

AOS-CX 10.14 Monitoring Guide

6300, 6400 Switch Series*

(excluding S3L75A, S3L76A, S3L77A)



a Hewlett Packard
Enterprise company

Copyright Information

© Copyright 2024 Hewlett Packard Enterprise Development LP.

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

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

Notices

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

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

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

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.

Contents	3
About this document	7
Applicable products	7
Latest version available online	7
Command syntax notation conventions	7
About the examples	8
Identifying switch ports and interfaces	8
Identifying modular switch components	9
Monitoring hardware through visual observation	10
Confirming normal operation of the switch by reading LEDs	10
Detecting if the switch is not ready for a failover event	11
Finding faulted components using the switch LEDs	11
IP Flow Information Export	13
Compatibility with Application Recognition and Traffic Insight	13
Flow monitors	13
Flow Records	13
Flow Exporters	14
Destinations	14
Configuring IP Flow Information Export on 6300, 6400, 8100, 8360 and 8325 Switches	14
Step one: Create Flow Records	15
Step two: Configure flow exporter(s)	16
Step three: Configure the monitor(s)	17
Step four: (Optional) Enable Application Recognition and apply a flow monitor to interfaces	18
Configuring IP Flow Information Export on 10000 Switch Series	18
Compatibility with Traffic Insight	19
Role-based IPFIX	19
FAQs and Troubleshooting	20
Flow monitoring commands	21
diag-dump ipfix basic	21
flow exporter	22
flow monitor	24
flow record	26
flow-tracking	28
ipv4 ipv6 flow monitor (interface)	31
ipv4 ipv6 flow monitor (role)	32
show flow exporter	33
show flow monitor	34
show flow record	35
show flow-tracking	38
show tech ipfix	39
Boot commands	41

boot fabric-module	41
boot line-module	42
boot management-module	43
boot management-module (recovery console)	44
boot set-default	45
boot system	46
show boot-history	48
Switch system and hardware commands	53
External storage	54
External storage commands	54
address	54
directory	55
disable	56
enable	56
external-storage	57
password (external-storage)	58
show external-storage	59
show running-config external-storage	60
type	60
username	61
vrf	62
IP-SLA	64
IP-SLA guidelines	64
Limitations with VoIP SLAs	65
IP-SLA commands	65
http	65
https	66
icmp-echo	68
ip-sla	68
ip-sla responder	69
show ip-sla responder	70
show ip-sla responder results	71
show ip-sla	72
start-test	76
stop-test	77
tcp-connect	77
udp-echo	78
udp-jitter-voip	79
vrf	80
L1-100Mbps downshift	82
Limitations with speed downshift	82
L1-100Mbps downshift commands	82
downshift enable	82
show interface	83
show interface statistics	89
show interface downshift-enable	92
show running-config interface	93
Mirroring	95
Mirror statistics	95
Classifier policies and mirroring sessions	95
VLAN as a source	96

Mirroring commands	96
clear mirror	96
clear mirror endpoint	97
comment	98
copy tcpdump-pcap	99
copy tshark-pcap	100
destination cpu	101
destination interface	101
destination tunnel	102
diagnostic	104
diag utilities tcpdump	105
disable	107
enable	108
mirror session	109
mirror endpoint	109
show mirror	111
show mirror endpoint	112
shutdown	113
source	114
source interface	115
source vlan	117
Monitoring a device using SNMP	120
Power-over-Ethernet	121
PoE commands	122
lldp dot3 poe	122
lldp med poe	123
power-over-ethernet	123
power-over-ethernet allocate-by	124
power-over-ethernet always-on	126
power-over-ethernet assigned-class	126
power-over-ethernet power-pairs	127
power-over-ethernet pre-std-detect	128
power-over-ethernet priority	129
power-over-ethernet quick-poe	130
power-over-ethernet threshold	131
power-over-ethernet trap	132
show lldp local	132
show lldp neighbor	133
show power-over-ethernet	134
Aruba AirWave	142
SNMP support and AirWave	142
SNMP on the switch	142
Supported features with AirWave and the AOS-CX switch	143
Configuring the AOS-CX switch to be monitored by AirWave	143
AirWave commands	144
logging	144
snmp-server community	146
snmp-server host	147
snmp-server vrf	149
snmpv3 context	149
snmpv3 user	150
Support and Other Resources	153

Accessing HPE Aruba Networking Support	153
Accessing Updates	154
Warranty Information	154
Regulatory Information	154
Documentation Feedback	154

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:

- HPE Aruba Networking 6300 Switch Series (JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL762A, R8S89A, R8S90A, R8S91A, R8S92A, S0E91A, S0X44A)
- HPE Aruba Networking 6400 Switch Series (R0X31A, R0X38B, R0X38C, R0X39B, R0X39C, R0X40B, R0X40C, R0X41A, R0X41C, R0X42A, R0X42C, R0X43A, R0X43C, R0X44A, R0X44C, R0X45A, R0X45C, R0X26A, R0X27A, JL741A, S0E48A, S0E48A #0D1, S1T83A, S1T83A #0D1)

Latest version available online

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

Command syntax notation conventions

Convention	Usage
<code>example-text</code>	Identifies commands and their options and operands, code examples, filenames, pathnames, and output displayed in a command window. Items that appear like the example text in the previous column are to be entered exactly as shown and are required unless enclosed in brackets ([]).
example-text	In code and screen examples, indicates text entered by a user.
Any of the following: ■ <code><example-text></code> ■ <code><example-text></code> ■ <i>example-text</i> ■ <i>example-text</i>	Identifies a placeholder—such as a parameter or a variable—that you must substitute with an actual value in a command or in code: ■ 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.

Convention	Usage
	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.
... 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 6300 Switch Series

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

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

On the 6400 Switch Series

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

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

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.

Confirming normal operation of the switch by reading LEDs

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



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

Procedure

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

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

Detecting if the switch is not ready for a failover event

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



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

Procedure

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

Finding faulted components using the switch LEDs

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



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

Procedure

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

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

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

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

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

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

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

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

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

IP Flow Information Export (IPFIX) is an embedded network flow analysis tool that compiles characteristic and measured properties of flows and sends flow reports to flow collectors. IPFIX is configurable via the command-line or REST interfaces. With IPFIX, customers configure flow records with match (key) fields and collection (non-key) fields. Match fields are the set of fields that define a flow, such as IP address or UDP port. Collection fields are the set of fields that identify information to collect for a flow, such as packet and byte counters.

A flow exporter defines where and how to export flow reports. Flow exporters are created as standalone entities in the **config** context to provide flow monitors the ability to export flow reports.

Compatibility with Application Recognition and Traffic Insight

The AOS-CX traffic insight feature allows monitoring of large amount of data that it collects from various flow exporters like IPFIX, and provides the ability to filter, aggregate, and sort the data based on user flow monitor requests. Traffic insight tracks different monitor requests simultaneously and provides monitor reports per request. If the **application recognition** feature is also enabled, then the application data and the flow properties collected by AR and IPFIX are exported to external or internal IPFIX collectors. For more information on configuring the Application Recognition and Traffic Insight features, refer to the *AOS-CX Security Guide*.

Flow monitors

A flow monitor is applied to an interface to perform network traffic monitoring. A flow monitor consists of a flow record, a flow cache, and optional flow exporters. A flow record must be created and assigned to the flow monitor for the monitoring process to function. Flow data is compiled from the network traffic on the interface and stored in the flow cache based on the match (key) and collect (non-key) fields in the flow record. Data from the flow cache is exported by the flow exporters assigned to the flow monitor. 6300 and 6400 series support a maximum of sixteen flow monitors with a limit of two flow exporters that can be applied to a single flow monitor.

Flow Records

A flow record defines match (key) fields and collection (non-key) fields. Match fields are the set of fields that define a flow, such as IP address or UDP port. Collection fields are the set of fields that identify information to collect for a flow, such as packet and byte counters. A maximum of sixteen flow records can be created.

There are six mandatory match fields, of which the IP match fields must be of the same type (IPv4 or IPv6).



A flow record is invalid if it does not contain one of the supported sets of match fields.

The supported sets of match fields are:

1. IPv4:

- IPv4 version
- IPv4 destination address
- IPv4 protocol
- Transport destination port
- Transport source port

2. IPv6:

- IPv6 version
- IPv6 destination address
- IPv6 protocol
- Transport destination port
- Transport source port

Flow Exporters

A flow exporter defines where and how to export flow reports. Flow exporters are created as standalone entities in the **config** context to provide flow monitors the ability to export flow reports. 6300 and 6400 Switch series support a maximum of sixteen flow monitors with a limit of two flow exporters that can be applied to a single flow monitor.

Destinations

The destination specifies where flow reports are sent. There are two possible types of destination for a flow exporter:

1. (default) Hostname or IP address of a device with an optional VRF
2. Traffic Insight instance

A flow exporter can only send flow reports to one destination. The destination type specifies which destination to use. If no destination type is specified, the default destination type is the hostname or IP address of a device with an optional VRF. (If a VRF is not specified, the default VRF will be used.) A destination of each type can be configured, but only the one corresponding to the destination type is used. If there is no destination corresponding to the destination type, then the flow exporter configuration is incomplete. If a new destination of a particular type is configured, it will replace the destination of that type that was previously configured.

Configuring IP Flow Information Export on 6300, 6400, 8100, 8360 and 8325 Switches

The following list describes the steps required to configure a IP flow information export (IPFIX) solution:

- Step one: Create flow records
- Step two: Configure flow exporter(s)
- Step three: Configure monitor(s)
- Step four: Apply a flow monitors to interface(s)

Step one: Create Flow Records

Flow Records are used to define the data that will be added to the IPFIX template. This example configures one record for IPv4 and one for IPv6.

```
switch(config)# flow record flowRecordv4
switch(config-flow-record)# match ipv4 protocol
switch(config-flow-record)# match ipv4 source add
switch(config-flow-record)# match ipv4 destination add
switch(config-flow-record)# match ipv4 version
switch(config-flow-record)# match transport destination port
switch(config-flow-record)# match transport source port
switch(config-flow-record)# collect counter bytes
switch(config-flow-record)# collect counter packets
switch(config-flow-record)# collect application name
switch(config-flow-record)# collect timestamp absolute first
switch(config-flow-record)# collect timestamp absolute last

switch(config)# flow record flowRecordv6
switch(config-flow-record)# match ipv6 protocol
switch(config-flow-record)# match ipv6 source add
switch(config-flow-record)# match ipv6 destination add
switch(config-flow-record)# match ipv6 version
switch(config-flow-record)# match transport destination port
switch(config-flow-record)# match transport source port
switch(config-flow-record)# collect counter bytes
switch(config-flow-record)# collect counter packets
switch(config-flow-record)# collect application name
switch(config-flow-record)# collect timestamp absolute first
switch(config-flow-record)# collect timestamp absolute last
```

Next, use the **show flow record** command to verify the configuration.

```
switch(config)# show flow record
```

```
-----
Flow record 'flowRecordv4'
```

```
-----
Match Fields
```

```
ipv4 destination address
ipv4 protocol
ipv4 source address
ipv4 version
transport destination port
transport source port
Collect Fields
application name
counter bytes
counter packets
timestamp absolute first
timestamp absolute last
```

```
-----
Flow record 'flowRecordv6'
```

```
-----
Match Fields
```

```
ipv6 destination address
ipv6 protocol
ipv6 source address
ipv6 version
transport destination port
transport source port
```

```
Collect Fields
application name
counter bytes
counter packets
timestamp absolute first
timestamp absolute last
```

Step two: Configure flow exporter(s)

In this step, you can define an exporter to send to an external destination by hostname or IP address, or to an internal destination such as Traffic Insight. The example below configures IPFIX to export data to an external address/hostname:

```
switch(config)# flow exporter flowExternal
switch(config-flow-exporter)# destination type hostname-or-ip-addr
switch(config-flow-exporter)# destination 11.1.1.1
switch(config-flow-exporter)# show flow exporter
-----
Flow exporter 'flowExternal'
-----
Status                : Accepted
Export Protocol        : ipfix
Destination Type       : Hostname or IP address
Destination            : 11.1.1.1
Transport Configuration
Protocol               : udp
Port                  : 4739
```

To configure IPFIX to export to Traffic Insight, first configure Traffic Insight.

```
switch(config)# traffic-insight TI
switch(config-ti-TI)# source ipfix
switch(config-ti-TI)# monitor topN type topN-flows
switch(config-ti-TI)# monitor dns type application-flows
switch(config-ti-TI)# enable
```

Next, configure the flow exporter for Traffic Insight

```
switch(config)# flow exporter flowExpTI
switch(config-flow-exporter)# export-protocol ipfix
switch(config-flow-exporter)# destination type traffic-insight
switch(config-flow-exporter)# destination traffic-insight TI
```

You can use the **show flow exporter** command to verify the flow exporter configuration for Traffic Insight

```
switch(config)# show flow exporter flowExpTI
-----
Flow exporter 'flowExpTI'
-----
Status                : Accepted
Export Protocol        : ipfix
Destination Type       : Traffic Insight
```

```
Destination          : TI
Transport Configuration
Protocol              : udp
Port                  : 4739
```

Finally, use the **show run traffic-insight** command to verify the Traffic Insight configuration:

```
switch(config)# show run traffic-insight
traffic-insight TI
enable
source ipfix
!
monitor topN type topN-flows entries 5
monitor appFlow type application-flows
```

