

AOS-CX 10.17 Monitoring Guide

6300, 6400 Switch Series



**Hewlett Packard
Enterprise**

Published: November 2025

Version: 1

Copyright Information

© Copyright 2025 Hewlett Packard Enterprise Development LP.

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

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



Notices

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

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

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

Acknowledgment

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

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

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

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

UNIX® is a registered trademark of The Open Group.

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

Contents

About this document	9
Applicable products	9
Latest version available online	9
Command syntax notation conventions	9
About the examples	10
Identifying switch ports and interfaces	10
Identifying modular switch components	11
Monitoring hardware through visual observation	12
Confirming normal operation of the switch by reading LEDs	12
Detecting if the switch is not ready for a failover event	13
Finding faulted components using the switch LEDs	13
IP Flow Information Export	15
Flow monitors	19
Flow Records	19
Flow Exporters	19
Destinations	20
Configuring IP Flow Information Export on 6300 and 6400 Switches	20
Role-based IPFIX	24
FAQs and Troubleshooting	25
Flow monitoring commands	25
alert-threshold aggregate active-flows percentage	26
alert-threshold endpoint active-flows	26
alert-threshold endpoint new-connections accepted	27
alert-threshold endpoint new-connections initiated	28
collect application tcp establishment-time	28
diag-dump ipfix basic	29
description	30
exporter	32
flow exporter	33
flow monitor	37
flow record	39
flow-tracking	43
ip ipv6 flow monitor (interface)	46
ipv4 ipv6 flow monitor (role)	47
show flow exporter	48
show flow monitor	50
show flow record	52
show flow-tracking	55
show running-config	56
show ip flow-table-utilization	58
show IPv4 flow endpoint statistics	59
show IPv6 flow session statistics	60
show tech ipfix	61
Queue Monitoring	63
Queue statistics history	63
Queue depth	63

Data retention limits	63
Using collected data	64
QoS commands	64
clear queue-monitor interface	64
queue-monitor	66
show interface queue-monitor	67
show interface queue-monitor status	75
show queue-monitor status	78
IP Flow Path Trace	79
Limitations, Conflicts or Exclusions	79
IP Flow Path Trace commands	81
IP Flow Path Trace	81
gRPC network management interface	85
gNMI prerequisites and setup	85
gNMI capabilities, limitations, and best practices	86
Getting started with gNMI	87
gNMI connection types and common attributes	88
Supported OpenConfig models	90
gNMI subscription modes	91
gNMI secure connectivity	92
gNMI use case 1: interface monitoring	92
gNMI use case 2: system monitoring	93
gNMI use case 3: platform and hardware monitoring	93
gNMI troubleshooting	94
gNMI diagnostics and monitoring	96
gNMI command syntax templates	98
gNMI implementation examples	98
gNMI commands	108
crypto pki application gnmi certificate	108
gnmi vrf	109
show gnmi	109
Boot commands	111
boot fabric-module	111
boot line-module	112
boot management-module	113
boot management-module (recovery console)	114
boot set-default	115
boot system	116
show boot-history	118
External storage	123
External storage commands	123
address	123
directory	124
disable	125
enable	125
external-storage	126
password (external-storage)	127
show external-storage	128
show running-config external-storage	128
type	129
username	130
vrf	131

IP-SLA	132
IP-SLA guidelines	132
Limitations with VoIP SLAs	133
IP-SLA commands	133
http	133
https	135
icmp-echo	137
ip-sla	138
ip-sla responder	139
show ip-sla all	141
show ip-sla responder	143
show ip-sla	144
start-test	146
stop-test	146
tcp-connect	147
udp-echo	148
udp-jitter-voip	150
vrf	152
L1-100Mbps downshift	153
Limitations with speed downshift	153
L1-100Mbps downshift commands	153
downshift enable	153
show interface	154
show interface statistics	161
show interface downshift-enable	163
show running-config interface	164
Mirroring	167
Mirror statistics	168
Classifier policies and mirroring sessions	168
Mirroring on a VSX configuration	168
VLAN as a source	169
Mirroring commands	169
clear mirror	169
clear mirror endpoint	170
comment	171
copy tcpdump-pcap	172
copy tshark-pcap	173
destination cpu	174
destination interface	174
destination tunnel	176
diagnostic	177
diag utilities tcpdump	178
disable	180
enable	181
mirror session	182
mirror endpoint	182
show mirror	184
show mirror endpoint	186
shutdown	187
source	188
source interface	189
source vlan	191

Monitoring a device using SNMP	194
Packet Capture	194
packet-capture commands	194
copy packet-capture	194
packet-capture delete-pcap	195
packet-capture enable	196
packet-capture	197
show packet-capture pcaps	199
show packet-capture session	200
Power-over-Ethernet	202
PoE commands	203
lldp dot3 poe	203
lldp med poe	204
power-over-ethernet	204
power-over-ethernet allocate-by	205
power-over-ethernet always-on	206
power-over-ethernet assigned-class	207
power-over-ethernet power-pairs	208
power-over-ethernet pre-std-detect	209
power-over-ethernet priority	210
power-over-ethernet quick-poe	211
power-over-ethernet threshold	212
power-over-ethernet trap	213
show lldp local	213
show lldp neighbor	214
show power-over-ethernet	215
Aruba AirWave	223
SNMP support and AirWave	223
SNMP on the switch	223
Supported features with AirWave and the AOS-CX switch	224
Configuring the AOS-CX switch to be monitored by AirWave	224
AirWave commands	225
logging	225
snmp-server community	227
snmp-server host	228
snmp-server vrf	230
snmpv3 context	230
snmpv3 user	231
Support and Other Resources	234
Accessing HPE Aruba Networking Support	234
Accessing Updates	235
Warranty Information	235
Regulatory Information	235
Documentation Feedback	235

Chapter 1

About this document

This document describes features of the AOS-CX network operating system. It is intended for administrators responsible for installing, configuring, and managing HPE Aruba Networking 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, S3L75A, S3L76A, S3L77A, S4P41A, S4P42A, S4P43A, S4P44A, S4P45A, S4P46A, S4P47A, S4P48A)
- 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> ■ example-text	Identifies a placeholder—such as a parameter or a variable—that you must substitute with an actual value in a command or in code: ■ For output formats where italic text cannot be displayed, variables are enclosed in angle brackets (<code>< ></code>). 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

Convention	Usage
	choose only one. Any spaces that are on either side of the vertical bar are included for readability and are not a required part of the command syntax.
{ }	Braces. Indicates that at least one of the enclosed items is required.
[]	Brackets. Indicates that the enclosed item or items are optional.
... 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 HPE Aruba Networking 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 HPE Aruba Networking 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.

Chapter 3

IP Flow Information Export

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 internal or external 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 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 the Traffic Insight features, refer to the *AOS-CX Security Guide*.

Compatibility with Application Recognition

If the **application recognition** feature is 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*.

Information Elements

The IPFIX Information Elements (IE) are entities that are defined and maintained by the Internet Assigned Numbers Authority (IANA). They are characterized by a unique piece of information they can provide about a flow. Information Elements may be either private or public. Private Information Elements are exported with a Private Enterprise Number (PEN).

AOS-CX can act as an intermediate collecting process for flow reports from hardware to append certain additional IPFIX information elements to the flow reports. When configured, the software will act as an intermediate exporting process to export the augmented flow reports to any configured flow exporters. AOS-CX supports the standard and private information elements shown in the tables below.



Not all switches support all information elements. View a list of information elements supported by your switch using the command **show flow monitor <monitor-name> information-elements**. All Standard Information Element registration information can be found on the IANA website, at <https://www.iana.org/assignments/ipfix/>.

