specified volume.

Parameter	Description
<VOLUME-NAME>	Specifies the external storage volume name that the show command will use.

Examples

```
switch# show external-storage

------------------------------------------------------------------------------
--
            Address      VRF      Username      Type       Directory     State
------------------------------------------------------------------------------
--
nfsvol      10.1.1.1     nas      ---           NFSv3      /home
operational
nfsfiles    20.1.1.1     nas      netstorage    NFSv4      /netstor      disabled
scpdev      nasserver    nas      scpstor       SCP        /scp
unaccessible
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show running-config external-storage

```
show running-config external-storage
```

Description

Shows the running configuration of the external storage.

Examples

```
switch# show running-config external-storage

external-storage nfsvol
      address   10.1.1.1
      vrf       nas
      type      nfsv4
      directoty /home
      enable
external-storage scpdev
      address   30.1.1.1
      vrf       nas
      username  switchuser
      password  ciphertext xxx
      type      scp
      directoty /home
      enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

type (external storage)

```
type {nfsv3 | nfsv4 | scp}
no type {nfsv3 | nfsv4 | scp}
```

Description

Sets the network attached storage access type for reaching the external storage volume.

The `no` form of this command deletes an external storage volume.

Parameter	Description
`nfsv3`	Specifies the NFSv3 network access protocol.
`nfsv4`	Specifies the NFSv4 network access protocol.
`scp`	Specifies the SCP network access protocol.

Examples

Creating the logfiles volume using NFSV4:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# type nfsv4
```

Clearing the external storage access type:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# no type nfsv4
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-external-storage-*<VOLUME-NAME>*	Administrators or local user group members with execution rights for this command.

username (external storage)

```
username <USER-NAME>
no username <USER-NAME>
```

Description

Sets the username for logging in to a network attached storage server.

The `no` form of this command clears a username.

Parameter	Description
<USER-NAME>	Specifies the username.

Examples

Creating a volume named logfiles with the user name nassuser:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# username nasuser
```

Clearing the user name nasuser from accessing the logfiles volume:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# no username nasuser
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config-external-storage-<VOLUME-NAME>`	Administrators or local user group members with execution rights for this command.

vrf (external storage)

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

Description

Setting a VRF to reach network attached storage.

The `no` form of this command clears access of a VRF to network attached storage.

Parameter	Description
<VRF-NAME>	Specifies the VRF name.

Examples

Creating the logfiles volume and setting a VRF named nas to access the network attached storage:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# vrf nas
```

Clearing access of a VRF named nas to the network attached storage:

```
switch(config)# external-storage logfiles
switch(config-external-storage-logfiles)# no vrf nas
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config-external-storage-<VOLUME-NAME>`	Administrators or local user group

Platforms	Command context	Authority
		members with execution rights for this command.

The IP Service Level Agreement (IP-SLA) is a feature that enables the measuring of network performance between two nodes in a network for different service level agreement parameters such as round-trip time (RTT), one-way delay, jitter, reachability, packet loss, and voice quality scores. These two nodes can span across area in access, distribution or core inside a LAN as well as across WAN between core to core or core to Data Centre switches. This feature helps you measure the SLA for different protocols or applications such as UDP echo, UDP jitter (for voice and video), TCP connect, HTTP, and ICMP echo. This guide provides details for managing and monitoring different types of IP-SLAs.

IP-SLA guidelines

- AOS-CX supports only SLA configuration through CLI and thresholds can be configured using NAE agents using WebUI/REST.
- AOS-CX supports only forever tests. On-demand tests are not supported.
- Maximum sessions: IP-SLA source 500, IP-SLA responder 100.
- NAE can effectively monitor a maximum of 300 parameters, reducing the maximum supported session by 300.
- NAE supports only syslog.
- NAE agents must be triggered for each IP-SLA test on every switch.
- If multiple IP addresses are received for a DNS query, DNS works with the first resolved IP.
- When the DNS server IP is not configured, the first DNS server in `resolve.conf` is used.
- The source interface/IP option is not applicable for SLAs configured on 'mgmt' VRF, as it has only one interface.
- A system time change because of NTP or a manual change causes an incorrect calculation.
- There is no interoperability of UDP echo SLA between AOS-CX and FlexFabric switches.
- Source IP and source port combination must be unique across SLA sessions in a same switch.
- Do not use the same source port across the source and responder sessions in a switch.
- NTP synchronization is a must for SLA types involving one-way delay such as UDP jitter VoIP.
- It is mandatory to set default CoPP to the max value when UDP jitter SLA is enabled otherwise 100% packet loss can be seen and `UDP-Jitter sla` probe will result in failure as seen in the following example.

```
copp-policy default
     class hypertext priority 6 rate 50000 burst 64
     default-class priority 6 rate 99999 burst 9999
```

- Deviations with respect to PVOS results: The packet losses due to internal switch-related issues like interface shutdown or interface flaps will not be considered as 'Probes Timed-out error', as the IP-SLA solution is to measure network performance and anomalies. Rather, this kind of packet loss will be counted in internal counters like 'Destination address unreachable'.

Limitations with VoIP SLAs

- A maximum of 80 concurrent VoIP SLAs can be scheduled in a 20 second slot.
- A single VoIP probe takes 20 seconds to complete.
- The default and minimum probe interval for VoIP SLA is 120 seconds.
- SLAs scheduled in the same slot, periodically sends 1000 probe packets for 120 seconds in 20 second intervals.
- Default 120 second probe interval is divided in to 6 slots of 20 seconds to avoid synchronization of all configured VoIP SLAs sending probes at the same time.
- SLAs started at the same time exceeding the concurrent limit of 80 must wait for the next 20 second VoIP slot to open before moving to 'running' state.
- The maximum number of VoIP SLAs supported is 80 X 6 slots = 480 SLAs.
- SLAs exceeding 480 will continue to remain in the 'waiting for VoIP slot' until any slot is freed by stopping the running SLA.
- To avoid high RTT, a single switch with more than 20 SLAs should not have single responder SLA.
- When IP is received dynamically (e.g. using DHCP) for interfaces other than management interface, IPSLA source or responder has to be configured only using interface name.

IP-SLA commands

http

```
http {get | raw} URL [source {<SOURCE-IPV4-ADDR> | <IFNAME>} source-port <PORT-NUM>]
     [proxy proxy-url] [cache disable] [name-server <IPV4-ADDR-DNS-SERVER>]
     [probe-interval <30-604800>] [version<VERSION-NUMBER>] [http-raw-request <RAW-
PAYLOAD>]
```

Description

Configures HTTP as the IP-SLA test mechanism. Requires destination URL and type of HTTP request (raw/get).