Step three: Configure the monitor(s)

First, configure an IPv4 flow monitor.

```
switch(config)# flow monitor flowMonv4
switch(config-flow-monitor)# record flowRecordv4
Switch (config-flow-monitor)# exporter flowExternal
switch(config-flow-monitor)# exit
```

Next, configure an IPv6 flow monitor.

```
switch(config)# flow monitor flowMonv6
switch(config-flow-monitor)# record flowRecordv6
switch(config-flow-monitor)# exporter flowExternal
switch(config-flow-monitor)# exit
```

Once both flow monitors are created, use the **show flow monitor** command to verify the flow monitor configurations.

```
switch(config-flow-monitor)# show flow monitor
```

```
-----
Flow monitor 'flowMonv4'
```

```
-----
Status                : Accepted
Flow Record            : flowRecordv4
Flow Exporter(s)       : flowExternal
Cache Configuration
Inactive Timeout       : 30
Active Timeout         : 1800
```

```
-----
Flow monitor 'flowMonv6'
```

```
-----
Status                : Accepted
Flow Record            : flowRecordv6
Flow Exporter(s)       : flowExternal
Cache Configuration
Inactive Timeout       : 30
Active Timeout         : 1800
```

Step four: (Optional) Enable Application Recognition and apply a flow monitor to interfaces



Enable Application Recognition only if you are using IPFIX to send an application ID. You do not need to enable Application Recognition for IPFIX to be able to report information to an external collector or for internal analytics reports

If you want to use IPFIX to send an application ID to the Application Recognition feature, you must first enable Application Recognition.

```
switch(config)# no ip source-lockdown resource-extended
switch(config)# app-recognition
switch(config-app-recognition)# enable
switch(config-app-recognition)# exit
```

Next, apply flow monitor to IPv4 and IPv6 interfaces

```
switch(config)# int 1/1/1-1/1/28
switch(config-if)# app-recognition enable
switch(config-if)# ip flow monitor flowMonv4 in
switch(config-if)# ipv6 flow monitor flowMonv6 in
switch(config-if)# exit
```

Finally, use the **show run interface** command to verify that the flow monitor was applied to interface.


```
switch(config-if)# show run int 1/1/1
interface 1/1/1
no shutdown
no routing
vlan access 1
app-recognition enable
ip flow monitor flowMonv4 in
ipv6 flow monitor flowMonv6 in
exit
```

Configuring IP Flow Information Export on 10000 Switch Series

With IPFIX on the 10000 Switch series, the following three distinct processes are involved:

- **Metering Process:** This process monitors an observation point within the network for flows and generates flow reports for the flows. An observation point is a location in the network where packets can be observed. The metering process passes the flow reports to the exporting process.
- **Exporting Process:** This process sends flow reports generated by the metering process to a collecting process.
- **Collecting Process:** This process receives flow reports from the exporting process and stores it or further processes it.


Figure 1 *IPFIX to Traffic Insight Workflow*



Metering and exporting processes are configured through AMD Pensando Policy and Services Manager (PSM), a cloud native application, that is connected to the device. IPFIX must be enabled under DSM context and IP source-interface must be specified for IPFIX via CLI or REST interface. The information for PSM and DSM configuration can be found in the DSS guide. For more information about configuration commands, see [AMD-PENSANDO DSS configuration guide](#).

The switch may optionally act as an intermediate collecting process for flow reports from DSMs to append certain additional IPFIX information elements to the flow reports. The switch will then act as an intermediate exporting process to export the augmented flow reports to any configured flow exporters.

Figure 2 IPFIX Flow Collector Workflow



To configure an intermediate collecting process on a 10000 Switch series with IPFIX, perform the following steps:

- Create a flow collector.
- Assign the address of a local loopback interface as the collector's **listen** address.
- Configure PSM with an export policy destined to the address of the local loopback interface.
- Add any additional information elements to be appended to the flow collector with the **append** keyword.

To configure an intermediate exporting process with IPFIX, perform the following steps:

- Create a flow exporter.
- Configure the flow exporter with a destination of a traffic insight instance.
- Assign the flow exporter to the previously-created flow monitor.

Finally, apply the collecting and exporting processes to the DSM by performing the following steps:

- Create a flow monitor.
- Assign the flow collector and flow exporter to the flow monitor.
- Apply the flow monitor within the DSM context on the switch.

Compatibility with Traffic Insight

The AOS-CX traffic insight feature allows monitoring of large amount of data that it collects from various flow exporters like IPFIX, and provides the ability to filter, aggregate, and sort the data based on user flow monitor requests. Traffic insight tracks different monitor requests simultaneously and provides monitor reports per request. For more information on configuring these features, refer to the *AOS-CX Security Guide*.

Role-based IPFIX

If IPFIX monitoring is configured on a port, a network administrator who wants to monitor traffic from clients must preconfigure IPFIX monitor before onboarding the clients. Also, if a client moves from one port to another, the monitor must be reconfigured accordingly. To reduce all these complexities, 6300 and 6400 switch series allow you to configure an IPFIX monitor on a port access role.

An IPFIX deployment with monitoring configured on a port uses TCAM to get flow statistics. Since TCAM is a high-demand resource which is shared across multiple protocols including the security protocols, security policy applications could fail if IPFIX consumes most of the TCAM resources. Port access clients can use a large amount of of TCAM resources for policy applications, making it difficult support port access clients with TCAM based IPFIX.

With role-based IPFIX, when a user plugs a device into to a specific port, their user role is applied to the port. Whenever the user plugs that device into to a different port again, the role is applied to that specific port. This gives mobility and accessibility to the user and provides permissions specific to the user's role when the user's device moves across ports. Role-based IPFIX enables colorless port behavior.

There are some mutually exclusive features that will disable role-based IPFIX when they are enabled. The mutually exclusive features are:

- UBT
- MAC Lockout
- Extended Router MAC
- Source Port Filter
- IP Lockdown
- L3VNI
- VSF stack

The following behaviors are observed in a deployment with role-based IPFIX:

- If a user applies the IPFIX monitor configuration on the port, it uses the existing TCAM infrastructure.
- If a port has both Role based IPFIX and Port based IPFIX, the Port based IPFIX implementation takes the priority.
- If IPFIX monitor configuration on a port moves from role based to TCAM based or vice-versa, existing IPFIX flows learned on that port are deleted.
- When the user role has an IPFIX configuration and fails to apply it, port access logs it, but it does not prevent the client from getting authorized.
- If port access is in the client mode and if there is more than one client with a conflicting IPFIX monitor configuration, traffic from all the clients will be exported through the IPFIX monitor configured for the first client.
- Role-based IPFIX and TCAM port-based IPFIX implementations can co-exist on a switch.
- Flows from unauthorized clients are not monitored.

These are some limitations to role-based IPFIX:

- Role based IPFIX can be enabled only through port access role assignment.
- Only Ingress flow monitoring is supported on both role-based and port-based IPFIX.

FAQs and Troubleshooting

- When IPFIX is used with Application Recognition, these features do not support LAGs or MCLAGs (VSX LAGs).

- The following messages are displayed to indicate an illegal argument:
 - % The flow exporter <EXPORTER-NAME> does not exist.
 - % The flow record <RECORD-NAME> does not exist.
 - % The flow monitor <MONITOR-NAME> does not exist.
 - Invalid destination IP address or hostname entered.
 - Unable to create the flow exporter. The maximum allowed number of flow exporters (<max>) has been reached.
 - Unable to create the flow record. The maximum allowed number of flow records (<max>) has been reached.
 - Unable to create the flow monitor. The maximum allowed number of flow monitors (<max>) has been reached.
 - Flow monitor cannot be applied while interface is part of LAG <LAG-NAME>.
 - Flow monitor could not be applied.
 - Flow monitor could not be unapplied

Flow monitoring commands

diag-dump ipfix basic

diag-dump ipfix basic

Description

Displays diagnostic information for IPFIX.

Examples

```
diag-dump ipfix basic
=====
[Start] Feature ipfix Time : Tue Apr 11 02:23:03 2023
=====
-----
[Start] Daemon ipfixd
-----
- IPFIX Record Cache dump -
- IPFIX Record ipfix -

....

:- IPFIX Monitor v6ti completed -
- End of IPFIX Monitor Cache dump -
-----
[End] Daemon ipfixd
-----
-----
[Start] Daemon ops-switchd
-----
Key format: <traffic_type>_<coalescence_id>_<agent_id>_<asic_port>
Key          TCAM Entry ID      Count
-----
1_1532781829_3_20      0xffff7c7e7a00      1
1_3217499901_1_12      0xffff91187580      1
1_3217499901_1_13      0xffff91183d80      1
```

```
1_3217499901_1_14          0xffff91186e80    1
```

```
....
```

```
-----  
[End] Daemon ops-switchd  
-----
```

```
=====
```

```
[End] Feature ipfix  
=====
```

```
Diagnostic-dump captured for feature ipfix
```

Command History

Release	Modification
10.11	Command introduced on 6300, 6400, 8100 and 8360 Switch series.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	Manager (#)	Administrators or local user group members with execution rights for this command.

flow exporter

```
flow exporter <name>  
  destination  
    <hostname> [vrf vrfname]  
    <ipaddr> [vrf vrfname]  
    <ip6addr> [vrf vrfname]  
  type traffic-insight  
no ..
```

Description

A flow exporter is the part of the IP Flow Information Export (IPFIX) feature that defines how a flow monitor exports flow reports. You can assign the same flow exporter configuration to more than one flow monitor. Each flow exporter includes a destination setting that identifies the device to which the flow reports are sent. 6300 and 6400 series support a maximum of sixteen flow monitors with a limit of two flow exporters that can be applied to a single flow monitor.

Parameter	Description
<name>	Name of the flow exporter, up to 64 characters.
dscp <0-63>	DSCP value to be used by the flow exporter. The default value is 0 .
export-protocol ipfix	Define an export protocol for the flow exporter. The default ipfix

Parameter	Description
	protocol is the only protocol currently available.
description <description>	A description of the flow exporter, up to 256 characters and spaces.
destination <hostname> <IPaddr> <ip6addr>	The exporter sends flow records to this destination. The destination can be defined as a hostname, or an IPv4 or IPv6 IP address.
[vrf vrfname]	You can optionally include the name of the destination VRF in the destination definition.
destination type {hostname-or-ip-addr traffic-insight}	The exporter sends flow reports to a traffic insight destination.
destination traffic-insight <name>	The exporter sends flow reports to a specific traffic insight destination.
no ..	Negate any configured parameter.
template data timeout <timeout>	A flow exporter template describes the format of exported flow reports. Therefore, flow reports cannot be decoded properly without the corresponding templates. This setting defines how often the flow exporter will resend templates to the flow monitor. The supported range is 1-86400 seconds, and the default is 600 seconds.
transport udp <port>	Transport protocol and port for sending flow record reports. The default port is port 4739,

Examples

The following example creates a flow exporter configuration named **exporter-1**.

```
switch(config)# flow exporter exporter-1
switch(config-flow-exporter)# dscp 34
switch(config-flow-exporter)# destination 192.0.2.1 vrf VRF1
switch(config-flow-exporter)# template data timeout 1200
switch(config-flow-exporter)# description Exports flows to 192.0.2.1
```

Related Commands

Command	Description
flow record	Define data to be included in a flow record by configuring flow record match and collect fields
flow monitor	Define a flow monitor configuration, including the flow exporter and flow record associated to that monitor.
show flow exporter	Display flow exporter configuration and status.

Command History

Release	Modification
10.11	Command introduced on 6300, 6400, 8100 and 8360 Switch series.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config config-flow-exporter	Administrators or local user group members with execution rights for this command.

flow monitor

```
flow monitor <name>
  exporter <name>
  cache timeout active|inactive <timeout>
  description <description>
  record <name>
```

Description

On a 6300 and 6400 Switch series, a flow monitor is the part of the IP Flow Information Export (IPFIX) feature that performs network monitoring for the selected interface. A flow monitor configuration consists of a flow record, a flow cache, and one or more associated flow exporters. A flow monitor compiles data from the network traffic on the interface and stores it in the flow cache in a format defined by the flow record. The flow exporters associated with the monitor then export data from the flow cache to the flow exporter destination.



6300 and 6400 series support a maximum of sixteen flow monitors with a limit of two flow exporters that can be applied to a single flow monitor. If no software augmentation of flows is required, there is no need to configure a flow collector or flow monitor.

Parameter	Description
<name>	Name of the flow monitor , up to 64 characters.
cache timeout active inactive <timeout>	Use the cache timeout parameter to define an active or inactive timeout for the flow monitor. A flow

Parameter	Description
	monitor closes a flow session that is active for longer than the active timeout or inactive for longer than the inactive timeout. The supported timeout range for both the active timeout and inactive timeout is 30-120 seconds, and the default is 30 seconds.
description	A description up to 256 characters long, including spaces.
exporter <name>	Assign a flow exporter to a flow monitor. Each flow monitor supports a maximum of two different flow exporters, sending flow records to up to two destinations.
record <name>	) Assigns a flow record to a flow monitor.