Standard Information Elements

octetDeltaCount

Private Information Elements

tlsClientVersion (ID 1029, 4823 Aruba)

Standard Information Elements

packetDeltaCount
protocolIdentifier
sourceTransportPort
sourceIPv4Address
ingressInterface
destinationTransportPort
destinationIPv4Address
sourceIPv6Address
destinationIPv6Address
sourceIPv6Address
destinationIPv6Address
classId
sourceMacAddress
vlanId
ipVersion
destinationMacAddress
applicationDescription
applicationId
applicationName
flowEndReason
flowStartMicroseconds
flowEndMicroseconds
dnsResponseCode
ingressPhysicalInterface
egressPhysicalInterface
applicationCategoryName

Private Information Elements

tlsServerVersion (ID 1030, 4823 Aruba)
ja3 (ID 1031, 4823 Aruba)
ja3s (ID 1032, 4823 Aruba)
supportedNextProtocol (ID 1033, 4823 Aruba)
CertificateIssuer (ID 1034, 4823 Aruba)
CertificateIssueDate (ID 1035, 4823 Aruba)
CertificateExpiryDate (ID 1036, 4823 Aruba)
applicationType (ID 1100, 4823 Aruba)
tcp3WayHsServerRespTime (ID 1101, 4823 Aruba)
tcp3WayHsClientRespTime (ID 1102, 4823 Aruba)
dropGroup1 (ID 1200, 14823 Aruba)
dropReason1 (ID 1201, 14823 Aruba)
dropStartMilliseconds1 (ID 1202, 14823 Aruba)
dropEndMilliseconds1 (ID 1203, 14823 Aruba)
dropGroup2 (ID 1210, 14823 Aruba)
dropReason2 (ID 1211, 14823 Aruba)
dropStartMilliseconds2 (ID 1212, 14823 Aruba)
dropEndMilliseconds2 (ID 1213, 14823 Aruba)
dropGroup3 (ID 1220, 14823 Aruba)
dropReason3 (ID 1221, 14823 Aruba)
dropStartMilliseconds3 (ID 1222, 14823 Aruba)
dropEndMilliseconds3 (ID 1223, 14823 Aruba)
dropGroup4 (ID 1230, 14823 Aruba)
dropReason4 (ID 1231, 14823 Aruba)
dropStartMilliseconds4 (ID 1232, 14823 Aruba)
dropEndMilliseconds4 (ID 1233, 14823 Aruba)
cngQueueThreshold (ID 490, 14823 Aruba)
cngQueueObservedTimeStamp (ID 491, 14823 Aruba)

About individual Information Elements

The following IEs are used for both IPv4 and IPv6 flows:

- cngSourceIP
- cngDestinationIP

They are always 16 bytes wide. The IP version of the flow is indicated in the **ipVersion** IE. For IPv4 addresses, the first 12 octets are zero and the IPv4 address is encoded in the last 4 octets. For IPv6 addresses, the full 16 bytes encode the address.

The following IEs reflect the hardware interface ID used by the associated flow:

- ingressPhysicalInterface
- egressPhysicalInterface

The hardware interface ID does not match the front-panel port number of the interface. The hardware interface ID can be mapped to the front-panel port number by using the REST API to query for **hw_intf_info** for all interfaces and saving the **switch_intf_id** value together with the associated front-panel name to form a mapping table.

Example using curl and jq:

```
bash
$ curl -X GET
  "{SWITCH}/rest/{VERSION}/system/interfaces?attributes=hw_
_intf,info,name&depth=2"
-H "x-csrf-token: {CSRF_TOKEN}"
-b {AUTH_COOKIE_FILE}
| jq 'to_entries
    | map({port: .key, switch_intf_id: (.value.hw_intf_info.switch_intf_id |
tonumber)})
    | sort_by(.switch_intf_id)
    | map({(.port): .switch_intf_id})
    | add'
{
  "1/1/6": 1,
  "1/1/2": 2,
  "1/1/1": 3,
  ...
}
```

The **IPFIX Drop Exceptions** feature is a network debugging and triaging tool supported on 6300 Switch series. It is configured using the **collect drop ingress-exceptions** command in the **config-flow-record** context. The private information elements 1200-1233 contain possible drop reason information for why dropped packets are seen for a flow. Up to four drop reasons are supported for each flow. Data in consecutive values (for example, 1200, 1201, 1202, 1203) indicate the data in those fields are associated with each other.

For each drop reason, AOS-CX reports the drop group name, drop reason name, and drop start and end timestamps in milliseconds.

Table 1: Drop reasons for the IPFIX Drop Exceptions feature

Drop reason	Example Cause	Supported Platforms
Unknown VLAN	The packet has an unknown VLAN	5420, 6200, 6300, 6400, 8100, 8325,

Drop reason	Example Cause	Supported Platforms
		8325P, 8325H, 8360, 10000 Switch series
MTU Failure	The packet size is larger than the IP Maximum Transmission Unit (MTU) configured in the interface.	5420, 6200, 6300, 6400, 8100, 8325, 8325P, 8325H, 8360 Switch series
Unroutable Unicast	The unicast packet can't be routed as destination address is not resolved.	5420, 6200, 6300, 6400, 8100, 8360 Switch series.
Unicast TTL expired	For IPv4: The unicast packet has a Time-to-Live (TTL) value of 0 or 1. For IPv6: The unicast packet has a hop limit value of 0 or 1.	5420, 6200, 6300, 6400, 8100, 8360 Switch series.
Security Ingress IP Lockdown	Packet's source IP address, VLAN, MAC or Interface does not match the IP binding entry.	5420, 6200, 6300, 6400, 8100, 8360 Switch series.
Security Ingress MAC Lockout	The packet is sourced from a locked-out MAC address.	5420, 6200, 6300, 6400, 8100, 8360 Switch series.
IPTCAM invalid IP address	Packet has invalid source or destination IP address. (For example, an invalid IPv4 address of 0.0.0.0 or 127.0.0.0 , or 240.0.0.0 , or an invalid IPv6 address of :: or ::1 .)	5420, 6200, 6300, 6400, 8100, 8360 Switch series.
Security learn	Packet is received on a security-enabled interface and is learned for the first time.	5420, 6200, 6300, 6400, 8100, 8360 Switch series.
Security drop	Packet is received on a security-enabled interface but is not allowed, as it is not authorized on the VLAN.	5420, 6200, 6300, 6400, 8100, 8360 Switch series.
Meter drop	The packet is dropped because of exceeding the rate limit configured on port or via a port-access policy.	5420, 6200, 6300, 6400, 8100, 8360 Switch series.
ASIC drop	The packet is dropped by ASIC because of parity error, parser error or invalid lookup entry.	5420, 6200, 6300, 6400, 8100, 8360 Switch series

For 6300 Switch series:

If there is no data in the **drop group name** or **drop reason name** fields, there have been less than four drop reasons for the flow. Every time the drop reason changes, the next set of drop reason IEs will be populated. Every time the cache inactive timeout value has passed between drops for the same drop reason, the next set of drop reason IEs will be populated. If more than four drop reasons are seen for a given flow, data for the oldest drop reason will be purged, and data for the new drop reason will be added in its place. Analyzing the timestamps for all drop reasons will indicate which reason is the newest.

For Switch Series:

If the **drop end** timestamp has the same value as the **drop start** timestamp, that means only one packet of the flow was dropped. If the drop end timestamp is different from the drop start timestamp, that means more than one packet was dropped within 30 seconds.

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. On the 5420, 6200, 6300, 6400, 8100, and 8360 switch series 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:

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

- IPv6 version
- IPv6 source address
- 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 is one possible destination for a flow exporter within the 6300 and 6400 switch series:

1. Hostname or IP address of a device with an optional VRF

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 and 6400 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)



IPv6 related commands are only applicable to switches that support IPv6 protocol.