Parameter	Description
`{get | raw}`	Selects HTTP request type as GET or RAW where the system will generate or provide HTTP payload.
`URL`	Specifies HTTP URL address of syntax. http://<HOST NAME/IP-ADDRESS>:<PORT>/<PATH>.
`source {<SOURCE-IPV4-ADDR> | <IFNAME>}`	Selects the source IPv4 address for SLA probes or the source interface to use for sending IP-SLA probes.
`source-port <PORT-NUM>`	Specifies the value of the source port for the IP-SLA probes.
`cache disable`	Selects cache option for the HTTP server. By default the option is enabled.
`name-server <IPV4-ADDR-DNS-SERVER>`	Specifies the IPv4 address of DNS server.
`probe-interval <PROBE-INTERVAL>`	Specifies the probe interval in seconds. Range: 30 to 604800.

Parameter	Description
`version <VERSION-NUMBER>`	Specifies the source interface to use for sending IP-SLA probes.
`http-raw-request <RAW-PAYLOAD>`	HTTP raw request. String.

Examples

```
switch(config-ipsla-1)# http get http://device.arubanetworks.com/root/home.html
    switch(config-ipsla-1)# http raw
http://device.arubanetworks.com/root/home.html
    switch(config-ipsla-1)# http 2.2.2.2 source 1/1/1
    switch(config-ipsla-1)# http http://device.arubanetworks.com  source 2.2.2.1
    switch(config-ipsla-1)# http http://device.arubanetworks.com/root/home.html
source-interface 1/1/1
    switch(config-ipsla-1)# http http://device.arubanetworks.com  name-server
10.10.10.2
    switch(config-ipsla-1)# http raw raw-request "GET /en/US/hmpgs/index.html
HTTP/1.0\r\n\r\n"
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config-ip-sla-<IP-SLA-NAME>`	Administrators or local user group members with execution rights for this command.

icmp-echo

```
icmp-echo {<DEST-IPV4-ADDR>|<HOSTNAME>} [source {<SOURCE-IPV4-ADDR> | <IFNAME>}]
     [name-server <IPV4-ADDR-DNS-SERVER>] [payload-size <PAYLOAD-SIZE>]
     [tos <TYPE-OF-SERVICE>] [probe-interval <PROBE-INTERVAL>]
```

Description

Configures ICMP echo as the IP-SLA test mechanism. Requires destination address for the IP-SLA test.

Parameter	Description
`{<DEST-IPV4-ADDR> | <HOSTNAME>}`	Selects the destination IPv4 address for the IP-SLA or the hostname of the destination.
`[source {<SOURCE-IPV4-ADDR> | <IFNAME>}]`	Selects the source IPv4 address for SLA probes or the source interface to use for sending IP-SLA probes.
`name-server <IPV4-ADDR-DNS-SERVER>`	Specifies the DNS server for destination hostname resolution.

Parameter	Description
payload-size *<PAYLOAD-SIZE>*	Specifies the payload size of an SLA probe. Range: 0 to 1440.
tos *<TYPE-OF-SERVICE>*	Specifies the type of serve to be used in the probe packets. Range: 0 to 255.
probe-interval *<PROBE-INTERVAL>*	Specifies the probe interval in seconds. Range: 5 to 604800.

Examples

```
switch(config)# ip-sla test
    switch(config-ip-sla-test)# icmp-echo 2.2.2.2
    switch(config-ip-sla-test)# icmp-echo 2.2.2.2 source 3.3.3.3
    switch(config-ip-sla-test)# icmp-echo 2.2.2.2 source 3.3.3.3 payload-size 400
    switch(config-ip-sla-test)# icmp-echo 2.2.2.2 source 3.3.3.3 payload-size 400
name-server 4.4.4.4
    switch(config-ip-sla-test)# icmp-echo 2.2.2.2 source 3.3.3.3 payload-size 400
name-server 4.4.4.4  probe-interval 80
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

ip-sla

ip-sla *<IP-SLA-NAME>*
no ip-sla *<IP-SLA-NAME>*

Description

Creates an IP Service Level Agreement (SLA) profile and switches to the config-ip-sla context.

The no form of this command deletes an IP-SLA profile. By default, all profile use the default VRF (default).

Parameter	Description
<IP-SLA-NAME>	Specifies an IP-SLA profile name. Length: 1 to 63 characters.

Examples

Creating an IP-SLA:

```
switch(config)# ip-sla 1
switch(config-ip-sla-1)#
```

Deleting an IP-SLA:

```
switch(config)# no ip-sla 1
switch(config)#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config	Administrators or local user group members with execution rights for this command.

ip-sla responder

```
ip-sla responder <SLA-NAME> {udp-echo | tcp-connect | udp-jitter-voip} <PORT-NUM>
     [source {<SOURCE-IPV4-ADDR> | <IFNAME>}][vrf <VRF-NAME>]
no ip-sla responder <SLA-NAME> {udp-echo | tcp-connect | udp-jitter-voip} <PORT-NUM>
     [source {<SOURCE-IPV4-ADDR> | <IFNAME>}][vrf <VRF-NAME>]
```

Description

Selects the IP-SLA responder. The responder can be configured for udp-echo, tcp-connect, udp-jitter-voip type. It requires the SLA name, SLA type, and port number as arguments. Source IP/interface ID is a must for type udp-jitter-voip and optional for other types.

The `no` form of this command removes the IP-SLA responder.

Parameter	Description
<SLA-NAME>	Specifies the SLA name.
udp-echo	Enables responder for udp-echo probes.
tcp-connect	Selects TCP connect as the IP-SLA test mechanism.
vrf *<VRF-NAME>*	Specifies the name of the VRF to use.
udp-jitter-voip	Selects VOIP jitter as the IP-SLA test mechanism.
<PORT-NUM>	Specifies the port number to listen for IP-SLA probes. Range: 1 to 65535.
[source {*<SOURCE-IPV4-ADDR>* | *<IFNAME>*}]	Selects the source IPv4 address for SLA probes or the source interface to use for sending IP-SLA probes.

Examples

```
switch(config)# ip-sla responder SLA1 udp-echo 8000 source 2.2.2.2
switch(config)# ip-sla responder SLA1 udp-echo 8000 source 1/1/1
```

```
switch(config)# no ip-sla responder SLA1 udp-echo 8000 source 2.2.2.2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config	Administrators or local user group members with execution rights for this command.

show ip-sla responder

`show ip-sla responder <SLA-NAME>`

Description

Shows the given IP-SLA responder configuration and operation status.