Examples

The following example creates a flow monitor configuration named **monitor-1**.

```
switch(config)# flow monitor monitor-1
switch(config-flow-monitor)# description Monitor for analyzing basic ipv4 traffic
switch(config-flow-monitor)# exporter flow-exporter-1
switch(config-flow-monitor)# exporter flow-exporter-2
switch(config-flow-monitor)# record flow-record-1
switch(config-flow-monitor)# cache timeout inactive 120
switch(config-flow-monitor)# cache timeout active 1500
```

The following workflow changes the flow record assigned to a flow monitor.

```
switch(config)# flow monitor flow-monitor-1
switch(config-flow-monitor)# record flow-record-2
```

Related Commands

Command	Description
flow exporter	Define how a flow monitor exports the flow reports.
flow record	Define data to be included in a flow record by configuring flow record match and collect fields
show flow monitor	Displays flow monitor configuration and status

Command History

Release	Modification
10.11	Command introduced on 6400, 6400, 8200 and 8360 Switch series.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config config-flow-monitor	Administrators or local user group members with execution rights for this command.

flow record

```

flow record <name>
  match
    ip|ipv6 {protocol|version}|{source|destination address}
    transport {source|destination} port
  collect
    application name
    application https url
    dns response-code
    counter {packets|bytes}
    timestamp absolute {first|last}
    description <description>

```

Description

Define data to be included in a flow record by configuring flow record match and collect fields.

A flow record defines match (key) fields and collection (non-key) fields. Customers configure flow records with **match** (key) fields and **collect** (non-key) fields. Match fields are the set of fields that define a flow, such as IP address or UDP port. Collect fields are the set of fields that identify information to collect for a flow, such as packet and byte counters.

Traffic with matching attributes (for example, traffic coming from the same interface, sent to the same destination with the same protocol) are classified as a single flow. Information for some or all of the matched settings can be collected and exported to a destination defined by the flow exporter assigned to the flow monitor.



Traffic must match a match rule definition before it can be collected and sent. You cannot collect and send data that is not matched.

Parameter	Description
<name>	Name of the flow monitor, up to 64 characters.
match	match traffic according to one or more of the following key attributes: ▪ ip: match traffic on an IPv4 network ▪ ipv6: match traffic on an IPv6 network ▪ protocol: Match traffic using the same IP protocol ▪ version: Match traffic using the same IP version ▪ source: Match traffic from the same source ▪ destination: Match traffic to the same destination ▪ address: Match traffic by source or destination IP address ▪ transport: Match traffic by source or destination transport type ▪ port: Match traffic by source or destination transport port

Parameter	Description
description	A description for the flow record up to 256 characters long, including spaces.
collect	Configures data fields to be included a flow record. ▪ application name: Specify the application name as a non-key field in a flow record. ▪ application https url: Specify the HTTP/HTTPS application URL as a non-key field in a flow record. ▪ dns response-code: Specify the DNS parameters and DNS response code as a non-key field in the flow record. ▪ counter packets: Collect counter data for packets in the flow ▪ counter bytes: Collect counter data for bytes in the flow ▪ timestamp absolute first: Collect absolute timestamp of the first packet observed. ▪ timestamp absolute last: Collect absolute timestamp of the last packet observed.

Examples

Adding IPv4 and transport match fields to **flow-record-1**:

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# match ip source address
switch(config-flow-record)# match ip destination address
switch(config-flow-record)# match ip protocol
switch(config-flow-record)# match ip version
switch(config-flow-record)# match transport source port
switch(config-flow-record)# match transport destination port
switch(config-flow-record)# description Record used for basic ipv4 traffic analysis
```

Removing the IPv4 destination match field from the **flow-record-1**:

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# no match ipv4 destination address
```

Adding counter and timestamp collect fields to **flow-record-1**:

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# collect counter packets
switch(config-flow-record)# collect counter bytes
switch(config-flow-record)# collect timestamp absolute first
switch(config-flow-record)# collect timestamp absolute last
```

Add a application name to **flow record 1** as a collect field:

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# collect application name
switch(config-flow-record)# collect application https url
switch(config-flow-record)# collect application dns response-code
switch(config-flow-record)# collect application tls-attributes
```

Related Commands

Command	Description
flow exporter	Define how a flow monitor exports the flow reports.
flow monitor	Define a flow monitor configuration, including the flow exporter and flow record associated to that monitor.
show flow record	Display flow record configuration and status.

Command History

Release	Modification
10.14	The ipv4 parameter is deprecated and replaced with ip .
10.13	Added application https url and dns response-code parameters.
10.11	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config config-flow-record	Administrators or local user group members with execution rights for this command.

flow-tracking

```
flow-tracking
  enable
  icmp-ageout
  interface-flow-limit
  no ...
  tcp-ageout
  track icmp
  udp-ageout
```

Description

Configures flow tracking for TCP and UDP flows, and optionally, ICMP flows. The **no** form of this command deletes the flow tracking configuration context.

In order to optimize the flow removal process, flows that have aged-out are flushed in batches. A flow that has aged out is flushed only when the next batch processes. This can cause some flows to stay inactive for a slightly longer time than the value configured here.

Parameter	Description
enable	Enables flow tracking.
icmp-ageout	Configures an age-out time for ICMP flows, in seconds. Range: 10-

Parameter	Description
	86400. Default: 15.
interface-flow-limit	Configures global concurrent flow limit for flow tracking enabled interfaces. Range: 64-25000. Default: none.
tcp-ageout	Configures age-out time for established TCP flows in seconds. Range: 120-86400. Default: 600.
track icmp	Enable tracking of ICMP flows, in addition to the TCP/UDP flows tracked by default.
udp-ageout	Configures age-out time for established UDP flows in seconds. Range: 30-86400. Default: 30.

Examples

Configuring flow tracking:

```
switch(config) # flow-tracking
switch(config-flow-tracking) #
```

Deleting flow tracking:

```
switch(config) # no flow-tracking
switch(config) #
```

Enabling flow tracking:

```
switch(config) # flow-tracking
switch(config-flow-tracking) # enable
```

Disabling flow tracking:

```
switch(config) # flow-tracking
switch(config-flow-tracking) # no enable
```

Configuring an established ICMP flow age-out to 600 seconds:

```
switch(config) # flow-tracking
switch(config-flow-tracking) # icmp-ageout 600
```

Removing an established ICMP flow age-out of 600 seconds:

```
switch(config) # flow-tracking
switch(config-flow-tracking) # no icmp-ageout 600
```

Configuring an established TCP flow age-out to 1000 seconds:

```
switch(config)# flow-tracking  
switch(config-flow-tracking)# tcp-ageout 1000
```

Removing an established TCP flow age-out of 1000 seconds:

```
switch(config)# flow-tracking  
switch(config-flow-tracking)# no tcp-ageout 1000
```

Configuring an established UDP flow age-out to 1000 seconds:

```
switch(config)# flow-tracking  
switch(config-flow-tracking)# udp-ageout 1000
```

Removing an established UDP flow age-out of 1000 seconds:

```
switch(config)# flow-tracking  
switch(config-flow-tracking)# no udp-ageout 1000
```

Configuring global level interface flow limit to 256 interfaces:

```
switch(config)# flow-tracking  
switch(config-flow-tracking)# interface-flow-limit 256
```

Removing global level interface flow limit to 256 interfaces:

```
switch(config)# flow-tracking  
switch(config-flow-tracking)# no interface-flow-limit 256
```

Related Commands

Command	Description
IP source lockdown resource extended	no ip source-lockdown resource-extended must be disabled to enable flow-tracking

Command History

Release	Modification
10.14	The track icmp parameter is introduced.
10.13	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config	Administrators or local user group members with execution rights for this command.

ipv4 | ipv6 flow monitor (interface)

[no] ip|ipv6 flow monitor <name> in

Description

Enable flow monitoring on inbound and outbound interfaces by assigning a flow monitor to that interface. Only physical interfaces and LAG interfaces can be monitored. A flow monitor cannot be applied to an interface that is part of a LAG. If an unsupported application is attempted, an error message will be displayed. If the flow monitor is associated with a flow record that contains application fields as collect fields, then Application Recognition should be enabled on the same interface.

The **[no]** form of command disables the flow monitoring.

Examples

Associate a flow monitor configuration named **flow-monitor-1** and **flow-monitor-2** for IPv4 or IPv6 traffic respectively on physical interface.

```
switch(config)# interface 1/1/1
switch(config-if)# ip flow monitor flow-monitor-1 in
switch(config-if)# ipv6 flow monitor flow-monitor-2 in
```

Associate a flow monitor configuration named **flow-monitor-3** and **flow-monitor-4** for IPv4 or IPv6 traffic respectively on a Lag interface.

```
switch(config)# interface lag 1
switch(config-lag-if)# ip flow monitor flow-monitor-3 in
switch(config-lag-if)# ipv6 flow monitor flow-monitor-4 in
```

Related Commands

Command	Description
flow exporter	Define how a flow monitor exports the flow reports.
flow record	Define data to be included in a flow record by configuring flow record match and collect fields
flow monitor	Define a flow monitor configuration, including the flow exporter and flow record associated to that monitor.

Command History

Release	Modification
10.11	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config config-flow-monitor	Administrators or local user group members with execution rights for this command.

ipv4 | ipv6 flow monitor (role)

[no] ip|ipv6 flow monitor <name>

Description

Enable flow monitoring on a role. The authorization status of a client does not depend on the flow monitor status for the client in hardware. The client will be authorized even if the system is not able to apply the flow monitor configuration in the role.

In case of multiple clients onboard to a port with varied flow monitor configurations, the flow monitor associated with the first authenticated client on the port will be applied for all the traffic on the port.

The **[no]** form of command removes the flow monitor from the role.

Parameter	Description
<name>	Name of the flow monitor , up to 64 characters.

Examples

Enable an IPv4 flow monitor on a role named role01

```
switch# config terminal
switch(config)# port-access role role01
switch(config-pa-role)# ip flow monitor flow-monitor-1
```

Enable an IPv6 flow monitor on a role named role01

```
switch# config terminal
switch(config)# port-access role role01
switch(config-pa-role)# ipv6 flow monitor flow-monitor-2
```

Related Commands

Command	Description
flow exporter	Define how a flow monitor exports the flow reports.
flow record	Define data to be included in a flow record by configuring flow record match and collect fields
flow monitor	Define a flow monitor configuration, including the flow exporter and flow record associated to that monitor.

Command History

Release	Modification
10.14	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config config-flow-monitor	Administrators or local user group members with execution rights for this command.

show flow exporter

`show flow exporter [<name>] [statistics]`

Description

Displays flow exporter statistics, configuration and status. When no exporter name is specified, the output of this command displays information for all flow exporters.

The output of this command can indicate the following status types:

- Accepted
- Rejected (Internal error: exporter does not exist)
- Rejected (Internal error: destination type does not exist)
- Rejected (Destination type is hostname or IP address, but no destination is specified)
- Rejected (Destination type is hostname or IP address, but the specified hostname or IP address is invalid)
- Rejected (Destination type is Traffic Insight, but no destination is specified)
- Rejected (Destination type is Traffic Insight, but the specified Traffic Insight instance does not exist)
- Rejected (Destination type is Traffic Insight, but the specified Traffic Insight instance is not enabled)
- Rejected (Destination type is Traffic Insight, but the specified Traffic Insight instance source is not IPFIX)
- Rejected (Internal error: destination type is Traffic Insight, but the specified Traffic Insight instance is invalid)

Parameter	Description
<name>	Name of the flow exporter.
statistics	The <code>statistics</code> parameter adds statistical information about the flow exporter to the output.

Examples

Display the configuration of a flow exporter named **exporter-1**.

```
switch# show flow exporter exporter-1
-----
Flow exporter 'exporter-1'
```

```

-----
Description          : Exports to the first collector
Status               : Accepted
Export Protocol      : ipfix
Destination Type     : Hostname or IP address
Destination          : 192.168.0.1
Transport Configuration
  Protocol           : UDP
  Port               : 9995

```

Display statistics information for all flow exporters

```
switch# show flow exporter exporter-1 statistics
```

```
-----
Flow exporter 'exporter-1'
-----
```

```
Reports sent          : 14961
-----
```

```
Flow exporter 'exporter-2'
-----
```

```
Reports sent          : 5
```

Related Commands

Command	Description
flow exporter	Define how a flow monitor exports the flow reports.

Command History

Release	Modification
10.11	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config config-flow-exporter	Administrators or local user group members with execution rights for this command.

show flow monitor

```
show flow monitor [<name>]
```

Description

Displays flow monitor configuration and status. When no monitor name is specified, the output of this command displays information for all flow monitors.

The output of this command can indicate the following status types:

- Accepted
- Rejected (Internal error: monitor does not exist)
- Rejected (The state of one or more of the assigned flow exporters is rejected)

Parameter	Description
<name>	Name of the flow monitor.

Examples

Display the configuration of a flow monitor named **flow-monitor-1**.

```
switch# show flow monitor monitor-1
-----
Flow monitor 'monitor-1'
-----
Description           : Used for IPv4 traffic analysis
Status                : Accepted
Flow Exporter(s)      : exporter-1, exporter-2
```

Related Commands

Command	Description
flow monitor	Define a flow monitor configuration, including the flow exporter and flow associated to that monitor.

Command History

Release	Modification
10.11	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config config-flow-exporter	Administrators or local user group members with execution rights for this command.

show flow record

```
show flow record [<name>]
```

Description

Display flow record configuration and status. When no record name is specified, the output of this command displays information for all flow records.