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 ip protocol
switch(config-flow-record)# match ip source address
switch(config-flow-record)# match ip destination address
switch(config-flow-record)# match ip 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)# collect application https url

switch(config)# flow record flowRecordv6
switch(config-flow-record)# match ipv6 protocol
switch(config-flow-record)# match ipv6 source address
switch(config-flow-record)# match ipv6 destination address
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
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
switch(config-flow-record)# collect application https url"
switch(config-flow-record)# collect application type
switch(config-flow-record)# collect egress interface
switch(config-flow-record)# collect egress queue
switch(config-flow-record)# collect egress vlan
switch(config-flow-record)# collect forwarding-status

```

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

```

-----
Flow record 'flowRecordv4'
-----
Description           : ipv4
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
  application https url
  timestamp absolute first
  timestamp absolute last
-----

Flow record 'flowRecordv6'
-----
Description           : ipv6
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
  application https url
  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 appFlow type application-flows
switch(config-ti-TI)# monitor dns type application-flows
switch(config-ti-TI)# monitor raw-mon type raw-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 running-config traffic-insight
traffic-insight TI
enable
source ipfix
!
monitor topN type topN-flows entries 5
monitor appFlow type application-flows
monitor raw type raw-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
```

A collecting process is not configured with IPFIX.

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, 5420, 6200, 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 on the 6200, 6300, and 6400 platforms. The mutually exclusive features are:

- UBT
- MAC Lockout
- Extended Router MAC

- Source Port Filter
- IP Lockdown
- L3VNI

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.
- Starting with 10.17, role-based IPFIX supports monitoring in both directions. If only ingress is enabled, it monitors traffic in the ingress direction. Port-based IPFIX continues to support only ingress flow monitoring.

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

alert-threshold aggregate active-flows percentage

[no] alert-threshold aggregate active-flows percentage <1-100>

Description

Sets the alert threshold as a percentage of maximum IP flow table entries per linecard or stack member. An alert is triggered when the number of active flows reaches or exceeds this percentage. The **show ip flow-table-utilization** command displays the IP flow table size on each LC or stack member.

Parameter	Description
<1-100>	Percentage of the maximum allowed aggregate active flows for alert.

Examples

Configuring the aggregate active flows alert threshold per linecard or stack member:

```
switch(config-flow-tracking)# alert-threshold aggregate active-flows percentage 50
```

Removing the aggregate active flows alert threshold:

```
switch(config-flow-tracking)# no alert-threshold aggregate active-flows percentage
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

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

alert-threshold endpoint active-flows

[no] alert-threshold endpoint active-flows <50-1000>

Description

Sets the alert threshold for the number of active flows per endpoint. An alert is triggered when an endpoint's active flow count reaches or exceeds this threshold.

Parameter	Description
<50-1000>	Endpoint session count alert threshold.

Examples

Configuring the endpoint active flows alert threshold:

```
switch(config-flow-tracking)# alert-threshold endpoint active-flows 50
```

Removing the endpoint active flows alert threshold:

```
switch(config-flow-tracking)#no alert-threshold endpoint active-flows
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

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

alert-threshold endpoint new-connections accepted

[no] alert-threshold endpoint new-connections accepted <10-500>

Description

Sets the alert threshold for the number of new connections accepted by an endpoint within a 30-second monitoring window. An alert is triggered when an endpoint's accepted connection count reaches or exceeds this threshold.

Parameter	Description
<10-500>	Endpoint new connections accepted alert threshold.

Examples

Configuring the endpoint new connections accepted alert threshold:

```
switch(config-flow-tracking)# alert-threshold endpoint new-connections accepted 50
```

Removing the endpoint new connections accepted alert threshold:

```
switch(config-flow-tracking)#no alert-threshold endpoint new-connections accepted
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

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

alert-threshold endpoint new-connections initiated

[no] alert-threshold endpoint new-connections initiated <10-500>

Description

Sets the alert threshold for the number of new connections initiated by an endpoint within a 30-second monitoring window. An alert is triggered when an endpoint's initiated connection count reaches or exceeds this threshold.

Parameter	Description
<10-500>	Endpoint new connections initiated alert threshold.

Examples

Configuring the endpoint new connections initiated alert threshold:

```
switch(config-flow-tracking)# alert-threshold endpoint new-connections initiated 50
```

Removing the endpoint new connections initiated alert threshold:

```
switch(config-flow-tracking)#no alert-threshold endpoint new-connections initiated
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

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

collect application tcp establishment-time

```
collect application tcp establishment-time
[no] collect application tcp establishment-time
```

Description

Configures collect (non-key) fields for a flow record when in the **config-flow-record** context. Supported on IPv4 and IPv6 flows.



Only one collect field can be specified per line, while a flow record can have multiple collect fields.

Parameter	Description
tcp	Specifies Transmission Control Protocol (TCP) parameters as a non-key field in a flow record.
establishment-time	Specifies TCP connection establishment time as a non-key field in a flow record.

Usage

- ARC inspects the first few packets of each flow and only those packets are copied to the switch.
- Time measurement for established TCP flows is possible during the TCP connection establishment phase, and all three packets are received in the order.

Examples

Adding a TCP connection establishment time to flow record **flow-record-1** as a collect field.

```
switch(config-flow-record) # collect application tcp establishment-time
```

Command History

Release	Modification
10.15.1000	Command introduced

Command Information

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

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.

description

description <DESCRIPTION>
no description <DESCRIPTION>

Description

Configures the description for the following telemetry options:

- flow monitor in the **config-flow-monitor** context
- flow congestion monitor in the **config-flow-congestion-monitor** context
- flow exporter in the **config-flow-exporter** context
- flow record in the **config-flow-record** context

The no form deletes a flow congestion monitor.

Parameter	Description
<code><DESCRIPTION></code>	Displays a string of 256 characters, maximum, including spaces.

Examples

Adding or modifying the description of flow monitor, **flow-monitor-1**:

```
switch(config)# flow monitor flow-monitor-1
switch(config-flow-monitor)# description Used for analyzing basic ipv4 traffic
```

Adding or modifying the description of flow record, **flow-record-1**:

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# description Used for basic traffic analysis
```

Adding or modifying the description of flow exporter, **flow-exporter-1**:

```
switch(config)# flow exporter flow-exporter-1
```

```
switch(config-flow-exporter)# description Exports flows to 10.2.3.45:2055
```

Removing the description of flow monitor, **flow-monitor-1**:

```
switch(config)# flow exporter flow-monitor-1
switch(config-flow-exporter)# no description
```

Command History

Release	Modification
10.16	Command introduced

Command Information

Platforms	Command context	Authority
6300 6400	config-flow-record config-flow-monitor config-flow-congestion-monitor config-flow-exporter	Administrators or local user group members with execution rights for this command.

exporter

```
exporter <EXPORTER-NAME>
no exporter <EXPORTER-NAME>
```

Description

Assigns a flow exporter to **flow monitor** in the **config-flow-monitor** context and **flow congestion monitor** in the **config-flow-congestion-monitor** context.

The no form removes the flow exporter from the selected monitor.

Parameter	Description
<EXPORTER-NAME>	Specifies the name of the flow exporter assigned to the monitor.