Parameter	Description
<SLA-NAME>	Specifies the SLA name.

Examples

```
switch(config)# show ip-sla responder SLA3

    SLA Name             : SLA3
    IP-SLA Type          : Udp-echo
    VRF                  : Default
    Responder Port       : 8000
    Responder IP         : 2.2.2.3
    Responder Interface  : 1/1/1
    Responder Status     : Running
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config	Administrators or local user group members with execution rights for this command.

show ip-sla responder results

```
show ip-sla responder <SLA-NAME> <SOURCE-IPV4-ADDR> <PORT-NUM> results
```

Description

Shows the given ip-sla responder statistics for a given source IP and port. This command is only applicable for the sources where source IP and port are configured.

Parameter	Description
<SLA-NAME>	Specifies the SLA name.
<SOURCE-IPV4-ADDR>	Specifies the source IPV4 address.
<PORT-NUM>	Specifies the port number. Range: 1 to 65535.

Examples

```
switch# show  ip-sla responder SLA1 2.2.2.1 8000 results

    IP-SLA Type          : Udp-echo
    VRF Name             : Default
    Source IP            : 2.2.2.1
    Source Port          : 8000
    Responder Port       : 8888
    Responder IP         : 2.2.2.3
    Responder Interface  :
    Responder Status     : Running
    Packets Received     : 2
    Packets Sent         : 2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config	Administrators or local user group members with execution rights for this command.

show ip-sla *<SLA-NAME>*

```
show ip-sla <SLA-NAME> results
```

Description

Shows the given IP-SLA source configuration and status.

Parameter	Description
<SLA-NAME>	Specifies the SLA name.
`results`	Shows the statistics calculated for an SLA type.

Examples

```
switch# show ip-sla xyz results

    IP-SLA session status
       IP-SLA Name                     : xyz
       IP-SLA Type                     : tcp-connect
       Destination Host Name/IP Address: 2.2.2.1
       Destination Port                : 8888
       Source IP Address/IFName        : 2.2.2.2
       Source Port                     : 5555
       Status                          : Running

    IP-SLA session cumulative counters
       Total Probes Transmitted        : 1
       Probes Timed-out                : 0
       Bind Error                      : 0
       Destination Address Unreachable : 0
       DNS Resolution Failures         : 0
       Reception Error                 : 0
       Transmission Error              : 0

    IP-SLA Latest Probe Results
       Last Probe Time                 : 2018 Jul 13 02:00:35
       Packets Sent                    : 1
       Packets Received                : 1
       Packet Loss in Test             : 0.0000%

    Minimum RTT(ms)                    : 0.7900
    Maximum RTT(ms)                    : 0.7900
    Average RTT(ms)                    : 0.7900
    DNS RTT(ms)                        : 0.0000
    TCP RTT(ms)                        : 0.9710

switch(config)# show ip-sla xyz
    IP-SLA Name             : xyz
    Status                  : scheduled
    IP-SLA Type             : tcp-connect
    VRF                     : ipslasrc
    Source Port             : 5555
    Source IP               : 2.2.2.2
    Source Interface        :
    Domain Name Server      :
    Probe interval(seconds) : 90


switch(config)# show ip-sla jitter-sla results
    IP-SLA session status
       IP-SLA Name                     : jitter-sla
       IP-SLA Type                     : udp-jitter-voip
       Destination Host Name/IP Address: 2.2.2.1
       Destination Port                : 8888
       Source IP Address/IFName        :
```

```
       Source Port                     : 5555
       Status                          : Running

    IP-SLA Session Cumulative Counters
       Total Probes Transmitted        : 1
       Probes Timed-out                : 0
       Bind Error                      : 0
       Destination Address Unreachable : 0
       DNS Resolution Failures         : 0
       Reception Error                 : 0
       Transmission Error              : 0

    IP-SLA Latest Probe Results
       Last Probe Time                 : 2018 Jul 13 02:02:48
       Packets Sent                    : 1
       Packets Received                : 1
       Packet Loss in Test             : 0.0000%

       Minimum RTT(ms)                 : 0.7900
       Maximum RTT(ms)                 : 0.7900
       Average RTT(ms)                 : 0.7900
       DNS RTT(ms)                     : 0.0000

       Min Positive SD                 : 1      Min Positive DS        : 2
       Max Positive SD                 : 1      Max Positive DS        : 2
       Positive SD Number              : 2      Positive DS Number     : 2
       Positive SD Sum                 : 2      Positive DS Sum        : 4
       Positive SD Average             : 5      Positive DS Average    : 5
       Min Negative SD                 : 1      Min Negative DS        : 1
       Max Negative SD                 : 1      Max Negative DS        : 1
       Negative SD Number              : 2      Negative DS Number     : 4
       Negative SD Sum                 : 2      Negative DS Sum        : 4
       Negative SD Average             : 5      Negative DS Average    : 5

       Max SD Delay                    : 0      Max DS Delay           : 0
       Min SD Delay                    : 0      Min DS Delay           : 0
       Average SD Delay                : 0      Average DS Delay       : 0

    Voice Scores:
       MOS  Score                      : 4.38   ICPIF                  : 0


switch(config)# show ip-sla m3op
    IP-SLA Name             : jitter-sla
    Status                  : Running
    IP-SLA Type             : udp-jitter-voip
    VRF                     : ipslasrc
    Source IP               : 2.2.2.2
    Source Interface        :
    Domain Name Server      :
    TOS                     : 10
    Probe Interval(seconds) : 90
    Advantage Factor        : 0
    Codec Type              : g711a


switch(config)# show ip-sla http-sla
    IP-SLA Name             : http-sla
    Status                  : Running
    IP-SLA Type             : http
    VRF                     : ipslasrc
    Source IP               : 2.2.2.2
```

```
    Source Interface        :
    Domain Name Server      : 10.10.10.2
    Probe Interval(seconds) : 90
    HTTP Request Type       : GET
    HTTP/HTTPS URL          : abcd.com/ws/home
    Cache                   : Enabled
    HTTP Proxy URL          :
    HTTP Version Number     : 1.1
    ```

IP-SLA status description
    ```
    | Status                  | Description                                          |
    |-------------------------|------------------------------------------------------|
    | Running                 | SLA is fully operational                             |
    | Bind Error              | Another service is using the same source port  |
    | Interface Down          | Interface status is not up
    | Dns Resolution Error    | Failed to resolve destination hostname         |
    | No Route                | No available route to the responder            |
    | Internal Error          | Unexpected error prevents SLA session          |
    | Disabled                | SLA is disabled                                |
    |Configuration Incomplete | Configuration is not complete to enable the SLA|
    ```
IP SLA session cumulative counters description
    ```
    | Status                          | Description
                            |
    |---------------------------------|-----------------------------------------
------------------------------|
    |Probes Timed-out                 | Total numbers of probes failed to receive
response.                       |
    |Bind Error                       | Total numbers of probes transmission failed
as source port not available.|
    |Destination Address Unreachable | Total numbers of probes transmission failed
due to route unavailable.    |
    |DNS Resolution Failures          | Total numbers of probes failed due to DNS
resolution failure.             |
    |Reception Error                  | Total numbers of probes failed due to
internal error in reception.       |
    |Transmission Error               | Total numbers of probes failed due to
internal errr in transmission.     |
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

start-test

`start-test`

Description

Starts the IP-SLA probes.

Examples