The output of this command can indicate the following status types:

- Accepted
- Rejected (Internal error: failed to process record)
- Rejected (Mix of IPv4 and IPv6 match fields is not allowed. Specify match fields of the same IP version (IPv4 or IPv6))

Parameter	Description
<name>	Name of the flow record.

Examples

Display the configuration of a flow record named **flow-record-1**.

```
switch# show flow record record-1
-----
Flow record  'record-1'
-----
Description           : Used for IPv4 traffic analysis
Status                : Accepted
Match Fields
  ipv4 destination address
  ipv4 protocol
  ipv4 source address
  ipv4 version
  transport destination port
  transport source port
Collect Fields
  application name
  counter bytes
  counter packets
```

Display the information of a specific flow record.

```
switch# show flow record record-1
-----
Flow record  'record-1'
-----
Description           : Used for IPv4 traffic analysis
Status                : Accepted
Match Fields
  ipv4 destination address
  ipv4 protocol
  ipv4 source address
  ipv4 version
  transport destination port
  transport source port
Collect Fields
  application name
  counter bytes
  counter packets
```

Display information for all flow records

```

switch# show flow record
-----
Flow record  'record-1'
-----
Description          : Used for IPv4 traffic analysis
Status               : Accepted
Match Fields
    ipv4 destination address
    ipv4 protocol
    ipv4 source address
    ipv4 version
    transport destination port
    transport source port
Collect Fields
    application name
    counter bytes
    counter packets
-----
Flow record  'record-2'
-----
Description          : Used for IPv6 traffic analysis
Status               : Accepted
Match Fields
    ipv6 destination address
    ipv6 protocol
    ipv6 source address
    ipv6 version
    transport destination port
    transport source port
Collect Fields
    application name
    counter bytes
    counter packets
...

```

Display information with no flow records configured

```

switch# show flow record
No flow records configured

```

Related Commands

Command	Description
flow record	Define data to be included in a flow record by configuring flow record match and collect fields

Command History

Release	Modification
10.11	Command introduced on 6400, 6400, 8100, and 8360 Switch series.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config config-flow-exporter	Administrators or local user group members with execution rights for this command.

show flow-tracking

show flow-tracking

Description

Displays flow-tracking and statistics collection configurations and status.

Examples

Display the configuration of role based flow tracking.

```
switch(config)# show flow-tracking
Flow Tracking Global Configuration
Configuration status          : Enabled
Operational status           : Enabled
Failure Reason                : NA
UDP Ageout                    : 30   (Seconds)
TCP Ageout                    : 600  (Seconds)
ICMP Ageout                   : 15   (Seconds)
Interface Flow limit          : None
Tracked Protocols             : TCP, UDP
Statistics Collection
  Configuration Status        : Enabled
  Operational Status          : Enabled
  Failure Reason              : NA
Flow Tracking Port Configuration
Interface      App Recognition  Reflexive ACL      IPFIX      Operation
Status
-----
-
1/1/1          Enabled          Disabled           Enabled     Enabled
1/1/2          Enabled          Disabled           Disabled    Enabled
1/1/3          Enabled          Disabled           Disabled    Enabled
1/1/4          Enabled          Disabled           Disabled    Enabled
1/1/5          Enabled          Disabled           Disabled    Enabled
1/1/6          Enabled          Disabled           Disabled    Enabled
1/1/7          Enabled          Disabled           Enabled     Enabled
1/1/8          Enabled          Disabled           Disabled    Enabled
1/1/9          Enabled          Disabled           Disabled    Enabled
1/1/10         Disabled          Disabled           Disabled
Disabled
1/1/13         Enabled          Disabled           Enabled     Enabled
1/1/14         Enabled          Disabled           Disabled    Enabled
1/1/15         Enabled          Disabled           Disabled    Enabled
1/1/16         Enabled          Disabled           Disabled    Enabled
1/1/17         Enabled          Disabled           Disabled    Enabled
1/1/18         Enabled          Disabled           Disabled    Enabled
1/1/19         Enabled          Disabled           Disabled    Enabled
1/1/20         Enabled          Disabled           Disabled    Enabled
1/1/21         Enabled          Disabled           Disabled    Enabled
1/1/23         Enabled          Disabled           Disabled    Enabled
1/1/24         Enabled          Disabled           Enabled     Enabled
```

1/1/28
Disabled

Disabled

Disabled

Disabled

Related Commands

Command	Description
IP source lockdown resource extended	IP source lockdown must be disabled with the no ip source-lockdown resource-extended command before enabling flow-tracking

Command History

Release	Modification
10.14	Added information related to role based IPFIX.
10.14	The output of this command includes ICMP ageout information.
10.13	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	config	Administrators or local user group members with execution rights for this command.

show tech ipfix

show tech ipfix

Description

Shows the IPFIX configuration settings.

If applicable source IP address or source interface is configured for the IPFIX protocol, then that configuration is used.

If a valid source is configured, the exporter sends flows to an external collector using the effective configured source IP address as the source IP address of the flow packets. In the context of this application, a valid source IP address is any IP address configured in the exporter's VRF namespace.

Examples

The example shows the IPFIX configuration settings.

```
switch#show tech ipfix
=====
Show Tech executed on Tue Apr 11 02:43:06 2023
=====
```

```

=====
[Begin] Feature ipfix
=====
*****
Command : show flow exporter
*****
-----
Flow exporter 'ipfix'
-----
Status                : Accepted
Export Protocol        : ipfix
Destination Type       : Traffic Insight
Destination            : t1
Transport Configuration
Protocol               : udp
Port                  : 4739
-----
Flow exporter 'V6E1'
-----

....

=====
[End] Feature ipfix
=====

```

Command History

Release	Modification
10.11	Command introduced on 6400, 6400, 8100, and 8360 Switch series.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	Manager (#)	Administrators or local user group members with execution rights for this command.

boot fabric-module

`boot fabric-module <SLOT-ID>`

Description

Reboots the specified fabric module.

Parameter	Description
<code><SLOT-ID></code>	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
6300 6400	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
<code><SLOT-ID></code>	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

This command is supported on switches that have multiple line modules.

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
6300 6400	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
<i>active</i>	Selects the active management module.
<i>standby</i>	Selects the standby management module.
<i><SLOT-ID></i>	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 is supported on switches that have multiple management modules.

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 Packard 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
6300 6400	Manager (#)	Administrators or local user group members with execution rights for this command.

boot management-module (recovery console)

boot management-module {local|remote}

Description

Reboots the specified management module by specified location (local or remote).

Parameter	Description
<local>	Reboots the local management module.
<remote>	Reboots the remote management module.

Usage

This command is supported on switches that have multiple management modules.

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 Packard 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

Booting a remote management module:

```
switch# boot management-module remote
There is no other management module installed.
Aborting.
switch#
```

Command History

Release	Modification
10.12	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400	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.
The system is going down for reboot.
```

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

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

Do you want to save the current configuration (y/n)? n

This will reboot the entire switch and render it unavailable
until the process is complete.
Continue (y/n)? y
The system is going down for reboot.
```

Canceling a system reboot:

```
switch# boot system

Do you want to save the current configuration (y/n)? n

This will reboot the entire switch and render it unavailable
```

```
until the process is complete.  
Continue (y/n)? n  
Reboot aborted.  
switch#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

show boot-history

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

Description

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

For switches that support line modules (such as 6400 switch series) including the **all** parameter displays information for the active management module and all available line modules.



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

Parameter	Description
all	Optional. Shows boot information for the active management module. For switches that support line modules, including this parameter displays information for and all available line modules.
vsf member <1-10>	Optional. Display boot history for the specified VSF member

Usage

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

The output of this command includes the following information:

Parameter	Description
Index	The position of the boot in the history file. Range: 0 to 3.
Boot ID	A unique ID for the boot . A system-generated 128-bit string.
Current Boot, up for <time>	For the current boot, the show boot-history command shows the number of seconds the module has been running on the current software.
<Timestamp>: boot reason	For previous boot operations, the show boot-history command shows the time at which the operation occurred and the reason for the boot. The reason for the boot is one of the following values: ▪ <DAEMON-NAME> crash: The daemon identified by <DAEMON-NAME> caused the module to boot. ▪ Kernel crash: The operating system software associated with the module caused the module to boot. ▪ Uncontrolled reboot: The reason for the reboot is not known. ▪ Reboot requested through database: The reboot occurred because of a request made through the CLI or other API. For details, see show boot-history

Table 1: Description of reboots handled through the database

Boot History String	Description
Reboot requested by user	A user requested a switch reboot through the CLI or web UI.
Reset button pressed	The switch detected a short-press of the reset button
Backplane fault	A backplane fault occurred.
Configuration change	A configuration change resulted in a reboot.
Configuration version migration	A configuration version migration occurred which required a reboot.
Console error	The console failed to start.
Fabric fault	A fabric fault occurred.
All line modules faulted	A zero line card condition occurred.
Redundancy switchover requested	A user requested a redundancy switchover.
Redundant Management communication timeout	The standby management module has taken over from an unresponsive active management module.
Redundant Management	A failure to elect a standby management module in the allotted

Boot History String	Description
election timeout	time.
Critical service fault (error)	A daemon critical to switch operation has stopped functioning. An extra error string may be present to describe the error in detail.
VSF autojoin renumber	Reset triggered by VSF autojoin.
VSF member renumbered	A user requested a renumber of a VSF member.
VSF switchover requested	A user requested a VSF switchover.
VSX software update	Reset triggered by a VSX software update.
Chassis critical temperature	Chassis operating temperature exceeded.
Chassis low critical temperature	Chassis temperature below the minimum operating threshold.
Chassis insufficient fans	Insufficient fans to cool the chassis.
Chassis unsupported PSUs/fans	Unsupported or misconfigured PSUs or system fans.
Management module critical temperature	Management module operating temperature exceeded.
ISSU SMM update	Standby management module reboot triggered by an In-Service Software Upgrade (ISSU).
ISSU switchover	Redundancy switchover triggered by an In-Service Software Upgrade.
ISSU aborted	Standby management module reset triggered by failure during an In-Service Software Upgrade.
Rollback timer expired	Reset triggered by the ISSU rollback timer expiring.

Examples

Showing the boot history of the active management module:

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

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

Index : 1
Boot ID : bfba9bc486304e57904ac717a0ccbdcd
02 Sep 23 02:55:33 : CPU request reset with 0x20201, Version: FL.10.14.0000-1619-
ga9ec1805bd442~dirty
02 Sep 23 02:55:33 : Switch boot count is 2

Index : 0
Boot ID : a88a71b7ca9a4574af7e3b811ddfdc7e
02 Sep 23 02:49:26 : Reboot requested by user, Version: FL.10.14.0000-1619-
```

```

ga9ec1805bd442~dirty
02 Sep 23 02:50:02 : Switch boot count is 1

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

```

(For 6400 Switch series) Showing the boot history of the active management module and all line modules:

```

switch#
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
...

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
...
```

In the event of a reset triggered by a power supply unit (PSU), or a PSU input fault, the output of this command also displays information about why the PSU initiated a reboot. The following example displays the boot history of a switch with a reboot initiated by a PSU.

```
switch# show boot-history

Management module
=====
Index : 2
Boot ID : a61ad00d10864c748bc7893a5d4af2e4
15 Dec 23 19:02:02 : Power on reset with 0x1, Version: FL.10.13.1000AF

15 Dec 23 19:02:02 : Switch boot count is 0
15 Dec 23 19:02:17 : PSU 1/1: Fault detected

Index : 1
Boot ID : 30d831bbfdfa425baf50a629ee01b185
15 Dec 23 19:01:58 : Power on reset with 0x1, Version: FL.10.13.1000AF
15 Dec 23 19:01:58 : Switch boot count is 0
```

The following example displays the boot history for the VSF member 2.

```
switch# show boot-history vsf member 2

Member-2
=====

Index : 0
Boot ID : df99026c194a44f1944a3e7685fb4d90
Current Boot, up for 3 hrs 31 mins 39 secs

Index : 3
Boot ID : 7bf4104903fe4ad1ba4bce40e8099c76
10 Aug 17 10:02:24 : Reboot requested through database
10 Aug 17 10:02:13 : Switch boot count is 2
```

Command History

Release	Modification
10.13.1000	The output of this command is enhanced to display additional information about the reason for the reboot, if available.
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 5

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

```
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
6300 6400	config-external-storage- <i><VOLUME-NAME></i>	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
<i><DIRECTORY-NAME></i>	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
6300 6400	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

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
6300 6400	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

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
6300 6400	config-external-storage- <i><VOLUME-NAME></i>	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
6300 6400	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
6300 6400	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
6300 6400	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
6300 6400	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

type

```
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.