Examples

Assigning flow exporter, **flow-exporter-1**, to flow monitor, **flow-monitor-1**:

```
switch(config)# flow monitor flow-monitor-1
switch(config-flow-monitor)# exporter flow-exporter-1
```

Removing a flow exporter assigned to flow monitor, **flow-monitor-1**:

```
switch(config)# flow monitor flow-monitor-1
switch(config-flow-monitor)# no exporter flow-exporter-1
```

Attempting to assign non-existent flow exporter, **flow-exporter-5**, to flow monitor, **flow-monitor-1**:

```
switch(config)# flow monitor flow-monitor-1
switch(config-flow-monitor)# exporter flow-exporter-5
Flow exporter 'flow-exporter-5' does not exist.
switch(config-flow-monitor)#
```

Assigning flow exporter, **flow-exporter-1**, to flow congestion monitor, **congestion-monitor-1**:

```
switch(config)# flow congestion-monitor congestion-monitor-1
switch(config-flow-congestion-monitor)# exporter flow-exporter-1
```

Removing a flow exporter assigned to flow congestion monitor:

```
switch(config)# flow congestion-monitor congestion-monitor-1
switch(config-flow-congestion-monitor)# no exporter flow-exporter-1
```

Attempting to assign non-existent flow exporter, **flow-exporter-5**, to flow congestion monitor, **congestion-monitor-1**:

```
switch(config)# flow congestion-monitor congestion-monitor-1
switch(config-flow-congestion-monitor)# exporter flow-exporter-5
Flow exporter 'flow-exporter-5' does not exist.
switch(config-flow-congestion-monitor)#
```

Assigning more than one flow exporter to flow monitor, **flow-monitor-2**:

```
switch(config)# flow monitor flow-monitor-2
switch(config-flow-monitor)# exporter flow-exporter-2
switch(config-flow-monitor)# exporter flow-exporter-3
```

Attempting to assign flow exporter, **flow-exporter-6**, to monitor, **flow-monitor-3**, when two exporters are assigned:

```
switch(config)# flow monitor flow-monitor-3
switch(config-flow-monitor)# exporter flow-exporter 4
switch(config-flow-monitor)# exporter flow-exporter 5
switch(config-flow-monitor)# exporter flow-exporter 6
Cannot assign more than two flow exporters to a flow monitor.
switch(config-flow-monitor)#
```

Command History

Release	Modification
10.16	Command introduced

Command Information

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

flow exporter

```
flow exporter <name>
  export-protocol ipfix
  description <description>
  destination
    <hostname>
    [vrf <vrfname>]
    <ipaddr>
```

```

    [vrf <vrfname>]
    <ip6addr>
    [vrf <vrfname>]
    type
    hostname-or-ip-addr
    traffic-insight}
    traffic-insight <instance-name>
no ...
template data timeout <timeout>
transport udp <port>

```

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

Parameter	Description
<name>	Name of the flow exporter, up to 64 characters.
export-protocol ipfix	Define an export protocol for the flow exporter. The default ipfix protocol is the only protocol currently available.
description <description>	A description of the flow exporter, up to 256 characters and spaces.
destination	Configure the export destination
<hostname>	The exporter sends flow records to the specified hostname destination. The hostname can be a string of up to 64 characters.
<IPaddr>	The exporter sends flow records to this IPv4 address destination.
<ip6addr>	The exporter sends flow records to this IPv6 address destination.
vrf <vrfname>	You can optionally include the name of the destination VRF in the destination definition on 5420, 6200, 6300, 6400, 8100, 8360, 9300S Switch series.
type	Configure the type of the destination.
hostname-or-ip-addr	Define the destination type as a hostname or IP address.
traffic-insight <name>	Define the destination type as a traffic insight instance.
no ...	Negate any configured parameter.
traffic-insight <INSTANCE-NAME>	Specify the a Traffic Insight instance to be used as the destination.
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.

Parameter	Description
transport udp <port>	Transport protocol and port for sending flow record reports. The default port is port 4739.

Usage

The following table shows the maximum supported of flow monitors and flow exporters for each switch model.

Switch	Maximum Flow Monitors	Maximum Flow Exporters
6300 6400	16	Two flow exporters can be applied to a single flow monitor.

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
```

The following example sets a Traffic Insight instance as the destination for a flow exporter

```
switch(config)# flow exporter exporter-3
switch(config-flow-exporter)# destination type traffic-insight
switch(config-flow-exporter)# destination traffic-insight instance-1
```

The following example adds a destination of each possible type and set **hostname-or-ip-addr** as the type to use:

```
switch(config)# flow exporter exporter-4
switch(config-flow-exporter)# destination collector-1
switch(config-flow-exporter)# destination traffic-insight instance-1
switch(config-flow-exporter)# destination type hostname-or-ip-addr
```

The following example sets an IPv4 address as the destination for a flow exporter

```
switch(config)# flow exporter exporter-1
switch(config-flow-exporter)# destination type hostname-or-ip-addr
switch(config-flow-exporter)# destination 192.168.0.1
```

The following example sets a hostname as the destination for a flow exporter

```
switch(config)# flow exporter exporter-1
switch(config-flow-exporter)# destination type hostname-or-ip-addr
switch(config-flow-exporter)# destination collector1
```

The following example sets an IPv6 address as the destination for a flow exporter

```
switch(config)# flow exporter exporter-1
switch(config-flow-exporter)# destination type hostname-or-ip-addr
switch(config-flow-exporter)# destination 2001:db87::8a2e:370a:7334
```

The following example sets an IPv4 address as the destination for a flow exporter with the VRF to which the IPv4 address belongs

```
switch(config)# flow exporter exporter-2
switch(config-flow-exporter)# destination type hostname-or-ip-addr
switch(config-flow-exporter)# destination 192.0.2.1 vrf VRF1
```

The following example sets a Traffic Insight instance as the destination for a flow exporter

```
switch(config)# flow exporter exporter-3
switch(config-flow-exporter)# destination type traffic-insight
switch(config-flow-exporter)# destination traffic-insight instance-1
```

The following example adds a destination of each possible type and set **hostname-or-ip-addr** as the type to use

```
switch(config)# flow exporter exporter-4
switch(config-flow-exporter)# destination collector-1
switch(config-flow-exporter)# destination traffic-insight instance-1
switch(config-flow-exporter)# destination type hostname-or-ip-addr
```

The following example removes the destination of type **traffic-insight** from a flow exporter:

```
switch(config)# flow exporter exporter-3
switch(config-flow-exporter)# no destination traffic-insight
```

The following example removes the destination of type **hostname-or-ip-addr** from a flow exporter

```
switch(config)# flow exporter exporter-1
switch(config-flow-exporter)# no destination
```

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 {inactive <timeout> | {active <timeout>
  description <description>
  record <name>
```

Description

On HPE Aruba Networking 5420, 6200, 6300, 6400, 8325, 8325H, 8325P, 8100, 8360, 9300, 9300S 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.



HPE Aruba Networking 5420, 6200, 6300, 6400, 8100, 8360 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. 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 <timeout>	Use the cache timeout parameter to define an active or inactive timeout for the flow monitor. A flow monitor closes a flow session that is active for longer than the active timeout or inactive for longer than the inactive timeout. The active timeout range is 30-604800. The default active time out value is 1800 and inactive timeout value is 30.
cache timeout inactive <timeout>	Use the cache timeout parameter to define an inactive timeout for the flow monitor. A flow monitor closes a flow session that is active for longer than the active timeout or inactive for longer than the inactive timeout. For 5420, 6200, 6300, 6400, 8100, 8360, 9300S Switch Series, the inactive timeout range is 30-604800. The default active time out value is 1800 and inactive timeout value is 30.
description	A description up to 256 characters long, including spaces.
exporter <name>	Assign a flow exporter to a flow monitor.
record <name>	(For HPE Aruba Networking 5420, 6200, 6300, 6400, 8100, 8325, 8325P, 8325H, 8360, 9300, 9300S Switch series) 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
```

Create more than the maximum number of allowed flow monitors

```
switch(config)# flow monitor monitor-1
switch(config)# flow monitor monitor-2
<--OUTPUT OMITTED FOR BREVITY-->
switch(config)# flow monitor monitor-16
switch(config)# flow monitor monitor-17
No more than 16 flow monitors can be configured. Another flow monitor
must be removed first.
```

Assign more than one flow exporter to flow monitor ****flow-monitor-2***

```
switch(config)# flow monitor flow-monitor-2
switch(config-flow-monitor)# exporter flow-exporter-2
switch(config-flow-monitor)# exporter flow-exporter-3
```

Add or modify the description of flow record ****flow-record-1****

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# description Used for basic traffic analysis
```

Add or modify the description of flow exporter ****flow-exporter-1****

```
switch(config)# flow exporter flow-exporter-1
switch(config-flow-exporter)# description Exports flows to 10.2.3.45:2055
```