```
switch(config)# ip-sla test
switch(config-ip-sla-test)# start-test
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

stop-test

```
stop-test
```

Description

Stops the IP-SLA probes.

Examples

```
switch(config)# ip-sla test
switch(config-ip-sla-test)# stop-test
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

tcp-connect

```
tcp-connect {<DEST-IPV4-ADDR> | <HOSTNAME>} <PORT-NUM> [source {<SOURCE-IPV4-ADDR> |
     <IFNAME>} [source-port <PORT-NUM>]] [name-server <IPV4-ADDR-DNS-SERVER>]
     [probe-interval <PROBE-INTERVAL>]
```

Description

Configures TCP connect as the IP-SLA test mechanism. Requires destination address/hostname and destination port for the IP-SLA of tcp-connect IP-SLA type.

Parameter	Description
{*<DEST-IPV4-ADDR>* | *<HOSTNAME>*}	Selects the destination IPv4 address for the IP-SLA or the hostname of the destination.
<PORT-NUM>	Destination port for the IP-SLA. Range: 1 to 65535.
[source {*<SOURCE-IPV4-ADDR>* | *<IFNAME>*}]	Selects the source IPv4 address for SLA probes or the source interface to use for sending IP-SLA probes.
[source-port *<PORT-NUM>*]	Specifies the port for the IP-SLA test.
[name-server *<IPV4-ADDR-DNS-SERVER>*]	Specifies the DNS server for destination hostname resolution.
[probe-interval *<PROBE-INTERVAL>*]	Probe interval in seconds. Range: 30 to 604800.

Examples

```
switch(config-ipsla-1)# tcp-connect 2.2.2.2 8080
       switch(config-ipsla-1)# tcp-connect 2.2.2.2 8080 source 2.2.2.1 source-port
6000
       switch(config-ipsla-1)# tcp-connect 2.2.2.2 8080 source 1/1/1 source-port
6000

       switch(config-ipsla-1)# tcp-connect https://device.arubanetworks.com 8080
       switch(config-ipsla-1)# tcp-connect https://device.arubanetworks.com 8080
source 2.2.2.1 source-port 6000
       switch(config-ipsla-1)# tcp-connect https://device.arubanetworks.com 8080
source 1/1/1 source-port 6000
       switch(config-ipsla-1)# tcp-connect https://device.arubanetworks.com 8080
name-server 10.10.10.2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

udp-echo

```
udp-echo {<DEST-IPV4-ADDR>|<HOSTNAME>} <PORT-NUM> [source {<SOURCE-IPV4-ADDR> |
     <IFNAME>} [source-port <PORT-NUM>]] [name-server <IPV4-ADDR-DNS-SERVER>] [payload-
size
     <PAYLOAD-SIZE>] [tos <TYPE-OF-SERVICE>] [probe-interval <PROBE-INTERVAL>]
```

Description

Configures UDP echo as the IP-SLA test mechanism. Requires destination address/hostname and destination port number for the IP-SLA of udp-echo SLA type.

Parameter	Description
{*<DEST-IPV4-ADDR>* | *<HOSTNAME>*}	Selects the destination IPv4 address for the IP-SLA or the hostname of the destination.
<PORT-NUM>	Specifies the destination port for the IP-SLA. Range: 1 to 65535.
[source {*<SOURCE-IPV4-ADDR>* | *<IFNAME>*}]	Selects the source IPv4 address for SLA probes or the source interface to use for sending IP-SLA probes.
[source-port *<PORT-NUM>*]	Specifies source port for the IP-SLA test. Range: 1 to 65535.
[name-server *<IPV4-ADDR-DNS-SERVER>*]	Specifies the DNS server for destination hostname resolution.
[payload-size *<PAYLOAD-SIZE>*]	Specifies the payload size of an SLA probe. Range: 28 to 1440.
[*<TYPE-OF-SERVICE>*]	Type of service. Range: 0 to 255.
probe-interval *<PROBE-INTERVAL>*	Probe interval in seconds. Range: 5 to 604800.

Examples

```
switch(config-ipsla-1)# udp-echo 2.2.2.2 8080
    switch(config-ipsla-1)# udp-echo 2.2.2.2 8080 source 2.2.2.1
    switch(config-ipsla-1)# udp-echo https://device.arubanetworks.com 8080
    switch(config-ipsla-1)# udp-echo 2.2.2.2 8080 source 1/1/1
    switch(config-ipsla-1)# udp-echo 2.2.2.2 8080 source 2.2.2.1 payload-size 50
    switch(config-ipsla-1)# udp-echo 2.2.2.2 8080 source 1/1/1 payload-size 50
    switch(config-ipsla-1)# udp-echo 2.2.2.2 8080 payload-size 50
    switch(config-ipsla-1)# udp-echo https://device.arubanetworks.com 8080 source
2.2.2.1
     payload-size 50
    switch(config-ipsla-1)# udp-echo https://device.arubanetworks.com 8080 source
1/1/1
     payload-size 50
    switch(config-ipsla-1)# udp-echo https://device.arubanetworks.com 8080
     name-server 10.10.10.2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

udp-jitter-voip