Parameter	Description
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
6300 6400	config-external-storage- <i><VOLUME-NAME></i>	Administrators or local user group members with execution rights for this command.

username

```
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
<i><USER-NAME></i>	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 nassuser
```

Clearing the user name nassuser 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
6300 6400	config-external-storage- <i><VOLUME-NAME></i>	Administrators or local user group members with execution rights for this command.

vrf

```
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
<i><VRF-NAME></i>	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
6300 6400	config-external-storage-<VOLUME-NAME>	Administrators or local user group 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

Parameter	Description
	604800.
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
6300 6400	config-ip-sla-<IP-SLA-NAME>	Administrators or local user group members with execution rights for this command.

https

```
https {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 <<PROBE-INTERVAL>>] [version <VERSION-NUMBER>] [https-raw-request
<RAW-PAYLOAD>]
no https {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 <<PROBE-INTERVAL>>] [version <VERSION-NUMBER>] [https-raw-request
<RAW-PAYLOAD>]
```

Description

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

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



For HTTPS IP-SLA sessions, it is not required to install a certificate on the switch.

Parameter	Description
{get raw}	Selects HTTPS request type as get or raw where the system will generate or provide HTTPS payload.
URL	Specifies HTTPS URL address of syntax. https://<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 HTTPS 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.
version <VERSION-NUMBER>	Specifies the source interface to use for sending IP-SLA probes.
https-raw-request <RAW-PAYLOAD>	HTTPS raw request. String.

Examples

```

switch(config-ipsla-1)# https get https://device.arubanetworks.com/root/home.html
switch(config-ipsla-1)# https get https://2.2.2.2 source 1/1/1
switch(config-ipsla-1)# https get https://device.arubanetworks.com source 2.2.2.1
switch(config-ipsla-1)# https get https://device.arubanetworks.com/root/home.html
source-interface 1/1/1
switch(config-ipsla-1)# https get https://device.arubanetworks.com name-server
10.10.10.2
switch(config-ipsla-1)# https raw https://device.arubanetworks.com/root/home.html
raw-request "GET /en/US/hmpgs/index.html"
switch(config-ipsla-1)# no https get https://2.2.2.2 source 1/1/1
switch(config-ipsla-1)# no https raw
https://device.arubanetworks.com/root/home.html raw-request "GET
/en/US/hmpgs/index.html"

```

Command History

Release	Modification
10.12.1000	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400	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.
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
6300 6400	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 64 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
6300 6400	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. Length: 1 to 64 characters.
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
6300 6400	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
```

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

```
SLA Name       : 1 (non-persistent)
SLA Type       : udp-echo
VRF Name       : default
Responder Port  : 10
Responder IP    :
Responder Interface :
Responder Status : Running
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400	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
6300 6400	config	Administrators or local user group members with execution rights for this command.

show ip-sla

```
show ip-sla {<SLA-NAME> [results] | all}
```

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.
all	Shows all ip-sla source configurations and status.

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)                  : 12
  Maximum RTT(ms)                  : 12
  Average RTT(ms)                  : 12
  DNS RTT(ms)                      : 0
  TCP RTT(ms)                      : 12

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)                 : 1
    Maximum RTT(ms)                 : 1
    Average RTT(ms)                 : 1
    DNS RTT(ms)                     : 0

    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 https-sla**

```

SLA Name         : https-sla
Status           : running
SLA Type         : https
VRF              : default
Source Port      : 1027
Source IP        : 1.1.1.1
Source Interface :
Domain Name Server :
Probe Interval(seconds) : 60
HTTPS Request Type : raw
HTTPS URL        : https://1.1.1.2
Cache            : Enabled
HTTPS Proxy URL  :
HTTP Version Number :

```

switch(config)# **show ip-sla all**

```

IP-SLA session status
IP-SLA Name           : 707 (non-persistent)
IP-SLA Type           : https
Destination Host Name/IP Address : NA
Destination Port      : NA
Source IP Address/IFName :
Source Port           :
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              : 2023 Jun 05 13:10:19
Packets Sent                  : 1
Packets Received              : 1
Packet Loss in Test           : 0.0000%
```

```
Minimum RTT(ms)              : 20
Maximum RTT(ms)               : 20
Average RTT(ms)               : 20
DNS RTT(ms)                   : 0
TCP RTT(ms)                   : 12
TLS RTT(ms)                   : 8
```

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<br>as source port not available.                  | Total numbers of probes transmission failed |
| Destination Address Unreachable<br>due to route unavailable. | Total numbers of probes transmission failed |
| DNS Resolution Failures<br>resolution failure.               | Total numbers of probes failed due to DNS   |
| Reception Error<br>internal error in reception.              | Total numbers of probes failed due to       |
| Transmission Error<br>internal errr in transmission.         | Total numbers of probes failed due to       |

## Command History

| Release          | Modification                                       |
|------------------|----------------------------------------------------|
| 10.12.1000       | Updated to display <b>https</b> as an IP-SLA type. |
| 10.07 or earlier | --                                                 |

## Command Information

| Platforms    | Command context             | Authority                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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.                                           |

| Parameter                       | Description                                    |
|---------------------------------|------------------------------------------------|
| [<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                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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. |

| Parameter                                | Description                                                                                              |
|------------------------------------------|----------------------------------------------------------------------------------------------------------|
| <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                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-ip-sla-<IP-SLA-NAME> | Administrators or local user group members with execution rights for this command. |

## vrf

```
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                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-ip-sla-<IP-SLA-NAME> | Administrators or local user group members with execution rights for this command. |

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

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

## Limitations with speed downshift

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

## L1-100Mbps downshift commands

### downshift enable

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

#### Description

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

#### Examples

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

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

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

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

When automatic downshift is enabled:

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

Disabling automatic speed downshift:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## show interface

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

## Description

Shows active configurations and operational status information for interfaces.

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

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

Showing interface information when it 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
Energy-Efficient Ethernet is enabledMDI mode: MDIX
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     |

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

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

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

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

Showing information when the interface is 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     |

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

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

| EEE    |        | PoE Power | Link        | Admin       | Speed       |        | Flow-Control |        |
|--------|--------|-----------|-------------|-------------|-------------|--------|--------------|--------|
| Port   | Type   | (Watts)   | Status      | Config      | Status      | Config | Status       | Config |
| Status | Config |           | State       | Information | Description |        |              |        |
| 1/1/1  | 1GbT   | --        | up          | up          | 1G          | auto   | off          | off    |
| on     | on     |           | 10M/100M/1G |             | --          |        |              |        |

Showing the monitor information:



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

| Interface 1/1/1 is up |          |          |               |
|-----------------------|----------|----------|---------------|
| Rate                  | RX       | TX       | Total (RX+TX) |
| Mbits / sec           | 30196.43 | 30196.43 | 60392.85      |
| MPkts / sec           | 58977.39 | 58977.40 | 117954.79     |
| Unicast               | 0.00     | 0.00     | 0.00          |
| Multicast             | 58977.39 | 58977.40 | 117954.79     |

|                  |              |              |               |
|------------------|--------------|--------------|---------------|
| Broadcast        | 0.00         | 0.00         | 0.00          |
| Utilization %    | 75.49        | 75.49        | 150.98        |
| Statistic        | RX           | TX           | Total (RX+TX) |
| -----            | -----        | -----        | -----         |
| Packets          | 4756527649   | 4756527865   | 9513055514    |
| Unicast          | 0            | 0            | 0             |
| Multicast        | 4756527649   | 4756527865   | 9513055514    |
| Broadcast        | 2            | 0            | 2             |
| Bytes            | 304417778668 | 304417795428 | 608835574096  |
| Jumbos           | 0            | 0            | 0             |
| Dropped          | 0            | 19028847730  | 19028847730   |
| Pause Frames     | 0            | 0            | 0             |
| Errors           | 0            | 0            | 0             |
| CRC/FCS          | 0            | n/a          | 0             |
| help: ?, quit: q |              |              |               |

```

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

```

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



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

```

switch(config-if)# show interface 1/1/1 human-readable
Interface 1/1/1 is up
Rate

Bits / sec 3M 3M 6M
Pkts / sec 316 316 633
Unicast 319 319 638
Multicast 0 0 0
Broadcast 0 0 0
Utilization % < 1 < 1 < 1
Statistic RX TX Total

Packets 577K 577K 1M
Unicast 577K 577K 1M
Multicast 0 51 51
Broadcast 0 15 15
Bytes 744M 745M 1G
Jumbos 0 0 0
Dropped 0 0 0
Filtered 0 0 0
Pause Frames 0 0 0
Errors 0 0 0
CRC/FCS 0 n/a 0
Collision n/a 0 0
Runts 0 n/a 0

```

## Showing information about extended counters:



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

```
switch(config-if)# show interface 1/1/17 extended

Interface 1/1/17

Statistics Value

Dot1d Tp Port In Frames 547
Dot1d Tp Port Out Frames 608
Dot3 In Pause Frames 0
Dot3 Out Pause Frames 0
Ethernet Stats Broadcast Packets 19
Ethernet Stats Bytes 40162
Ethernet Stats Packets 342
...

Error-Statistics Value

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

## Showing interface link-status:

```
switch# show interface link-status

Port Type Physical Link Last
Link State Transitions Change

1/1/1 1G-BT down 0 --
1/1/2 1G-BT up 1 1 minute ago (Fri Mar 09
12:36:56 UTC 2018)
1/1/3 1G-BT up 1 1 minute ago (Fri Mar 09
12:36:56 UTC 2018)
1/1/4 -- down 0 --
1/1/5 -- down 0 --
```

## Showing interface loopback 1 link-status:

```

Port Type Physical Link Last
Link State Transitions Change

```

```

loopback1 -- up -- --
```

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

```

Port Type Physical Link State Link Transitions Last Change

1/1/2 1G-BT up 1 1 minute ago (Fri Mar 09
12:36:56 UTC 2018)
1/1/3 1G-BT up 1 1 minute ago (Fri Mar 09
12:36:56 UTC 2018)
```

Showing interface link-status:

```
switch# show interface link-status

Port Type Physical Link State Link Transitions Link Flaps Ignored Last Change

1/1/1 1G-BT down 0 0 --
1/1/2 1G-BT up 1 0 1 minute ago
(Fri Mar 09 12:36:56 UTC 2018)
1/1/3 1G-BT up 1 0 1 minute ago
(Fri Mar 09 12:36:56 UTC 2018)
1/1/4 -- down 0 0 --
1/1/5 -- down 0 0 --
```

## Command History

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

## Command Information

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

## show interface statistics

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

```
show interface lag [<LAG-ID>] statistics monitor [non-zero] [human-readable]
show interface lag [<LAG-ID>] error-statistics [non-zero] [human-readable]
show interface lag [<LAG-ID>] error-statistics monitor [non-zero] [human-readable]
show interface vxlan <VXLAN-ID> statistics [non-zero] [human-readable]
```

## Description

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

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

## Examples

Showing statistics of all interfaces:

```
show interface statistics
```

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

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

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

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

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

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

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

Showing statistics of a single interfaces:

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

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

Showing statistics of all members of a LAG interface:

```
show interface lag1 statistics
```

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

Showing error statistics of all interfaces:

```
show interface error-statistics
```

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

Showing monitor statistics:



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

Showing monitor error statistics in human-readable format:

## Command History

| Release          | Modification                    |
|------------------|---------------------------------|
| 10.11            | Added moitor parameter.         |
| 10.10            | Added human-readable parameter. |
| 10.07 or earlier | --                              |

## Command Information

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

## show interface downshift-enable

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

### Description

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

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

### Examples

Showing automatic downshift information:

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

Port Downshift Speed
 Enabled | Active Status | Config

1/1/1 yes yes 100M-FDx auto
1/1/2 yes no 1G auto
1/1/3 yes no 100M-FDx 100M-FDx
1/1/4 no no -- auto
```

Showing automatic downshift information on per interface:

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

Port Downshift Speed
 Enabled | Active Status | Config

1/1/2 yes no 1G auto
```

### Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

### Command Information

| Platforms    | Command context | Authority                                                                                                                                                              |
|--------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6300<br>6400 | config          | Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only. |

## show running-config interface

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

### Description

Displays active configurations of various switch interfaces.

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

### Examples

Showing 1/1/2 interface configuration:

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

interface 1/1/2
 no shutdown
 description DC-23
 exit
```

Showing loopback interfaces configured:

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

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

Showing loopback interfaces not configured:

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

No loopback interfaces configured.
```

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                                                                                                              |
|--------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6300<br>6400 | config          | Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only. |

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

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

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



---

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

---

## Mirror statistics

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

## Classifier policies and mirroring sessions

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

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

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

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

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

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

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

## VLAN as a source

AOS-CX allows configuration of VLAN as a mirroring source. When a VLAN source is configured in the 'rx' direction, all packets are mirrored as they are received in the switch. When a VLAN source is configured in 'tx' direction, all packets are mirrored as they are transmitted out of the switch.

More than one source VLAN can be configured in a mirror session. Each such VLAN may specify its own direction.

There is a limit of 1024 source VLANs in each direction of a given mirror session.

Same VLANs can be configured as a mirror source for multiple sessions.

Note: When changing a source VLAN in an enabled mirror session (that is, adding, changing direction, or removing), mirrored packets being transmitted out the mirror destination port from other mirror sources may be briefly interrupted during the reconfiguration.

Direction of an existing source VLAN can be updated in one of two ways:

1. Reenter the `source vlan` command with the new preferred direction.
2. Use the `no` form of the command with a direction (rx or tx) to selectively remove the specified direction. Specifying the last remaining direction for that VLAN will remove the VLAN from the configuration entirely.

For packets bridged through the switch:

If the mirror is configured in 'both' direction, two copies of packets are mirrored, otherwise one copy of the packet will be mirrored.

For routed packets:

- If the mirror is configured in the 'rx' direction, packets are mirrored in the pre-routed form with the destination MAC address as the switch address.
- If the mirror is configured in the 'tx' direction, packets are mirrored in the post-routed form with the source MAC as the switch address. Destination MAC is the nexthop gateway or station.
- If the mirror is configured in the 'both' direction, one copy of the packet will be mirrored.

Control plane packets generated by the switch's CPU are processed both in the ingress and the egress packet processing pipeline. The following are the behaviors for mirroring with VLAN as source:

- If the mirror is configured in the 'rx' or 'tx' direction, the packets are mirrored to the mirror destination.
- If the mirror is configured in the 'both' direction, two copies of the packets are mirrored to the mirror destination.

## 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. |

## clear mirror endpoint