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 <record-name>
  match
    ip {source|destination} address
    ip protocol|version
    ipv6 {source|destination} address
    ipv6 protocol|version
    transport {source | destination} port
  collect
    application name
    application https url
    application dns response-code
    application tls-attributes
    forwarding-status
    drop ingress-exceptions
    egress vlan
    egress interface
    egress queue
    counter {packets|bytes}
    tcp establishment-time
    transport {source | destination} port
    timestamp absolute first
    timestamp absolute last
  description <description>

```

Description

Creates or modifies a flow record and switches to the **config-flow-record** context for the flow record. 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.



A maximum of one flow record can be created for 5420, 8325, 8325H, 8325P, 930 and 9300S Switch series. For 6200, 8360, 8100, 6300, and 6400 Switch series, a maximum of 16 flow records can be created.



It is advised to configure collect egress interface and queue also when threshold is being configured.

Parameter	Description
<record-name>	Name of the flow monitor, up to 64 characters.
match	match traffic according to one or more of the following key attributes: ▪ ip source: Match traffic from the same IPv4 source. ▪ ip destination: Match traffic to the same IPv4 destination. ▪ ip protocol: Match traffic using the same IP version ▪ ip version: Match traffic using the same IP protocol ▪ ipv6 source: Match traffic from an IPv6 source. ▪ ipv6 destination: Match traffic to an IPv6 destination. ▪ ipv6 protocol: Match traffic using the same IPv6 version ▪ ipv6 version: Match traffic using the same IPv6 protocol ▪ transport {source destination} port: Match traffic by source or destination transport port
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. ▪ application dns response-code: Specify the DNS parameters and DNS response code as a non-key field in the flow record. ▪ application tls-attributes: Specifies TLS Attributes as a non-key field in a flow record ▪ application tcp establishment-time: Specifies TCP connection establishment time as a non-key field in a flow record ▪ drop ingress-exceptions: On 5420, 6200, 6300, 6400, 8100, 8325, 8325P, 8325H, 9300 and 10000 Switch series, this enables the IPFIX drop-reason feature, which specifies drop ingress-exceptions as a non-key field in a flow record. ▪ counter bytes: Collect counter data for bytes in the flow. Byte count represents the number of incoming bytes since the previous report. ▪ counter packets: Collect counter data for packets in the flow. Packet count represents the number of incoming packets since the previous report. ▪ egress interface: Specifies an egress interface as a non-key field in a flow record. ▪ egress queue: Specifies an egress queue as a non-key field in a flow record. ▪ egress vlan: Specifies an egress VLAN ID as a non-key field in a flow record. ▪ fowarding status: Specifies forwarding status as a non-key field in a flow record ▪ timestamp absolute first: Collect absolute timestamp of the first packet observed. ▪ timestamp absolute last: Collect absolute timestamp of the last packet observed. ▪ tcp establishment-time: Specifies TCP connection establishment time as a non-key field in a flow record.

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

```

The following example adds egress packet metadata collect fields to **flow record flow-record-1**. Egress interface details will not be available until the MAC entry is not programmed after ARP resolution. As a result, the egress interface information is exported only after the next active/inactive timeout after the MAC entry is programmed.

```

switch(config)# flow record flow-record-1
switch(config-flow-record)# collect egress interface
switch(config-flow-record)# collect egress queue

```

The following example adds a TCP connection establishment time to flow record **flow-record-1** as a collect field. During this process, ARC inspects the first few packets of each flow and only those packets are copied to the switch. Time measurement for established TCP flows is possible during the TCP connection establishment phase, and all three packets are received in the order.

```

switch(config-flow-record)# collect application tcp establishment-time

```

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

```

Creating more than the maximum number of allowed flow records:

```

switch(config)# flow record record-1
switch(config)# flow record record-2
No more than 1 flow record can be configured. Another flow record
must be removed first.

```

Adding IPv4 match fields to flow record **flow-record-1** using the **ip** keyword:

```

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

```

Adding IPv6 match fields to flow record ****flow-record-2****

```
switch(config)# flow record flow-record-2
switch(config-flow-record)# match ipv6 source address
switch(config-flow-record)# match ipv6 destination address
switch(config-flow-record)# match ipv6 protocol
switch(config-flow-record)# match ipv6 version
```

Adding IPv6 collect fields to flow record **flow-record-2**:

```
switch(config)# flow record flow-record-2
switch(config-flow-record)# collect application tcp establishment-time
switch(config-flow-record)# collect egress-vlan
switch(config-flow-record)# collect egress interface
switch(config-flow-record)# collect forwarding-status
switch(config-flow-record)# collect egress queue
```

Removing the IPv4 destination address match field from flow record **flow-record-1** using the **ip** keyword:

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

Removing the IPv4 destination address match field from flow record **flow-record-1** using the **ipv4** keyword:

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

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

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# no match transport destination port
```

Adding queue congestion threshold to **flow-record-1** as a collect field:

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# collect egress queue threshold
```



To export queue threshold attribute, capture-flows threshold profile should be configured on the egress interfaces. If threshold profile is not configured, congested flows can't be detected and the IEs values are exported as 0. If threshold profile is configured, but collect field is not configured, these IEs are not exported. So in either case, congested flows are not known. For more information on features that use this command, refer to the Monitoring Guide for your switch model.

Command History

Release	Modification
10.16	The drop ingress-exceptions parameter was added for 6300 and

Release	Modification
	8325 Switch series.
10.15.1000	The collect tcp establishment-time parameter is introduced on 6200, 630 and 6400 Switch series.
10.15	The egress vlan, egress interface and egress queue parameters were added.
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 on 6300, 6300, 8100 and 8360 Switch Series.

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
  flow
    statistics
      enable
      ingress-only
  tcp-ageout
  track
    icmp
  udp-ageout
  no ...

```

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 to 86400. Default: 15.

Parameter	Description
interface-flow-limit	Configures global concurrent flow limit for flow tracking enabled interfaces. Range:For 6300 and 6400 Switch series, 64 to 25000.. Default: none.
flow	Configures flow parameters.
statistics	Configures flow statistics.
enable	Enables flow statistics collection.
ingress-only	Restricts flow tracking to the ingress direction. NOTE: The application visibility feature requires traffic flows in both the ingress and egress directions. Configuring this option prevents features that depend on bi-directional flows from functioning.
tcp-ageout	Configures age-out time for established TCP flows in seconds. Range: 120 to 86400. Default: 600.
track icmp	Enables 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 to 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
10.17	The flow statistics ingress-only parameter was introduced.
10.14	The flow statistics enable parameter is introduced.
10.14	The track icmp parameter is introduced.
10.13	Command introduced. on 6300 and 6400 Switch Series

Command Information

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

ip|ipv6 flow monitor (interface)

```
[no] ip|ipv6 flow monitor (interface)
```

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

Enable a flow monitor configuration named **flow-monitor-1** for IPv4 traffic on a physical interface.

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

Associate a flow monitor configuration named **flow-monitor-2** for IPv4 traffic on a LAG interface.

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

Associate a flow monitor configuration named **flow-monitor-3** for IPv6 traffic on a physical interface.

```
switch(config)# interface 1/1/1
switch(config-if)# ip flow monitor flow-monitor-3 in
```

Associate a flow monitor configuration named **flow-monitor-4** for IPv6 traffic on a physical interface.

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

Command History

Release	Modification
10.11	Command introduced on 6300, 6400 and 8325 Swtich 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.