```
udp-jitter-voip {<DEST-IPV4-ADDR> | <HOSTNAME>} <PORT-NUM> [codec-type <CODEC-TYPE>]
     [advantage-factor <VALUE>] [source {<SOURCE-IPV4-ADDR> | <IFNAME>} [source-port
<PORT-NUM>]]
     [name-server <IPV4-ADDR-DNS-SERVER>][probe-interval <PROBE-INTERVAL>] [tos <TYPE-OF-
SERVICE>]
```

Description

Configure UDP jitter voip as the IP-SLA test mechanism. Requires destination address/hostname and source address/interface for the IP-SLA of udp-jitter-voip IP-SLA type.

Parameter	Description
{*<DEST-IPV4-ADDR>*|*<HOSTNAME>*}	Selects the destination IPv4 address for the IP-SLA or the hostname of the destination.
<PORT-NUM>	Selects the port number for the IP-SLA. Range: 1 to 65535.
[codec-type *<CODEC-TYPE>*]	Selects the codec-type for the Voip IP-SLA test.
[advantage-factor *<ADVANTAGE-FACTOR>*]	Selects the value for the advantage factor. Default value is 0.
[source {*<SOURCE-IPV4-ADDR>* | *<IFNAME>*}]	Selects the source IPv4 address for SLA probes or the source interface to use for sending IP-SLA probes.
[source-port *<PORT-NUM>*]	Specifies the value of source port for the IP-SLA probes.
[name-server *<IPV4-ADDR-DNS-SERVER>*]	Specifies the DNS server for destination hostname resolution.
tos *<TYPE-OF-SERVICE>*	Specifies the type of service. Range: 0 to 255.
probe-interval *<PROBE-INTERVAL>*	Specifies the probe interval in seconds. Range: 120 to 604800.

Examples

```
switch(config-ipsla-1)# udp-jitter-voip  2.2.2.2 8080 advantage-factor 10 codec-
type g711a
    switch(config-ipsla-1)# udp-jitter-voip  2.2.2.2 8080 advantage-factor 10
codec-type g711a source 2.2.2.1
    switch(config-ipsla-1)# udp-jitter-voip  https://device.arubanetworks.com 8080
advantage-factor 10 codec-type g711a
    switch(config-ipsla-1)# udp-jitter-voip  2.2.2.2 8080 advantage-factor 10
codec-type g711a source 1/1/1
    switch(config-ipsla-1)# udp-jitter-voip  https://device.arubanetworks.com 8080
advantage-factor 10 codec-type g711a source 2.2.2.1
    switch(config-ipsla-1)# udp-jitter-voip  https://device.arubanetworks.com 8080
```

```
advantage-factor 10 codec-type g711a source 1/1/1
    switch(config-ipsla-1)# udp-jitter-voip  https://device.arubanetworks.com 8080
advantage-factor 10 codec-type g711a name-server 10.10.10.2 probe-interval 120
source 10.1.1.1 source-port 8888 tos 10
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

vrf (ip sla)

```
vrf <VRF-NAME>
no vrf [<VRF-NAME>]
```

Description

Configures the VRF on which the SLA will send or receive packets. By default, the default VRF is used.

The no form of the command removes VRF from SLA.

Parameter	Description
<VRF-NAME>	Specifies a VRF name. Length: Default: default.

Examples

```
switch(config-ip-sla-test)# vrf ipslasrc
```

```
switch(config-ip-sla-test)# no vrf
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

show interface

```
show interface [<IFNNAME>|<IFRANGE>] [brief | physical | extended [non-zero]]
show interface [lag | loopback | tunnel | vlan ] [<ID>] [brief | physical]
show interface [lag | loopback | tunnel | vlan ] [<ID>] [extended [non-zero]]
```

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.
non-zero	Shows only non zero statistics.
LAG	Shows LAG interface information.
LOOPBACK	Shows loopback interface information.
TUNNEL	Shows tunnel interface information.
VLAN	Shows 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	Shows the VXLAN interface information.
<VXLAN-ID>	Specifies the VXLAN interface identifier. Default: 1

Examples

The following example shows when the interface is configured as a route-only port:

```
switch# show interface 1/1/1

Interface 1/1/1 is up
 Admin state is up
 Link state: up for 2 days (since Sun Jun 21 05:30:22 UTC 2020)
 Link transitions: 1
 Description: backup data center link

 Hardware: Ethernet, MAC Address: 70:72:cf:fd:e7:b4
 MTU 1500
 Type 1GbT
```

```
 Full-duplex
 qos trust none
 Speed 1000 Mb/s
 Auto-negotiation is on
 Flow-control: off
 Error-control: off

 L3 Counters: Rx Enabled, Tx Enabled
 Rate collection interval: 300 seconds

 Rates                          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

 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
 L3 Packets                      0                    0                    0
 L3 Bytes                        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
   Other                         0                    0                    0
```

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
```

When the interface is shut down during a VSX split:

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

Interface 1/1/1 is down
 Admin state is up
 State information: Disabled by VSX
 Link state: down for 3 days (since Tue Mar 16 05:20:47 UTC 2021)
 Link transitions: 0
 Description:

 Hardware: Ethernet, MAC Address: 04:09:73:62:90:e7
```

```
MTU 1500
Type SFP+DAC3
Full-duplex
qos trust none
Speed 0 Mb/s
Auto-negotiation is off
Flow-control: off
Error-control: off
VLAN Mode: native-untagged
Native VLAN: 1
Allowed VLAN List: 1502-1505
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                            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
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
10.07 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.

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.

Mirroring and sFlow

The mirroring feature (when mirroring received traffic) and the sFlow sampling feature both require the receive (rx) capability of a given port. If both features are configured and enabled to use the receive capability on the same port, only one of the features can perform its task.

This interaction does not affect transmit (tx) mirroring because sFlow does not use the transmit (tx) capability of a port.

Behavior if sFlow is enabled and mirror enable is attempted

If sFlow is enabled on a port and a mirroring session specifies the same port as a 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.
- To enable the mirroring session, first you must disable sFlow on that port.

Behavior if mirroring is enabled and sFlow enable is attempted

If a mirroring session specifies a port as a source of received traffic (the source is configured with a direction of `rx` or `both`), and you attempt to enable sFlow on the same port:

- Mirroring on that port continues.
- No error or warning message is returned when sFlow is enabled, but sFlow sampling on that port does not occur.

When sFlow is enabled on a port but sampling is not occurring, the `show sflow <INTERFACE-NAME>` command shows that sFlow is enabled but shows a value of 0 (zero) for the number of samples.