```
clear mirror endpoint [<NAME>]
```

## Description

Clears mirror endpoint statistics for all configured mirror endpoints. The optional parameter can be added to clear a specific mirror endpoint.

| Parameter | Description                                                   |
|-----------|---------------------------------------------------------------|
| <NAME>    | Specifies name of the mirror endpoint instance to be cleared. |

## Examples

Clearing statistics for all configured mirror endpoints:

```
switch# clear mirror endpoint
```

Clearing mirror statistics for mirror endpoint test:

```
switch# clear mirror endpoint test
```

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context             | Authority                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | Operator (>) or Manager (#) | Administrators or local user group members with execution rights for this command. |

## comment

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

## Description

Specifies a comment for the mirroring session.

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

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

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

## Usage

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

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

## Examples

Adding a comment to a mirror session:

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

Removing the comment from mirror session 3:

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

Adding a comment to a mirror endpoint:

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

Replacing the existing comment for mirror endpoint:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms     | Command context                                      | Authority                                                                          |
|---------------|------------------------------------------------------|------------------------------------------------------------------------------------|
| All platforms | config-mirror-<SESSION-ID><br>config-mirror-endpoint | 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 password:
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 successfully.
```

## Command History

| Release | Modification       |
|---------|--------------------|
| 10.08   | Command introduced |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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                                                                                                                                                          |
|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>&lt;REMOTE-URL&gt;</code>   | Specifies the capture file on a remote TFTP or SFTP server. The URL syntax is:<br>{tftp://   sftp://<USER>@} {<IP>   <HOST>} [:<PORT>]<br>[;blocksize=<SIZE>]/<FILE> |
| <code>vrf &lt;VRF-NAME&gt;</code> | 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                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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                                                                          |
|--------------|----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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

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

*On the 6400 Switch Series, interface identification differs.*

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

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

Replacing the existing destination with different interface:

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

Removing a destination:

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

| Switch | Destination interface limit per mirror session (4 possible sessions) |
|--------|----------------------------------------------------------------------|
| 6300   | 64                                                                   |
| 6400   | 64                                                                   |

## 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> source <SOURCE-IPv4-ADDR>
dscp <DSCP-VALUE> vrf <VRF-NAME>
```

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.

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

## Examples

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

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

Replacing the existing tunnel destination:

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

Replacing the existing destination with a different DSCP value:

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

Removing the destination:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context            | Authority                                                                          |
|--------------|----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | Manager (#)     | Administrators or local user group members with execution rights for this command. |

## disable

disable

### Description

Disables the mirroring session specified by the current command context.

### Usage

By default, mirroring sessions are disabled.

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

### Example

Disabling a mirroring session:

```

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

```

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

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

## enable

enable

## Description

Enables the mirroring session for the current command context.

## Usage

By default, mirroring sessions are disabled.

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

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

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



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

## Example

*On the 6400 Switch Series, interface identification differs.*

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 | <code>config-mirror-&lt;SESSION-ID&gt;</code> | 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                                     |
|---------------------------------|-------------------------------------------------|
| <code>&lt;SESSION-ID&gt;</code> | 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 | <code>config</code> | Administrators or local user group members with execution rights for this command. |

## mirror endpoint

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

### Description

Creates the specified mirror endpoint or enters its context if it already exists. The specifics of a mirror endpoint are created or altered while in the mirror endpoint context and the mirror endpoint is enabled

or disabled from this context. It may be possible to support different encapsulations by different ASICs. For example, UDP for PVOS compatibility. Termination of GRE encapsulation is also supported. The **no** form of this command removes an existing mirror endpoint. An enabled mirror endpoint is automatically disabled first before removal.

| Parameter | Description                     |
|-----------|---------------------------------|
| <NAME>    | Specifies mirror endpoint name. |

## Examples

Creating a mirror endpoint named test :

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

Deleting mirror endpoint named test:

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

Configuring a mirror endpoint named test :

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



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



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

## Command History

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

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | 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                                                                                                                                                                                                                      |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>&lt;SESSION-ID&gt;</b> | Specifies the session identifier. Range: 1 to 4                                                                                                                                                                                  |
| <b>vsx-peer</b>           | 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

*On the 6400 Switch Series, interface identification differs.*

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. |

## show mirror endpoint

show mirror endpoint [<NAME>]

### Description

Shows a list of all configured mirror endpoints, their Admin Status and their Operation Status.

The optional parameter will display the details of the specified mirror endpoint if it exists.

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

## Examples

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

```
switch# show mirror endpoint
Name Admin Status Operation Status

test enable enabled
monitor disable disabled
```

Showing the details of enabled mirror endpoint test:

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



"Output Packets" in "show mirror endpoint [name]" is only supported for statistics.

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context             | Authority                                                                          |
|--------------|-----------------------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | Operator (>) or Manager (#) | Administrators or local user group members with execution rights for this command. |

## shutdown

```
shutdown
no shutdown
```

## Description

Enables mirror endpoint from its default disabled state. To verify the mirror endpoint was successfully activated, run the `show mirror endpoint NAME` command and verify that the **Admin Status** and

**Operational Status** has changed from disabled to enabled. If the status value remains disabled, consult the system logs to determine the reason for activation failure. To disable the mirror endpoint, first disable the remote mirror session on the switch that's originating the data. Next, use the `shutdown` command to disable the mirror endpoint.

## Examples

Enabling a mirror endpoint:

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

Disabling a mirror endpoint:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config          | Administrators or local user group members with execution rights for this command. |

## source

```
source <SOURCE-IP> destination <DESTINATION-IP> id <1-4294967295> [vrf <VRF_NAME>] [type {gre}]
no source
```

## Description

Configures tunnel parameters of the mirror endpoint. Configuring a tunnel parameter to a mirror endpoint will replace the existing configuration. By default the VRF is **default**, users can also explicitly provide a custom VRF. The default tunnel type is considered to be GRE and users also have the option to explicitly give type as GRE.

The **no** form removes the tunnel parameters of the mirror endpoint.

| Parameter        | Description                                                     |
|------------------|-----------------------------------------------------------------|
| <SOURCE-IP>      | Specifies L3 encapsulated IPv4 source in the form A.B.C.D.      |
| <DESTINATION-IP> | Specifies L3 encapsulated IPv4 destination in the form A.B.C.D. |
| id               | Specifies tunnel identifier from the encapsulated packet.       |
| <VRF_NAME>       | Specifies the name of VRF for which the tunnel belongs to.      |

## Examples

Configuring a tunnel parameter to a mirror endpoint:

```
switch(config-mirror-endpoint-test)# source 1.1.1.1 destination 7.7.7.7 id 1 vrf
default type gre
```

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config          | Administrators or local user group members with execution rights for this command. |

## source interface

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

## Description

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

| Parameter   | Description                                                                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <PORT-NUM>  | Specifies a physical port on the switch. Use the format <b>member/slot/port</b> (for example, <b>1/3/1</b> ).                                       |
| <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 per mirror session (4 possible sessions) |
|--------|-----------------------------------------------------------------|
| 6300   | 64                                                              |
| 6400   | 64                                                              |

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



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

## Examples

Configuring a mirrored traffic source interface:

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

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

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

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

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

Removing the first source interface:

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

Configuring a source interface to mirror received packets only:

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

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

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

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

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

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

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

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

## source vlan

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

## Description

Mirroring with VLAN as a source is supported in the following traffic directions:

- **both** - traffic received and transmitted
- **rx** - only received traffic
- **tx** - only transmitted traffic

More than one source VLAN can be configured in a mirror session. Each such VLAN may specify its own direction.

There is a limit of 1024 source VLANs for a given mirror session. There is also a limit of 4096 source VLANs across all mirror sessions.

Same VLAN can be configured as a mirror source for multiple sessions.



---

When changing a source VLAN in an enabled mirror session (i.e. adding, changing direction, or removing) mirrored packets being transmitted out of the mirror destination port from other mirror sources may be briefly interrupted during the reconfiguration.

---

Direction of an existing source VLAN can be updated in one of two ways.

- Reenter the **source vlan <VLAN-NUM> <direction>** command with the new preferred direction.
- Use the **no source vlan <VLAN-NUM> <direction>** form of the command with a direction (**rx** or **tx**) to selectively remove the specified direction.

Specifying the last remaining direction for that VLAN will remove the VLAN from the configuration entirely.

Mirroring allows configuration of VLAN as a source. When VLAN source is configured in the **rx** direction, all packets are mirrored as they are received in the switch. When VLAN source is configured in **tx** direction, all packets are mirrored as they are transmitted out of the switch.

For packets bridged through the switch:

- If the mirror is configured in 'both' direction, two copies of packets are mirrored, otherwise one copy of the packet will be mirrored.

For routed packets:

- If the mirror is configured in **rx** direction, packets are mirrored in the pre-routed form with the Destination MAC address as the switch address.
- If the mirror is configured in **tx** direction, packets are mirrored in post-routed form with the source MAC as the switch address. Destination MAC is the nexthop gateway or station.
- If the mirror is configured in **both** direction, one copy of the packet will be mirrored.

Control plane packets generated by the switch's CPU are processed both in the ingress and the egress packet processing pipeline. The following are the behavior for mirroring with VLAN as source:

- If the mirror is configured in the **rx** or **tx** direction, the packets are mirrored to the mirror destination.
- If the mirror is configured in the **both** direction, two copies of the packets are mirrored to the mirror destination.

The **no** form command will cease mirroring traffic from the specified source VLAN and remove the source from the mirror configuration.

| Parameter | Description                                                                                       |
|-----------|---------------------------------------------------------------------------------------------------|
| VLAN-NUM  | Selects the VLAN number.                                                                          |
| direction | Specifies the direction of mirroring. <b>tx</b> (transmit), <b>rx</b> (receive), or <b>both</b> . |

## Examples

Creating a mirror session and adding a VLAN as a source of traffic in both directions on that port:

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

Creating a mirror session and adding two VLANs as sources of traffic in both directions:

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

Configuring the source in session 2 to receive by specifying the source interface configuration:

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

Removing the first source interface in session 2 entirely, and removing the transmit direction from the other so that mirroring only occurs in the receive direction:

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

Showing maximum of 1024 mirror source VLANs allowed:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config          | Administrators or local user group members with execution rights for this command. |



## Chapter 10

# Monitoring a device using SNMP

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

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

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

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



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

## PoE commands

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

### lldp dot3 poe

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

#### Description

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

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

#### Examples

*On the 6400 Switch Series, interface identification differs.*

Enabling 802.3 TLV list in LLDP:

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

Disabling 802.3 TLV list in LLDP:

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

#### Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

#### Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## lldp med poe

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

### Description

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

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

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

### Examples

*On the 6400 Switch Series, interface identification differs.*

Enabling and disabling LLDP MED PoE:

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

Enabling and disabling LLDP MED PoE priority override:

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

### Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

### Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet

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

## Description

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

## Examples

*On the 6400 Switch Series, interface identification differs.*

Enabling per-interface power distribution:

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

Disabling per-interface power distribution:

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

Showing Quick PoE enabled:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet allocate-by

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

## Description

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

Allocate Power will be based on PD class. Reserve power is based on PD Class. By default, power allocation is by usage.

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

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

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

## Usage

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

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

## Examples

*On the 6400 Switch Series, interface identification differs.*

Configuring the power allocation method:

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

Resetting power allocation method:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet always-on

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

### Description

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

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

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

### Examples

Enabling per-interface power distribution:

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

Disabling per-interface power distribution:

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

### Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

### Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config          | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet assigned-class

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

### Description

Limit PoE power based on the assigned class. When an user assigns a maximum class to an interface, the PSE will limit the maximum power delivered to the PD up to a total power draw not exceeding the PSE assigned-class power. Power demotion occurs when a PD requested class is higher than the PSE assigned class, permitting the PD to receive power and operate in a reduced power mode. PoE ports

cannot set an assigned class when Quick PoE is enabled on the sybsystem. The default assigned class is 4 for 2-pair capable PSE and 6 for 4-pair capable PSE.

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

## Examples

*On the 6400 Switch Series, interface identification differs.*

Setting PoE assigned class:

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

Resetting PoE assigned class to default:

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

Showing Quick PoE enabled:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet power-pairs