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
```

Command History

Release	Modification
10.14	Command introduced on 6300 and 6400 Switch series.

Command Information

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

show flow exporter

```
show flow exporter [<name>] [statistics] [vsx-peer]
```

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)
- Rejected (Internal error: the specified destination VRF is invalid)
- Rejected (Internal error: the specified destination VRF is not ready)

Parameter	Description
<name>	Name of the flow exporter.
statistics	Adds statistical information about the flow exporter to the output.
vsx-peer	Displays flow collector configuration for the VSX peer.

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 statistics
```

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

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

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

```
-----
Reports sent          : 5
-----
```

Display information with no flow exporters configured

```
switch# show flow exporter
No flow exporters configured.
switch# show flow exporter statistics
No flow exporters configured.
```

Display a flow exporter's information with TI as a destination

```
switch# show flow exporter exporter-5
```

```
-----
Exporter Name         : exporter-5
-----
```

```
Description          : Exporter configured with TI as the destination
Status               : Rejected (Destination type is Traffic Insight, but the
specified Traffic Insight instance does not exist)
Export Protocol       : ipfix
Destination Type      : Traffic Insight
Destination           : instance-1
Transport Configuration
  Protocol            : UDP
  Port                : 2055
```

Display information for a flow exporter that has multiple destinations of different types configured and *hostname-or-ip-addr* is specified as the destination type to use

```
switch# show flow exporter exporter-6
```

```
-----
Flow exporter 'exporter-6'
```

```
-----
Description          : TI and hostname configured, but type is hostname-or-ip-addr
Status               : Accepted
Export Protocol       : ipfix
```

```
Destination Type      : Hostname or IP address
Destination           : collector-1
Destination VRF       : mgmt
Transport Configuration
Protocol              : UDP
Port                  : 4821
```

Command History

Release	Modification
10.11	Command introduced on 6300, 6400, 8100 and 8630 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.

show flow monitor

```
show flow monitor [statistics]
show flow monitor <MONITOR-NAME>[statistics]
```

Description

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

Parameter	Description
<name>	Name of the flow monitor.
statistics	For 5420, 6200, 6300, 6400, 8100, 8325, 8325P, 8325H, 8360, 9300, 9300S Switch series, include the statistics parameter to display additional flow and cache statistics.

Usage

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

- Possible status types for all switches
 - Accepted
 - Rejected (Internal error: monitor does not exist)
 - Rejected (The state of one or more of the assigned flow exporters is rejected)
- Possible Status types for 5420 ,6200, 6300, 6400, 8100, 8325 ,8325P, 8325H, 8360, 9300, 9300S Switch series:

- Rejected (A record must be assigned to the monitor, but no record is assigned)
- Rejected (The state of the assigned record is rejected)
- Rejected (Internal error: failure in processing the record configuration)

The possible statistics for a flow monitor are:

Statistics name	Meaning	Switches that support this statistic
Current Entries	Current number of flows in the flow cache for this flow monitor	6300 6400 (v2 profile only)
Flows Added	Total number of flows added to the flow cache for this flow monitor since it was created	6300 6400 (v2 profile only)
Total Flows Terminated	Total number of flows removed from the flow cache for this flow monitor since it was created due to any flow end reason	6300 6400
Flows Aged	Number of flows removed from the flow cache for this flow monitor since it was created due to active or inactive timeout	6300 6400
Inactive Timeout	Number of flows removed from the flow cache for this flow monitor since it was created due to inactive cache timeout.	6300 6400
Forced End	Number of flows removed from the flow cache for this flow monitor since it was created due to some external event. For example, the shutdown of a flow monitor.	6300 6400

Examples

Display information for a flow monitor on 6200, 6300, 6400, 8100 or 8360 Switch series:

```
switch# show flow monitor 'monitor-1'
-----
Flow monitor 'monitor-1'
-----
Description          : Used for IPv4 traffic analysis
Status               : Accepted
Flow Record          : record-1
Flow Exporter(s)     : exporter-1, exporter-2
Cache Configuration
Inactive Timeout     : 1800
Active Timeout       : 300
-----
```

Display information and statistics for a flow monitor on 5420, 6200, 6300, 6400, 8100 or 8360 Switch series:

```
show flow monitor statistics
```

```
-----  
Flow monitor 'monitor-1'  
-----
```

```
Current Entries      : 2  
Flows Added          : 6  
Total Flows Terminated : 4  
  Flows Aged         : 2  
    Active Timeout   : 1  
    Inactive Timeout  : 1  
End of Flow Detected : 2  
Forced End           : 0  
Flows Aged           : 4
```



The flow monitor statistics counters will be reset to zero after a VSF ISSU switchover.

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-monitor	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



IPv6 related commands are only applicable to switches that support IPv6 protocol.

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
  transport destination port
  transport source port
Collect Fields
  application name
  counter bytes
  counter packets
  application https URL
  drop ingress-exceptions
```

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

  transport destination port
  transport source port
Collect Fields
  application name
  application https url
  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
```

```

        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 for a specific flow record

```

switch# show flow record record-3
-----
Flow record  'record-3'
-----
Description          : Used for IPv4 traffic analysis
Status               : Rejected (Incomplete match fields. The mandatory match
fields are: version, source address,
destination address, protocol, transport destination port, and transport source
port.)
Match Fields
        ipv4 destination address
        ipv4 protocol
        ipv4 source address
        ipv4 version
Collect Fields
        application name
        counter bytes
        counter packets

```

Display information with no flow records configured

```

switch# show flow record
No flow records configured

```

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.

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 on a 6300 or 6400 Switch series.

```
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
  Ingress-only                : 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
10.17	Added Ingress-only information for Statistics Collection.
10.14	Added role based IPFIX and ICMP ageout information.
10.13	Command introduced on 6300 and 6400 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.

show running-config

```
show running-config
    flow [ | record | interface [lag]][vsx-peer]
```

Description

Shows all IPFIX flow configuration.



To show the VSX peer switch configuration, insert **vsx-peer** at the end.

Parameter	Description
flow	Displays IPFIX flow configuration
record	Displays flow record configuration.
interface	Displays interfaces with flow monitors applied.
lag	Displays LAG interfaces with flow monitors applied.
vsx-peer	Displays VSX peer switch information.

Examples

Showing current IPFIX flow configurations:

```
switch# show running-config flow
flow exporter exporter-1
  description This is an exporter
  destination 40.0.0.2 vrf default
  transport udp 2055
flow exporter exporter-2
  destination 4000:0:0:0:0:0:2 vrf default
flow record record-1
  description This is a record
  match ipv4 destination address
  match ipv4 protocol
  match ipv4 source address
  match ipv4 version
  match transport destination port
  match transport source port
  collect counter bytes
  collect counter packets
  collect drop ingress-exceptions
  collect timestamp absolute first
  collect timestamp absolute last
flow record record-2
  description This is a record
  match ipv6 destination address
  match ipv6 protocol
  match ipv6 source address
  match ipv6 version
  match transport destination port
  match transport source port
  collect counter bytes
  collect counter packets
  collect timestamp absolute first
  collect timestamp absolute last
flow monitor monitor-1
  description This is a monitor
  cache timeout active 30
  exporter exporter-1
  record record-1
flow monitor monitor-2
  description This is a monitor
  cache timeout active 30
  exporter exporter-2
  record record-2
interface lag 1
  ip flow monitor monitor-1 in
interface 1/1/1
  ip flow monitor monitor-1 in
  ipv6 flow monitor monitor-2 in
```

Showing current IPFIX flow monitor configurations:

```
switch# show running-config flow monitor
flow monitor monitor-1
  description This is a monitor
  cache timeout active 30
  exporter exporter-1
  record record-1
flow monitor monitor-2
```

```
description This is a monitor
cache timeout active 30
exporter exporter-2
record record-2
```

Showing current IPFIX flow congestion monitor configurations:

```
switch# show running-config flow congestion-monitor
flow congestion-monitor cng-mon-1
  description This is a monitor
  exporter exporter-1
  ingress-interface 1/1/10
  ingress-interface 1/1/11
```

Showing operational status of the current configuration:

```
switch# show running-config
flow exporter exporter-1
  description This is an exporter
  destination 40.0.0.2
  transport udp 2055
flow congestion-monitor monitor-1
  exporter exporter-1
  ingress-interface 1/1/1
interface 1/1/1
  ip-all flow congestion-monitor m out queue 0-7
    ! disabled - Blocked by a higher precedence feature
interface 1/1/2
  ip-all flow congestion-monitor m out queue 0-7
    ! disabled - Blocked by a higher precedence feature
```

Command History

Release	Modification
10.16	Command introduced

Command Information

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

show ip flow-table-utilization

```
show ip flow-table-utilization
```

Description

Displays the current IP flow table utilization per linecard (LC), including both the maximum capacity and the current number of active flows.