To activate sFlow sampling on that port, you must do the following:

1. Disable the mirroring session on the port.
2. Disable sFlow on the port.
3. Enable sFlow on the port.

Behavior when the startup configuration has both sFlow and rx mirroring enabled on the same port

If the startup configuration has the same port configured with both sFlow enabled and as a source of received traffic in an enabled mirroring session:

- During a boot or management module failover operation, it is not possible to predict whether the receive capability of the port will be assigned to the sFlow feature or to the mirroring feature.
- To ensure that the feature that you want is used on a specific port, after the boot operation or management module failover operation completes, disable both features on that port and then enable the feature you want to use.

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.

Mirror statistics are reset for a Mirror-to-CPU session on 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.

Any subsequent configuration changes to either the enabled mirroring session or the classifier policy can affect the behavior of the network monitoring that occurs. For examples, see Scenario 1 and Scenario 2.

Scenario 1

1. Mirroring session 1 is configured with destination interface 1/1/10 and source interface 1/1/5 in the `both` direction, then the session is enabled.
2. Mirroring session 2 is configured with destination interface 1/1/20, then the session is enabled.
3. Policy `Policy_2` is configured with a class matching OSPF traffic from any source IPv4 address to

any destination IPv4 address and an action of `mirror`, specifying mirroring session 2.

4. `Policy_2` is applied to interface 1/1/5 in the inbound direction.

This sequence of actions creates a situation where the interface 1/1/5 is effectively configured as a source for two separate enabled mirroring sessions. This configuration is not permitted if you attempt to configure and enable the two mirroring sessions through the CLI. However, mirroring may occur for both sessions because policies with mirror actions have priority over basic mirroring sessions.

In this example:

- Because of `Policy_2`, all OSPF traffic ingressing interface 1/1/5 is mirrored to 1/1/20, which is the destination interface of mirroring session 2.
- After `Policy_2` is applied, and because of the mirroring session 1 is enabled, all non-OSPF traffic ingressing interface 1/1/5 is mirrored to 1/1/10, which is the destination interface of mirroring session 1.
- Because `Policy_2` does not match egressing traffic, and because mirroring session 1 is enabled, all traffic egressing interface 1/1/5 is mirrored to 1/1/10, which is the destination interface of mirroring session 1.

Scenario 2

1. Mirroring session 2 is configured with destination interface 1/1/20 and source interface 1/1/3, then the session is enabled.
2. Policy `Policy_2` is configured with a class matching OSPF traffic from any source IPv4 address to any destination IPv4 address and an action of `mirror` specifying mirroring session 2.
3. `Policy_2` is applied to interface 1/1/5 in the inbound direction.

In this scenario, a single mirroring session is configured with a source interface and is configured as the target of the mirror action of a policy applied to a different source interface. In this example, the destination interface 1/1/20 receives traffic from interface 1/1/3 and from interface 1/1/5.

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.

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
```

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.

copy tcpdump-pcap

`copy tcpdump-pcap <FILE-NAME> <REMOTE-URL>`

Description

Saves packet capture files to external storage.

Parameter	Description
<FILE-NAME>	Specifies the packet capture file to save.
<REMOTE-URL>	Specifies the external storage to which the packet capture file will be saved.

Usage

Only four files can be saved at any point on the switch. Packet capture files are not saved after a failover or reboot. View a list of saved files using `diag utilities list-files`.

Examples

Saving my_capture_file.pcap to sftp://root@10.0.0.2/file.pcap:

```
switch# copy tcpdump-pcap my_capture_file.pcap sftp://root@10.0.0.2/file.pcap
root@10.0.0.2's passowrd:
Connected to 10.0.0.2.
sftp > put my_capture_file.pcap file.pcap
Uploading my_capture_file.pcap to /root/file.pcap
my_capture_file.pcap                                  100%   156   219.8KB/s   00:00
Copied successfuly.
```

Command History

Release	Modification
10.08	Command introduced

Command Information

Platforms	Command context	Authority
8400	Manager (#)	Administrators or local user group members with execution rights for this command.

copy tshark-pcap

```
copy tshark-pcap <REMOTE-URL> [vrf <VRF-NAME>]
```

Description

Copies the tshark capture data to a file on a TFTP or SFTP server.

Parameter	Description
<REMOTE-URL>	Specifies the capture file on a remote TFTP or SFTP server. The URL syntax is: {tftp:// | sftp://*<USER>*@} {*<IP>*|*<HOST>*} [:*<PORT>*] [;blocksize=*<SIZE>*]/*<FILE>*
`vrf` *<VRF-NAME>*	Specifies the name of a VRF. Default: default.

Example

Copying the capture data to a file on SFTP server 10.0.0.2:

```
switch# copy tshark-pcap sftp://root@10.0.0.2/file.pcap

root@10.0.0.2's password:
Connected to 10.0.0.2.
sftp> put packets.pcap file.pcap
Uploading packets.pcap to /root/file.pcap
packets.pcap                                  100%  156   219.8KB/s   00:00
Copied successfully.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	Manager (#)	Administrators or local user group members with execution rights for this command.

destination cpu

```
destination cpu
no destination cpu
```

Description

The command causes the mirror session to transmit mirrored packets to the switch CPU. This destination may be configured for multiple sessions, however only one such configured session may be active at a given time.

The diagnostic utility Tshark may be used to view and capture packets transmitted to the CPU through this route. Ctrl+C must be entered to terminate a Tshark capture session. More details can be found in the *Supportability Guide*.

The `no` form of this command will immediately stops mirroring traffic to the CPU, but will not remove any sources from the mirror configuration.

Examples

Configuring a mirror session with CPU as the destination.

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

Removing the destination entirely.

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

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 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

Supported mirror destinations: Layer 2 or Layer 3 Ethernet ports, LAGs, tunnel, or CPU as a Mirror Destination. A port that is already a member of a LAG is not a valid mirror destination.

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.

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
```

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

```
destination tunnel <TUNNEL-IPV4-ADDR> source <SOURCE-IPv4-ADDR>
   dscp <DSCP-VALUE> vrf <VRF-NAME> id <SPAN-ID>

no 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.

Multiple Mirror Sessions can be enabled with the same source/destination IP address pair if the span IDs are different for sessions. By default it is assigned 0 if not specified.

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.
<SPAN-ID>	Specifies the span ID for the ERSPAN session and during a failover. Range: 0 to 10.