```
power-over-ethernet power-pairs {alt-a | alt-a-and-alt-b}
no power-over-ethernet power-pairs {alt-a | alt-a-and-alt-b}
```

### Description

Configures the four-pair capable switch to operate in a mode, that restricts the power delivery for class 0 to class 4 single signature devices to operate only on ALT-A power pair.



When configured, a warning message is displayed. User must accept the warning by entering **Y** to enable the mode.

The **no** form of this command resets the power pairs to default PoE pairs.

| Parameter       | Description                                                                                           |
|-----------------|-------------------------------------------------------------------------------------------------------|
| alt-a           | Delivers power only on the ALT-A pair.                                                                |
| alt-a-and-alt-b | Delivers power on the ALT-A and ALT-B pairs. This is the default configuration on all PoE interfaces. |

## Usage

IEEE 802.3bt devices such as four-pair (class 5 and higher) and dual signature powered devices require power on both pairs. However, there is no such restriction on IEEE 802.3af (class 0 to class 3) and IEEE 802.3at (class 4) powered devices not to draw power on both pairs if the overall consumption does not violate the power class limit. For such powered devices, a **power-pairs** configuration is provided to configure the 4-pair capable switch to restrict power on only one power pair.

## Examples

Configuring PoE power pairs:

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

This setting configures the interface to deliver power only on the ALT-A cable pair when a Class 0-4 device is connected. Devices that require power on all pairs may not operate correctly.

Continue (y/n)? y

Resetting the PoE power pair to default:

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

This setting configures the interface to deliver power on the ALT-A and ALT-B cable pairs. This is the default and most devices work properly with this setting, however some older Class 0-4 devices may not operate correctly.

Continue (y/n)? y

## Command History

| Release | Modification       |
|---------|--------------------|
| 10.09   | Command Introduced |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet pre-std-detect

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

## Description

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

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

## Examples

*On the 6400 Switch Series, interface identification differs.*

Enabling standard device detection:

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

Disabling standard device detection:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet priority

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

## Description

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

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

## Examples

Configuring PoE priority:

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

Resetting the PoE priority to default:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet quick-poe

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

### Description

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

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

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

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

### Examples

*On the 6400 Switch Series, interface identification differs.*

Enabling and disabling quick PoE:

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

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

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet threshold

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

### Description

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

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

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

### Examples

Setting the power-over-ethernet threshold:

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

Resetting the power-over-ethernet threshold to default:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config          | Administrators or local user group members with execution rights for this command. |

## power-over-ethernet trap

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

### Description

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

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

### Examples

Enabling SNMP trap generation for PoE:

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

Disabling SNMP trap generation for PoE:

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context | Authority                                                                          |
|--------------|-----------------|------------------------------------------------------------------------------------|
| 6300<br>6400 | config-if       | Administrators or local user group members with execution rights for this command. |

## show lldp local

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

### Description

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

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

## Examples

*On the 6400 Switch Series, interface identification differs.*

Showing LLDP local device:

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

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

PoE Plus Information

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context             | Authority                                                                                                                                                              |
|--------------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6300<br>6400 | Operator (>) or Manager (#) | Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only. |

## show lldp neighbor

show lldp neighbor [<INTERFACE-ID>]

## Description

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

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

## Examples

*On the 6400 Switch Series, interface identification differs.*

Showing LLDP neighbor information when there is only one neighbor:

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

```

Port : 1/1/10
Neighbor Entries : 1
Neighbor Entries Deleted : 0
Neighbor Entries Dropped : 0
Neighbor Entries Aged-Out : 0
Neighbor Chassis-Name : 84:d4:7e:ce:5d:68
Neighbor Chassis-Description : ArubaOS (MODEL: 325), Version Aruba IAP
Neighbor Chassis-ID : 84:d4:7e:ce:5d:68
Neighbor Management-Address : 169.254.41.250
Chassis Capabilities Available : Bridge, WLAN
Chassis Capabilities Enabled :
Neighbor Port-ID : 84:d4:7e:ce:5d:68
Neighbor Port-Desc : eth0
TTL : 120
Neighbor Port VLAN ID :
Neighbor PoEplus information : DOT3
Neighbor Device Type : TYPE2 PD
Neighbor Power Priority : Unkown
Neighbor Power Source : Primary
Neighbor Power Requested : 25.0 W
Neighbor Power Allocated : 0.0 W
Neighbor Power Supported : No
Neighbor Power Enabled : No
Neighbor Power Class : 5
Neighbor Power Paircontrol : No
Neighbor Power Pairs : SIGNAL

```

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

| Platforms    | Command context             | Authority                                                                                                                                                              |
|--------------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6300<br>6400 | Operator (>) or Manager (#) | Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only. |

## show power-over-ethernet

6300 Switch Series:

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

6400 Switch Series:

```
show power-over-ethernet [<MODULE-ID>] [brief]
```

6300, 6400 Switch Series:

```
show power-over-ethernet [<IFRANGE>] [brief]
```

## Description

Displays the status information of the full system. Displays the brief status of all port or given port if parameter brief is used. Displays the detailed status of given port.

| Parameter   | Description                                              |
|-------------|----------------------------------------------------------|
| <MODULE-ID> | Displays detailed status for the given module.           |
| <IFRANGE>   | Port identifier range.                                   |
| <IFNAME>    | Display the detailed status of given port.               |
| brief       | Display the brief status of all ports or the given port. |

## Examples

Showing sample output for show power-over-ethernet on standalone box with VSF capability:

```
switch# show power-over-ethernet

System Power Status for member 1

Configured Power Status : No redundancy
Operational Power Status : No redundancy
Total Available Power : 740 W
Total Failover Pwr Avl : 0 W
Total Redundancy Power : 0 W
Total Power Drawn : 0 W +/- 6W
Total Power Reserved : 0 W
Total Remaining Power : 740 W
Trap Threshold : 80 %
Trap Enabled : Yes
Always-on PoE Enabled : 1/1
Quick PoE Enabled : None
```

```
Internal Power
PS Total Power
(Watts) Status

1 0 Absent
2 740 Ok
```

```
System Power Status for member 2

Configured Power Status : No redundancy
Operational Power Status : No redundancy
Total Available Power : 600 W
Total Failover Pwr Avl : 0 W
Total Redundancy Power : 0 W
Total Power Drawn : 0 W +/- 6W
Total Power Reserved : 0 W
Total Remaining Power : 600 W
Trap Threshold : 80 %
Trap Enabled : Yes
Always-on PoE Enabled : None
Quick PoE Enabled : None
```

```
Internal Power
PS Total Power
(Watts) Status

1 0 Absent
2 600 Ok
```

Showing sample output for power-over-ethernet member:

```
switch# show power-over-ethernet member 1

System Power Status for member 1

Configured Power Status : No redundancy
Operational Power Status : No redundancy
Total Available Power : 740 W
Total Failover Pwr Avl : 0 W
Total Redundancy Power : 0 W
Total Power Drawn : 0 W +/- 6W
Total Power Reserved : 0 W
Total Remaining Power : 740 W
Trap Threshold : 80 %
Trap Enabled : No
Always-on PoE Enabled : 1/1
Quick PoE Enabled : 1/1

Internal Power
 Total Power
PS (Watts) Status

1 0 Absent
2 740 Ok
```

Showing sample output for power-over-ethernet brief in a VSF stack:

```
switch# show power-over-ethernet brief

Status and Configuration Information for PoE

Member 1 Power Status
 Available: 370 W Reserved: 55.60 W Remaining: 314.40 W
 Always-on PoE Enabled: 1/1
 Quick PoE Enabled: None

PoE Pwr Power Pre-std Alloc PSE Pwr PD Pwr PoE Port PD Cls Type
Port En Priority Detect Act Rsrvd Draw Status Sign --- ---

1/1/1 Yes Low Off Class 0.0 W 0.0 W Denied None 4 2
1/1/2 Yes Critical Off Usage 1.6 W 1.5 W Delivering* Single 0 1
1/1/3 Yes High Off Class 54.0 W 25.5 W Delivering*^ Dual 1/3 3
1/1/4 No Low On Usage 0.0 W 0.0 W Disabled None N/A N/A

Member 2 Power Status
 Available: 600 W Reserved: 0.00 W Remaining: 600 W
 Always-on PoE Enabled: None
 Quick PoE Enabled: None

PoE Pwr Power Pre-std Alloc PSE Pwr PD Pwr PoE Port PD Cls Type
Port En Priority Detect Act Rsrvd Draw Status Sign --- ---

2/1/1 Yes Low Off Class 0.0 W 0.0 W Searching None N/A N/A
2/1/2 Yes Critical Off Usage 0.0 W 0.0 W Searching None N/A N/A
2/1/3 Yes High Off Class 0.0 W 0.0 W Searching None N/A N/A
2/1/4 No Low On Usage 0.0 W 0.0 W Disabled None N/A N/A

*This port may go down in the event of a PSU failure.
^This port is power demoted due to user config or power availability.
```

Showing sample output for power-over-ethernet brief for a Chassis system:

```
switch# show power-over-ethernet brief
```

Status and Configuration Information for PoE

Power Status

Available: 370 W Reserved: 55.60 W Remaining: 314.40 W

Always-on PoE Enabled: 1/1,1/3,1/4,1/7

Quick PoE Enabled: None

| PoE Port | Pwr En | Power Priority | Pre-std Detect | Alloc Act | PSE Pwr Rsrvd | PD Pwr Draw | PoE Port Status | PD Sign | Cls | Type |
|----------|--------|----------------|----------------|-----------|---------------|-------------|-----------------|---------|-----|------|
| 1/1/1    | Yes    | Low            | Off            | Class     | 0.0 W         | 0.0 W       | Denied          | None    | 4   | 2    |
| 1/1/2    | Yes    | Critical       | Off            | Usage     | 1.6 W         | 1.5 W       | Delivering*     | Single  | 0   | 1    |
| 1/1/3    | Yes    | High           | Off            | Class     | 54.0 W        | 25.5 W      | Delivering^     | Dual    | 1/3 | 3    |
| 1/1/4    | No     | Low            | On             | Usage     | 0.0 W         | 0.0 W       | Disabled        | None    | N/A | N/A  |

\*This port may go down in the event of a PSU failure.

^This port is power demoted due to user config or power availability.

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

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

Status and Configuration Information for port 1/1/1

Member 1Power Status

Available: 370 W Reserved: 55.60 W Remaining: 314.40 W

Always-on PoE Enabled: 1/1

| PoE Port | Pwr En | Power Priority | Pre-std Detect | Alloc Act | PSE Pwr Rsrvd | PD Pwr Draw | PoE Port Status | PD Sign | Cls | Type |
|----------|--------|----------------|----------------|-----------|---------------|-------------|-----------------|---------|-----|------|
| 1/1/1    | Yes    | Low            | Off            | Class     | 0.0 W         | 0.0 W       | Denied          | None    | 4   | 2    |

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

For 6300 Switch series:

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

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

Member 1Power Status

Available: 370 W Reserved: 55.60 W Remaining: 314.40 W

Always-on PoE Enabled: 1/1

Quick PoE Enabled: None

| PoE Port | Pwr En | Power Priority | Pre-std Detect | Alloc Act | PSE Pwr Rsrvd | PD Pwr Draw | PoE Port Status | PD Sign | Cls | Type |
|----------|--------|----------------|----------------|-----------|---------------|-------------|-----------------|---------|-----|------|
| 1/1/1    | Yes    | Low            | Off            | Class     | 0.0 W         | 0.0 W       | Denied          | None    | 4   | 2    |
| 1/1/2    | Yes    | Critical       | Off            | Usage     | 1.6 W         | 1.5 W       | Delivering*     | Single  | 0   | 1    |

For 6400 Switch series:

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

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

```

Power Status
Available: 360 W Reserved: 0.00 W Remaining: 360.00 W
Always-on PoE Enabled: 1/1
Quick PoE Enabled: None

```

| PoE Port | Pwr En | Power Priority | Pre-std Detect | Alloc Act | PSE Pwr Rsrvd | PD Pwr Draw | PoE Port Status | PD Sign | Cls | Type |
|----------|--------|----------------|----------------|-----------|---------------|-------------|-----------------|---------|-----|------|
| 1/1/1    | Yes    | Low            | Off            | Usage     | 0.0 W         | 0.0 W       | Searching       | N/A     | N/A | N/A  |
| 1/1/2    | Yes    | Low            | Off            | Usage     | 0.6 W         | 0.0 W       | Searching       | N/A     | N/A | N/A  |

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

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

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

Showing sample output for power-over-ethernet brief for a missing member:

```
switch# show power-over-ethernet member 3 brief
```

```
Member 3 is not physically present.
```

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

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

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

Showing power-over-ethernet port with dual signature PD connected:

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

```
Status and Configuration Information for port 1/1/1*
```

|                       |              |                        |              |
|-----------------------|--------------|------------------------|--------------|
| Power Enable          | : Yes        | PD signature           | : Dual       |
| PoE PairA Status      | : Delivering | PoE PairB Status       | : Delivering |
| Alloc-by Configured   | : Class      | Alloc-by Actual        | : Class      |
| User Profile Priority | : High       | Port Config Priority   | : Low        |
| Port Priority         | : High       | Pre-std Detect         | : Disabled   |
| PD Type               | : Type3      | User Assigned Class    | : Class6     |
| PairA Requested Class | : Class1     | PairB Requested Class  | : Class4     |
| PairA Assigned Class  | : Class1     | PairB Assigned Class   | : Class4     |
| Fault Status PairA    | : None       | Fault Status PairB     | : None       |
| PD Class Override     | : Disabled   | Power Pairs Configured | : alt-a      |
|                       |              | Power Pairs Applied    | : alt-a-and- |