Examples

Configuring the endpoint new connections accepted alert threshold:

```
switch# show ip flow-table-utilization
Line Modules
=====
Slot name | flow-table capacity | flow-count
-----|-----|-----
1/1       | 28672               | 15000
1/3       | 28672               | 5000
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

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

show IPv4 flow endpoint statistics

```
show ip flows endpoint <A.B.C.D> [vrf <vrf_name>]
```

Description

Displays the number of active flows of VRF with name **vrf_name**, the flows observed in the last 30 seconds, and details about new connections initiated and accepted by this endpoint during the same period. If the **vrf_name** is not specified, default VRF active flows are displayed.

Parameter	Description
IPv4 address	Show endpoint statistics for the IPv4 address. <i>Required</i>
vrf	Display the routes in the VRF. <i>Optional</i>
vrf_name	Specify the VRF name. <i>Required</i>

The output displays the following information:

- **VRF:** The specified vrf name.
- **Endpoint address:** The specified IPv4 address.
- **Active flows:** Number of active flows for the endpoint.
- **Flows seen last 30sec:** Number of flows observed in the last 30 seconds.
- **No of new connections (last-30sec):**
 - **Initiated:** Number of new connections initiated by the endpoint in the last 30 seconds.
 - **Accepted:** Number of new connections accepted for the endpoint in the last 30 seconds.

Examples

This example shows the number of active flows of the default VRF:

```
switch# show ip flows endpoint 192.168.1.1

VRF                                     : default
Endpoint address                       : 192.168.1.1
Active flows                           : 10
Flows seen last 30sec                  : 20
Number of new connections (last-30sec):
  Initiated                           : 2
  Accepted                             : 0
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

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

show IPv6 flow session statistics

```
show ipv6 flows endpoint <A.B.C.D> [vrf <vrf_name>]
```

Description

Displays the number of active flows of VRF with name **vrf_name**, the flows observed in the last 30 seconds, and details about new connections initiated and accepted by this endpoint during the same period. If the **vrf_name** is not specified, default VRF active flows are displayed.

Parameter	Description
IPv6 address	Show endpoint statistics for the IPv6 address.

Parameter	Description
	<i>Required</i>
vrf	Display the routes in the VRF. <i>Optional</i>
vrf_name	Specify the VRF name. <i>Required</i>

The output displays the following information:



- **VRF:** The specified vrf name.
- **Endpoint address:** The specified IPv6 address.
- **Active flows:** Number of active flows for the endpoint.
- **Flows seen last 30sec:** Number of flows observed in the last 30 seconds.
- **No of new connections (last-30sec):**
 - **Initiated:** Number of new connections initiated by the endpoint in the last 30 seconds.
 - **Accepted:** Number of new connections accepted for the endpoint in the last 30 seconds.

Examples

This example shows the number of active flows of the specified VRF:

```
switch# show ipv6 flows endpoint 2001::1 vrf red

VRF                               : red
Endpoint address                  : 2001::1
Active flows                      : 10
Flows seen last 30sec             : 20
Number of new connections (last-30sec):
  Initiated                      : 2
  Accepted                       : 0
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400 (v2 profile only)	Manager (#)	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, that configuration is used.

For 6200, 6300, 6400, 8360, 8100 and 9300S Switch Series, 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.

AOS-CX provides multiple features that assist with monitoring queuing behavior such as Queue Monitoring.

The Queue Monitoring feature allows the switch to collect queue statistics at specific time intervals, including queue depth, and display this information as a **queue statistics history**. This data provides a history of activity for each monitored queue and can be viewed directly on the switch. These aggregate statistics can be used to monitor network health, troubleshoot network issues, and identify normal network behavior patterns and performance anomalies. The feature is disabled by default, but can be enabled on a maximum of 52 interfaces per switch. For more information on enabling this feature, see [queue-monitor](#).

Queue statistics history data can be displayed in the command-line interface as a list or a histogram reflecting the last eight hours of queue statistics history, or as a table showing the previous five minutes of queue statistics history.

As performance anomalies are identified, the interface queue experiencing the anomaly can further analyzed to identify the cause of the behavior. For example, traffic patterns for flows competing for network bandwidth can be examined to determine if corrective action is needed to address potential problems.

If no queue depth is specified, the output of the **show interface queue-monitor** command displays statistics for queues with a minimum depth of one kbyte, which includes all queues which have reached a depth greater than zero.

Queue statistics history

Queue monitoring stores collected queue depth data in a time-series database within the switch. This data, accessible on the switch for reviewing the queue statistics history, helps with identifying network issues or behavior patterns.

Queue depth

Queue depth is the memory size required to buffer packets in a queue during congestion. The value changes depending on the arrival rate of the packets destined to egress a given queue, and the rate at which that queue is able to transmit packets.

Data retention limits

Data retention limits are enforced within the switch by reducing time precision of the stored data within specific windows of time. The *retention duration* for queue statistics data is derived from the frequency of the polling interval.

Table 1: Queue monitoring data retention schedule

Polling interval	Maximum statistic age
1 seconds	8 hours
5 seconds	8 hours
10 seconds	8 hours
30 seconds	8 hours
60 seconds	8 hours



The switch has limited storage for retaining information. For longer monitoring durations, use an external monitor with the queue monitoring feature to collect and store data.

Using collected data

Queue monitoring provides the ability to look into switch congestion-related performance issues and identify when problems have occurred. The methods available for reviewing performance issues are:

- Congestion events
- Per-interface congestion histograms
- Per-interface congestion-detailed view

As anomalies are identified, the interface queue that experienced the anomaly is the starting point for identifying the cause of the behavior. Examine traffic patterns for flows competing for network bandwidth to determine if corrective action is needed to address potential problems.

QoS commands

clear queue-monitor interface

```
clear queue-monitor interface [<IFNAME>|<IFRANGE>]
```

Description

Clear the data collected on an interface since enabling queue monitoring. If no interfaces are specified, perform a clear of data collected on all interfaces with queue monitoring enabled.

Parameter	Description
<IFNAME>	Name of the interface. Format: <MEMBER>/<SLOT>/<PORT>
<IFRANGE>	Range of interfaces. Format: <MEMBER>/<SLOT>/<PORT>-<MEMBER>/<SLOT>/<PORT>

Examples

Clearing queue monitor data collected on interface 1/1/1:

```
config)# clear queue-monitor interface 1/1/1
Warning: clearing collected queue monitor statistics will be reflected
in all CLI sessions, any agents running in the analytics engine, and any
external systems monitoring switch statistics.
Continue (y/n)? y
```

Clearing queue monitor data collected on interfaces 1/1/1 and 1/1/2:

```
config)# clear queue-monitor interface 1/1/1-1/1/2
Warning: clearing collected queue monitor statistics will be reflected
in all CLI sessions, any agents running in the analytics engine, and any
external systems monitoring switch statistics.
Continue (y/n)? y
```

Clearing all queue monitor data collected on all interfaces:

```
config)# clear queue-monitor interface
Warning: clearing collected queue monitor statistics will be reflected
in all CLI sessions, any agents running in the analytics engine, and any
external systems monitoring switch statistics.
Continue (y/n)? y
```

Related Commands

Command	Description
queue-monitor	This command enables the queue monitoring feature, allowing the switch to collect queue statistics at 10-second time intervals. These aggregate statistics can be used to monitor network health, troubleshoot network issues, and identify normal network behavior patterns and performance anomalies. This feature is disabled by default, but can be enabled on a maximum of 52 interfaces per switch.

Command History

Release	Modification
10.17	Command introduced on 5420, 6200, 6300 and 6400 Switch series. An additional parameter for slot <value> is available only on the 6400 Switch series.

Command Information

Platforms	Command context	Authority
6300M/F (except for JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL669A, JL670A, JL762A, JL761A, JL760A) 6400v2	config	Administrators or local user group members with execution rights for this command.

queue-monitor