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
```

Replacing the existing destination with a different VRF:

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

Removing the destination:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`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.

diag utilities tcpdump

```
diag utilities tcpdump [command <TEXT> | delete file <FILE-NAME> | list-files |
  vrf <VRF-NAME> | count <COUNT-NUM> | proto <PROTO-NUM> | host-ip <IP-ADDR> | source-ip
  <IP-ADDR> | destination-ip <IP-ADDR> | host-port <PORT> | source-port <PORT> |
  destination-port <PORT> | verbosity <LEVEL> | print <DATA> | ethernet-type <ETH-NUM>]
```

Description

Captures traffic received or transmitted over a network.

Parameter	Description
`command <TEXT>`	Captures packets based on a specified tcpdump command string.
`delete file <FILE-NAME>`	Deletes specified tcpdump list files.
`list-files`	Lists all the tcpdump capture files saved on the device.
`vrf <VRF-NAME>`	Captures packets on the specified VRF. If no VRF is named, the default is used.
`count <COUNT-NUM>`	Runs the tcpdump command until the specified number of packets are captured. Range: 1-2147483647.
`proto <PROTO-NUM>`	Captures packets of a particular type based on IP protocol number. Range: 0-255.
`host-ip <IP-ADDR>`	Captures packets matching with the source or destination IP address.
`source-ip <IP-ADDR>`	Captures packets from the specified IP address.
`destination-ip <IP-ADDR>`	Captures packets sent to the specified IP address.
`host-port <PORT>`	Captures packets matching with the source or destination port.
`source-port <PORT>`	Captures packets from the specified IP port.
`destination-port <PORT>`	Captures packets sent to the specified IP port.
`verbosity <LEVEL>`	Captures packets of the specified verbosity. Range: level1-level4. If no verbosity is specified, the default is level1.
`print <DATA>`	Captures the data of each packet. The maximum is 262144 bytes

Parameter	Description
`ethernet-type <ETH-NUM>`	Captures packets based on the particular ethernet type. Range: 0-65535.

Usage

- When using the `command` option, the only traffic captured will be packets that have been mirrored to the CPU.
- When using the `command` option, command line sanitization is performed to prevent options that may cause harm or security issues. The following options are blocked:
  - -i/--interface
  - -Z
  - -B/--buffer-size
  - -C
  - -W
  - -Z/--relinquish privileges
- Non-word operators such as "&" or "|" are not allowed. Use boolean keywords such as "and," "or," and "not."
- When using `command -r` to read a file, do not provide any directory path characters. Use list-files command to get the list of file names currently saved on the device, and then use those file names.
- A total of four files can be saved at any given point on the device. Packet capture files are not saved after a failover or reboot, but can be saved to external storage using the `copy tcpdump-pcap` command.

Examples

Inspecting traffic mirrored to the CPU via tcpdump and saving the output to my_capture_file.pcap:

```
switch# diag utilities tcpdump command -c 2 -x -w my_capture_file.pcap
Inspecting traffic mirrored to the CPU via tcpdump until Ctrl-C is entered.
2 packets captured
2 packets received by filter
0 packets dropped by kernel
Ending traffic capture.
```

Listing saved capture files:

```
switch# diag utilities tcpdump list-files
my_capture_file.pcap
```

Reading my_capture_file.pcap:

```
switch# diag utilities tcpdump command -r my_capture_file.pcap
reading from file /tmp/tcpdump/my_capture_file1.pcap, link-type EN10MB (Ethernet)
  1  11:59:34.047867 IP6 localhost.40318 > localhost.ntp: NTPv2, Reserved, length
12
        0x0000:  0000 0304 0006 0000 0000 0000 0000 86dd  ................
        0x0010:  600a 7e47 0014 1140 0000 0000 0000 0000  `.~G...@........
        0x0020:  0000 0000 0000 0001 0000 0000 0000 0000  ................
```

```
        0x0030:  0000 0000 0000 0001 9d7e 007b 0014 0027  .........~.{...'
        0x0040:  1601 0001 0000 0000 0000 0000            ............
  2  11:59:34.047915 IP6 localhost.ntp > localhost.40318: NTPv2, Reserved, length
12
        0x0000:  0000 0304 0006 0000 0000 0000 0000 86dd  ................
        0x0010:  6b8d 23c5 0014 1140 0000 0000 0000 0000  k.#....@........
        0x0020:  0000 0000 0000 0001 0000 0000 0000 0000  ................
        0x0030:  0000 0000 0000 0001 007b 9d7e 0014 0027  .........{.~...'
        0x0040:  d681 0001 c016 0000 0000 0000
```

Removing my_capture_file.pcap:

```
switch# diag utilities tcpdump delete-file my_capture_file.pcap
Successfully removed file
```

Command History

Release	Modification
10.08	Command introduced

Command Information

Platforms	Command context	Authority
8400	Manager (#)	Administrators or local user group members with execution rights for this command.

disable (mirror session)

`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-<SESSION-ID>`	Administrators or local user group members with execution rights for this command.

enable (mirror session)

`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.
- To enable the mirroring session, first you must disable sFlow on the port.

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-<SESSION-ID>`	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
<SESSION-ID>	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)#
```

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.

show mirror

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

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
`vsx-peer`	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Usage

Admin Status indicates the configured status. 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.

Operation Status 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
 Output Packets: 0
 Output Bytes: 0
switch#
```

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.

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:
`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
8400	256

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.

You can configure the same source interface in multiple mirroring sessions, but only one of those mirroring sessions can be enabled at a time.

Classifier policies with mirror actions can also be used to match and mirror network traffic. Although mirror actions of classifier policies must specify an enabled mirroring session, the traffic matching and mirroring actions are separate from and take priority over basic mirroring sessions. For example, mirroring session 1 might monitor a source interface, but a classifier policy might match some traffic from that same source interface and direct it to the destination interface of a different mirroring session. In this situation, only the traffic that is not matched by the policy is considered for matching by mirroring session 1.

If an interface is in active use by the sFlow feature, then that interface cannot be used as source of received traffic (configured as a source destination with a direction of `rx` or `both`) in an enabled mirroring session. If you want to use this interface as a source of received traffic in a mirroring session, you must disable sFlow on the interface before you enable the mirroring session on the same interface.

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.

source vlan

Syntax

```
source vlan <VLAN-NUM> {rx | tx | both}
no source vlan <VLAN-NUM> [rx | tx | both]
```

Description

Adds or removes VLAN as a source of traffic to be mirrored. More than one source VLAN can be configured in a mirror session. Each VLAN may specify its own direction.

The `no` version of the command ceases mirroring traffic from the specified source VLAN and removes the source from the mirror configuration.

There is a limit of 1024 source VLANs in each direction of a given mirror session. The same VLAN can be configured as a mirror source for multiple sessions.

Command context

`config`

Parameters

<VLAN-NUM>

Configured VLAN number.

`rx`

Mirror only received traffic.

`tx`

Mirror only transmitted traffic.

`both`

Mirror both received and transmitted traffic.

Authority

Administrators or local user group members with execution rights for this command.

Example

Create a mirror session and add VLAN 10 as a source of traffic in both directions on that port.