```
alt-b
```

```
PoE Counter Information
```

|                        |     |                      |     |
|------------------------|-----|----------------------|-----|
| Over Current Cnt PairA | : 0 | MPS Absent Cnt PairA | : 0 |
| Power Denied Cnt PairA | : 0 | Short Cnt PairA      | : 0 |
| Over Current Cnt PairB | : 0 | MPS Absent Cnt PairB | : 0 |
| Power Denied Cnt PairB | : 0 | Short Cnt PairB      | : 0 |

#### Power Information

|                       |          |                    |          |
|-----------------------|----------|--------------------|----------|
| PSE Voltage           | : 56.3 V | PSE Reserved power | : 34.0 W |
| PD Current Draw       | : 4.1 A  | PD Power Draw      | : 24.6 W |
| PD Average Power Draw | : 24.0 W | PD Peak Power Draw | : 25.1 W |

#### LLDP Information

|                                    |                 |
|------------------------------------|-----------------|
| MED Override                       | : Enabled       |
| MED Priority                       | : High          |
| PSE TLV Configured                 | : dot3, med     |
| PSE TLV Sent Type                  | : dot3-ext      |
| PD TLV Sent Type                   | : med, dot3-ext |
| DS PSE Allocated Power Value Alt A | : 2.5 W         |
| DS PD Requested Power Value Mode A | : 2.5 W         |
| DS PSE Allocated Power Value Alt B | : 25.0 W        |
| DS PD Requested Power Value Mode B | : 25.0 W        |

Showing power-over-ethernet port with single signature PD connected:

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

#### Status and Configuration Information for port 1/1/9\*

|                       |              |                         |                   |
|-----------------------|--------------|-------------------------|-------------------|
| Power Enable          | : Yes        | PD signature            | : None            |
| PoE Port Status       | : Delivering | PD Type                 | : Type3           |
| Alloc-by Configured   | : Usage      | Alloc-by Actual         | : Usage           |
| User Profile Priority | : High       | Port Config Priority    | : Low             |
| Port Priority         | : High       | Pre-std Detect          | : Disabled        |
| PD Requested Class    | : Class1     | PSE Assigned Class      | : Class1          |
| Fault Status          | : None       | User set Assigned Class | : Class6          |
| PD Class Override     | : Disabled   | Power Pairs Configured  | : alt-a-and-alt-b |
|                       |              | Power Pairs Applied     | : alt-a-and-alt-b |

#### PoE Counter Information

|                  |     |                |     |
|------------------|-----|----------------|-----|
| Over Current Cnt | : 0 | MPS Absent Cnt | : 0 |
| Power Denied Cnt | : 0 | Short Cnt      | : 0 |

#### Power Information

|                       |          |                    |         |
|-----------------------|----------|--------------------|---------|
| PSE Voltage           | : 56.3 V | PSE Reserved power | : 8.6 W |
| PD Current Draw       | : 1.1 A  | PD Power Draw      | : 8.6 W |
| PD Average Power Draw | : 8.0 W  | PD Peak Power Draw | : 9.1 W |

#### LLDP Information

|                           |            |
|---------------------------|------------|
| LLDP Detect               | : Disabled |
| PSE TLV Configured        | : N/A      |
| PSE TLV Sent Type         | : N/A      |
| PD TLV Sent Type          | : N/A      |
| PSE Allocated Power Value | : 0.0 W    |
| PD Requested Power Value  | : 0.0 W    |

Showing power-over-ethernet for a port range:

switch# **show power-over-ethernet 1/1/3-1/1/4**

Status and Configuration Information for port 1/1/3

|                       |              |                         |                   |
|-----------------------|--------------|-------------------------|-------------------|
| Power Enable          | : Yes        | PD signature            | : None            |
| PoE Port Status       | : Delivering | PD Type                 | : Type3           |
| Alloc-by Config       | : Usage      | Alloc-by Actual         | : Usage           |
| User Profile Priority | : High       | Port Config Priority    | : Low             |
| Port Priority         | : High       | Pre-std Detect          | : Disabled        |
| PD Requested Class    | : Class1     | PSE Assigned Class      | : Class1          |
| Fault Status          | : None       | User set Assigned Class | : Class6          |
| PD Class Override     | : Disabled   | Power Pairs Configured  | : alt-a-and-alt-b |
|                       |              | Power Pairs Applied     | : alt-a-and-alt-b |

PoE Counter Information

|                  |     |                |     |
|------------------|-----|----------------|-----|
| Over Current Cnt | : 0 | MPS Absent Cnt | : 0 |
| Power Denied Cnt | : 0 | Short Cnt      | : 0 |

Power Information

|                       |          |                    |         |
|-----------------------|----------|--------------------|---------|
| PSE Voltage           | : 56.3 V | PSE Reserved power | : 8.6 W |
| PD Current Draw       | : 1.1 A  | PD Power Draw      | : 8.6 W |
| PD Average Power Draw | : 8.0 W  | PD Peak Power Draw | : 9.1 W |

LLDP Information

|                           |            |
|---------------------------|------------|
| LLDP Detect               | : Disabled |
| PSE TLV Configured        | : N/A      |
| PSE TLV Sent Type         | : N/A      |
| PD TLV Sent Type          | : N/A      |
| PSE Allocated Power Value | : 0.0 W    |
| PD Requested Power Value  | : 0.0 W    |

Status and Configuration Information for port 1/1/4\*

|                       |              |                         |            |
|-----------------------|--------------|-------------------------|------------|
| Power Enable          | : Yes        | PD signature            | : None     |
| PoE Port Status       | : Delivering | PD Type                 | : Type3    |
| Alloc-by Config       | : Usage      | Alloc-by Actual         | : Usage    |
| User Profile Priority | : High       | Port Config Priority    | : Low      |
| Port Priority         | : High       | Pre-std Detect          | : Disabled |
| PD Requested Class    | : Class1     | PSE Assigned Class      | : Class1   |
| Fault Status          | : None       | User set Assigned Class | : Class6   |
| PD Class Override     | : Disabled   | Power Pairs Configured  | : alt-a    |
|                       |              | Power Pairs Applied     | : alt-a    |

PoE Counter Information

|                  |     |                |     |
|------------------|-----|----------------|-----|
| Over Current Cnt | : 0 | MPS Absent Cnt | : 0 |
| Power Denied Cnt | : 0 | Short Cnt      | : 0 |

Power Information

|                       |          |                    |         |
|-----------------------|----------|--------------------|---------|
| PSE Voltage           | : 56.3 V | PSE Reserved power | : 4.3 W |
| PD Current Draw       | : 1.1 A  | PD Power Draw      | : 4.3 W |
| PD Average Power Draw | : 4.0 W  | PD Peak Power Draw | : 4.3 W |

LLDP Information

```
LLDP Detect : Disabled
PSE TLV Configured : N/A
PSE TLV Sent Type : N/A
PD TLV Sent Type : N/A
PSE Allocated Power Value : 0.0 W
PD Requested Power Value : 0.0 W
```

## Command History

| Release          | Modification                                                                                   |
|------------------|------------------------------------------------------------------------------------------------|
| 10.09            | Added power-pairs configuration in the <b>show power-over-ethernet &lt;IFRANGE&gt;</b> output. |
| 10.07 or earlier | --                                                                                             |

## Command Information

| Platforms    | Command context             | Authority                                                                                                                                                              |
|--------------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6300<br>6400 | Operator (>) or Manager (#) | Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only. |

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

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

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

## SNMP support and AirWave

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

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

### SNMP on the switch

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

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

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

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

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

## Supported features with AirWave and the AOS-CX switch

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

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

## Configuring the AOS-CX switch to be monitored by AirWave

### Prerequisites

Aruba AirWave is active on the network.

### Procedure

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

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

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

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

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

```
switch(config)# snmpv3 user Admin auth sha auth-pass ciphertext
AQBapZHf2d20GYr/xcGUzYzm0zjNf/4VKHtSqBNImqtFYbJYCgAAALkGFJVcSp3nZ3o=
priv des priv-pass ciphertext
AQBapb0H2poBQKXPoVsC9L9qzZyfJQnzR7hmTr7LGsOsI7K3CgAAAKP98Rq2jfTrFwQ=

switch(config)# snmpv3 context Admin
```

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

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

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

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

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

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

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

## AirWave commands

### logging

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

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

### Description

Enables syslog forwarding to a remote syslog server.

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

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

## Examples

Enabling the syslog forwarding to remote syslog server 10.0.10.2:

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

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

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

Disabling syslog forwarding to a remote syslog server:

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

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

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

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

## snmp-server community

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

### Description

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

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

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

## Examples

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

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

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

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

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

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

## snmp-server host

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

## Description

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

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



Configuring **snmpv3 informs** is not supported.

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

| Parameter          | Description                                                                                                                                                                      |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| trap version v3    | Specifies the trap notification type for SNMPv3.                                                                                                                                 |
| user <NAME>        | Specifies the SNMPv3 user name to be used in the SNMP trap notifications.                                                                                                        |
| community <STRING> | Specifies the name of the community string to use when sending trap notifications. Range: 1 - 32 printable ASCII characters, excluding space and question mark. Default: public. |
| <UDP-PORT>         | Specifies the UDP port on which notifications are sent. Range: 1 - 65535. Default: 162.                                                                                          |
| vrf <VRF-NAME>     | Specifies the name of the VRF on which to send the notifications.                                                                                                                |

## Examples

```

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

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

```

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

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

## snmp-server vrf

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

### Description

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

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

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

### Example

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

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

### Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

### Command Information

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

## snmpv3 context

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

### Description

Creates an SNMPv3 context on the specified VRF.

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

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

## Examples

Creating an SNMPv3 context named **newContext**:

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

Creating an SNMPv3 context named **newContext** on VRF **myVrf** and with community string **private**.

```
switch(config)# snmpv3 context newContext vrf myVrf community private
```

Removing the SNMPv3 context named **newContext** on VRF **myVrf**:

```
switch(config)# no snmpv3 context newContext vrf myVrf
```

## 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. |

## snmpv3 user

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

## Description

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

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

| Parameter                                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>&lt;NAME&gt;</code>                                          | Specifies the SNMPv3 username. Range 1 - 32 printable ASCII characters, excluding space and question mark.                                                                                                                                                                                                                                                                                                                          |
| <code>auth &lt;AUTH-PROTOCOL&gt;</code>                            | Specifies the authentication protocol used to validate user logins. Available options are: <code>md5</code> or <code>sha</code> .                                                                                                                                                                                                                                                                                                   |
| <code>auth-pass {plaintext   ciphertext} &lt;AUTH-PWORD&gt;</code> | Specifies the SNMPv3 user password. Range for <code>plaintext</code> is 8 - 32 printable ASCII characters, excluding space and question mark.<br>Range for <code>ciphertext</code> is 1 - 120 printable ASCII characters. This option is only used when copying user configuration settings between switches. It enables you to duplicate a user's configuration on another switch without having to know their password.           |
| <code>priv &lt;PRIV-PROTOCOL&gt;</code>                            | Specifies the SNMPv3 security protocol (encryption method). Available options are: <code>aes</code> or <code>des</code> .                                                                                                                                                                                                                                                                                                           |
| <code>priv-pass {plaintext   ciphertext} &lt;PRIV-PWORD&gt;</code> | Specifies the SNMPv3 user privacy passphrase. Range for <code>plaintext</code> is 8 - 32 printable ASCII characters, excluding space and question mark.<br>Range for <code>ciphertext</code> is 1 - 120 printable ASCII characters. This option is only used when copying user configuration settings between switches. It enables you to duplicate a user's configuration on another switch without having to know their password. |

## Examples

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

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

Removing an SNMPv3 user named `Admin`:

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

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

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

Copying an SNMP user from switch 1 to switch 2.

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

```
switch1(config)# snmpv3 user Admin auth sha auth-pass plaintext mypassword
priv des priv-pass plaintext myprivpass
switch1(config)# exit
switch1# show running-config
Current configuration:
!
!Version AOS-CX TL.10.00.0003-8017-gdeb0606~dirty
!
!
!
snmpv3 user Admin auth sha auth-pass ciphertext
AQBapZHf2d20GYr/xcGUzYzm0zjNf/4VKHtSqbNImqtFYbJYCgAAALkGFJVcSp3nZ3o=
priv des priv-pass ciphertext
AQBapb0H2poBQKXPoVsC9L9qzZyfJQnzR7hmTr7LGsOsI7K3CgAAAKP98Rq2jfTrFwQ=
ssh server vrf mgmt
!
!
!
!
interface mgmt
 no shutdown
 ip dhcp
 vlan 1
```

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

```
switch1(config)# snmpv3 user Admin auth sha auth-pass
ciphertextAQBapZHf2d20GYr/xcGUzYzm0zjNf/4VKHtSqbNImqtFYbJYCgAAALkGFJVcSp3nZ3o=priv
des priv-pass ciphertext
AQBapb0H2poBQKXPoVsC9L9qzZyfJQnzR7hmTr7LGsOsI7K3CgAAAKP98Rq2jfTrFwQ=
```

## Command History

| Release          | Modification |
|------------------|--------------|
| 10.07 or earlier | --           |

## Command Information

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

### Accessing HPE Aruba Networking Support

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

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                           | <a href="https://community.arubanetworks.com/">https://community.arubanetworks.com/</a>                                                                                                 |
| HPE Aruba Networking Hardware Documentation and Translations Portal | <a href="https://www.arubanetworks.com/techdocs/hardware/DocumentationPortal/Content/home.htm">https://www.arubanetworks.com/techdocs/hardware/DocumentationPortal/Content/home.htm</a> |

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

## Accessing Updates

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

Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

To subscribe to eNewsletters and alerts:

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

## Warranty Information

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

## Regulatory Information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at <https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

### Additional regulatory information

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

## Documentation Feedback

HPE Aruba Networking is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback ([docsfeedback-switching@hpe.com](mailto: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.