```
switch(config)# mirror session 1
switch(config-mirror-1)# source vlan 10 both
```

Create a second mirror session and add VLAN 10 as a transmit sources of traffic and VLAN 20 in both receive and transmit directions.

```
switch(config)# mirror session 2
switch(config-mirror-2)# source vlan 10 tx
switch(config-mirror-2)# source vlan 20 both
```

Reconfigure the source in session 2 to be receive only by respecifying the source interface configuration.

```
switch(config-mirror-2)# source vlan 10 rx
```

From the second session, remove the first source interface entirely and remove the transmit direction from the other so that mirroring only occurs in the receive direction.

```
switch(config-mirror-2)# no source vlan 10
switch(config-mirror-2)# no source vlan 20 tx
```

Message received when trying to add more than 1024 mirror source VLANs

```
switch(config-mirror-2)# source vlan 2000 rx
The maximum number of source VLANs per mirror session is 1024 in each direction
```

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.

Ports default to an unsplit state. When a port is 'split', the split interfaces become active and can be configured independently. For example, when a 40G QSFP+ port is split four ways, each split interface behaves like a separate 10G SFP+ port. The split interfaces have the same name as the base port with an added suffix to represent their lane of the breakout cable or optical channel on SR4 optics. Splitting an interface removes most of the port's configuration settings and makes it inactive. The port will no longer appear in many show interface commands and most configuration commands are not allowed; the split interface name must be used.

The same thing happens in reverse when an interface is unsplit. However, note that the 'split' and 'no split' commands are always performed in the unsplit port's context.

Limitations with breakout cable support

- The 8400 switch does not support DAC breakout cables, only optical breakout cables.
- The JL365A Aruba 8400X 8p 40G QSFP+ Adv module does not support Priority-Based Flow Control (PFC) on split ports.
- The JL366A Aruba 8400X 6p 40G/100G QSFP28 Adv module does not support 100G breakout cables; it only supports split ports at the 40G speed (into 4x10G links).

Breakout cable support commands

split

```
split [confirm]
no split [confirm]
```

Description

Splits a port into multiple interfaces. Only ports capable of supporting breakout cables or SR4/eSR4 optics can be split.

Parameter	Description
`confirm`	Specifies the confirmation of port splitting.

Usage

The splittable ports for all models are shown in the table below:

Part Number (PN)	Description	Port info
Aruba 8400X modules ▪ JL365A	Aruba 8400X 8p 40G QSFP+ Adv Mod	1-8 (40G)

Part Number (PN)	Description	Port info
▪ JL366A	Aruba 8400X 6p 40G/100G QSFP28 Adv Mod	1-6 Only capable of 40G split into 4 x 10G JL366A modules do not have 25G MACs to support split 100G

Examples

Splitting an interface:

```
switch(config-if)# interface 1/1/1
switch(config-if)# split

This command will disable the specified port, clear its configuration,
and split it into multiple interfaces. The split interfaces will not
be available until the next system or line module reboot.

Continue (y/n)? y

switch(config-if)# show interface brief
--------------------------------------------------------------------------------
----
Port     Native  Mode   Type         Enabled Status  Reason             Speed
Desc             VLAN
(Mb/s)
--------------------------------------------------------------------------------
----
1/1/1:1  --      routed QSFP+DA3x4   yes     down    Split reboot pending   --   -
-
1/1/1:2  --      routed QSFP+DA3x4   yes     down    Split reboot pending   --   -
-
1/1/1:3  --      routed QSFP+DA3x4   yes     down    Split reboot pending   --   -
-
1/1/1:4  --      routed QSFP+DA3x4   yes     down    Split reboot pending   --   -
-
```

Unsplitting a port on a switch that requires a reboot:

```
switch(config)# interface 1/1/1
switch(config-if)# no split
This command will disable the split interfaces for this port and clear
their configuration. The port will not be available until the next
system or line module reboot.

Continue (y/n)? y
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8400	`config-if`	Administrators or local user group members with execution rights for this command.

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.

Accessing Aruba Support

Aruba Support Services	https://www.arubanetworks.com/support-services/
Aruba Support Portal	https://asp.arubanetworks.com/
North America telephone	1-800-943-4526 (US & Canada Toll-Free Number) +1-408-754-1200 (Primary - Toll Number) +1-650-385-6582 (Backup - Toll Number - Use only when all other numbers are not working)
International telephone	https://www.arubanetworks.com/support-services/contact-support/

Be sure to collect the following information before contacting Support:

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Other useful sites

Other websites that can be used to find information:

Airheads social forums and Knowledge Base	https://community.arubanetworks.com/
AOS-CX Switch Software Documentation Portal	https://www.arubanetworks.com/techdocs/AOS-CX/help_portal/Content/home.htm
Aruba Hardware Documentation and Translations Portal	https://www.arubanetworks.com/techdocs/hardware/DocumentationPortal/Content/home.htm

Aruba software	https://asp.arubanetworks.com/downloads
Software licensing	https://lms.arubanetworks.com/
End-of-Life information	https://www.arubanetworks.com/support-services/end-of-life/
Aruba Developer Hub	https://developer.arubanetworks.com/

Accessing Updates

You can access updates from the Aruba Support Portal or the HPE My Networking Website.

Aruba Support Portal

https://asp.arubanetworks.com/downloads

If you are unable to find your product in the Aruba Support Portal, you may need to search My Networking, where older networking products can be found:

My Networking

https://www.hpe.com/networking/support

To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

https://support.hpe.com/portal/site/hpsc/aae/home/

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HP Passport set up with relevant entitlements.

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://asp.arubanetworks.com/notifications/subscriptions (requires an active Aruba Support Portal (ASP) account to manage subscriptions). Security notices are viewable without an ASP 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

Aruba 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

Aruba 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.