```
queue-monitor
no queue-monitor
```

Description

This command enables the queue monitoring feature, allowing the switch to collect queue statistics. then display this information as a queue statistics history. These aggregate statistics can be used to monitor network health, troubleshoot network issues, and identify normal network behavior patterns and performance anomalies. This feature is disabled by default, but can be enabled on a maximum of 52 interfaces per switch.

Examples

Enabling the queue monitoring feature on interface 1/1/1:

```
(config)# interface 1/1/1
(config-if)# queue-monitor
```

Disabling the queue monitoring feature on interface 1/1/1:

```
(config)# interface 1/1/1
(config-if)# no queue-monitor
```

Related Commands

Command	Description
show queue-monitor status	This command is used to view the global state of the queue monitor feature. Output includes the memory consumed by the feature, the statistics being monitored for collection, and the interfaces currently enabled for monitoring.

Command History

Release	Modification
10.17	Command introduced on 5420, 6200, 6300 and 6400 Switch series.

Command Information

Platforms	Command context	Authority
6300M/F (except for JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL669A, JL670A, JL762A, JL761A, JL760A) 6400v2	config-if	Administrators or local user group members with execution rights for this command.

show interface queue-monitor

```
show interface {<IFNAME>} queue-monitor
  histogram [filter [since <1-10000> days|hours|minutes|seconds] [type {queue-depth [slot
    <value>}}]
  list [filter [since <1-10000> days|hours|minutes|seconds] [threshold {kbps <1-
    4294967295>}|{kbytes <1-4294967295>}|{packets <1-4294967295>}}] [type {queue-depth
    [slot <value>}}]] [slot <slot-id>]
  table [filter [threshold {kbps <1-4294967295>}|{kbytes <1-4294967295>}|{packets <1-
    4294967295>}}] [type {queue-depth [slot <value>}}]]
```

Description

Display queue statistics history in the CLI as either a list or a histogram, or as a table, showing the previous five minutes of queue statistics history.

Parameter	Description
interface [<IFNAME>]	Specifies the name of an Ethernet port or LAG on the switch. If the optional <IFNAME> parameter is omitted, the output of this command displays all queue statistic data for all interfaces. Format: <MEMBER>/<SLOT>/<PORT> or lag <ID>
histogram list table	Specifies the output format for queue monitor output. ■ histogram - Queue statistics history is presented in histogram format for the last 8 hours.

Parameter	Description
	- list - Queue statistics history is presented in the list format for the last 8 hours - table - Queue statistics history is presented in tabular format for the last 5 minutes.
filter	The list format is able to display all of the collected queue statistics history data when no filters are applied. To limit the information displayed in the output of this command, use the filter parameters to filter the output to the format and data you require.
since <1-10000> days hours minutes seconds	Show events no older than a specified age, measured in seconds, days, hours or minutes.
type {queue-depth [slot <value>]}	Filter queue monitoring data by the specified filter type. If you do not specify a filter type, the type defaults to queue-depth. To see filtered output for other statistics types, you must specify the type. queue-depth: Show queue depth data. For chassis-based platforms such as 6400 switch series, include the optional slot <value> parameters to view queue depth for individual slots. NOTE: The slot <value> parameter is supported only on 6400 switches.
threshold {kbps kbytes packets <1-4294967295>}	Show data that exceeds a specified threshold in kilobytes, number of packets, or kilobits per second.

Examples

Displaying the queue statistics history data for all monitored interfaces in list format:

```
switch# show interface queue-monitor list
Time Range: 8 hours (2023-08-07 09:11:30 to 2023-08-08 17:11:30 (UTC+00:00))
```

Time	Interface	Queue	Statistic	Value
2023-08-08 17:11:10	1/1/8	4	Queue Depth	12345 kbytes
2023-08-08 17:11:10	1/1/9	4	Queue Depth	5432 kbytes
2023-08-08 17:11:20	1/1/8	4	Queue Depth	12345 kbytes
2023-08-08 17:11:20	1/1/9	4	Queue Depth	5432 kbytes
2023-08-08 17:11:30	1/1/8	4	Queue Depth	12345 kbytes
2023-08-08 17:11:30	1/1/9	4	Queue Depth	5432 kbytes

The following example displays a portion of output of the **show interface 1/1/8 queue-monitor** command, showing queue statistics history data in tabular format for interface 1/1/8 with a min-depth of 2 kbytes.

```
switch# show interface 1/1/8 queue-monitor table min-depth kbytes 2
1/1/8 Statistics History
Collection Interval: 10 seconds
Time Range: 2023-08-08 17:07:00 to 2023-08-08 17:12:00 (UTC)
```

Min-Depth Filter: 2 KB

Time (s)	Queue Depth - bytes							
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
-300	-	-	-	-	-	-	-	-
-295	-	-	-	-	-	-	-	-
-290	-	-	-	-	-	-	-	-
-285	-	-	-	-	-	-	-	-
-280	-	-	-	-	-	-	-	-
-275	-	-	-	-	-	-	-	-
-270	-	-	-	-	-	-	-	-
-265	-	-	-	-	-	-	-	-
-260	-	-	-	-	-	-	-	-
-255	-	-	-	-	-	-	-	-
-250	-	2078	-	-	-	-	-	-
-245	-	9914	-	-	-	-	-	-
-240	-	171102	-	-	-	-	-	-
-235	-	996262	-	-	-	-	-	-
-230	-	5444	-	-	-	-	-	-
-225	-	-	-	-	-	-	-	-
-220	-	-	-	-	-	-	-	-
-215	-	-	-	-	-	-	-	-
-210	-	-	-	-	-	-	-	-
-205	-	-	-	-	-	-	-	-
-200	-	-	-	-	-	-	-	-
...								

Displaying the queue statistics history data for interface 1/1/8 in histogram format.

```
switch(config)# show interface 1/1/8 queue-monitor histogram
1/1/8 Statistics History
Time Range: 2023-08-08 09:12:15 to 2023-08-08 17:12:15 (UTC)
Samples per Queue: 13
```

Depth (KB)	Count at Depth							
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
0	13	13	13	13	7	13	13	10
1-64	0	0	0	0	6	0	0	3
65-128	0	0	0	0	0	0	0	0
129-256	0	0	0	0	0	0	0	0
257-512	0	0	0	0	0	0	0	0
513-1024	0	0	0	0	0	0	0	0
1025-2048	0	0	0	0	0	0	0	0
2049-4096	0	0	0	0	0	0	0	0
4097-8192	0	0	0	0	0	0	0	0
8193-16384	0	0	0	0	0	0	0	0
16385-32768	0	0	0	0	0	0	0	0
32769-65536	0	0	0	0	0	0	0	0

Output of the global **show** command when queue monitoring is not enabled on any interfaces.

```
switch(config)# show interface queue-monitor list
Queue monitoring is not enabled on any interfaces.
```

Output of the **show** command when the interface requested does not have queue monitoring enabled.

```
switch(config)# show interface 1/1/8 queue-monitor table
Interface 1/1/8 does not have queue monitoring enabled.
```

Output of the global **show** command with the list presentation specified when there is no relevant data to show. Note: the table and histogram presentation formats output all enabled interface data.

```
switch(config)# show interface queue-monitor list
Enabled interfaces have not collected any data.
```

Output of the **show** command for an interface with the list presentation specified but there is no relevant data to show.

```
switch(config)# show interface 1/1/8 queue-monitor list
Interface 1/1/8 has not collected any data.
```

On a 6400v2 Switch series, the following command will display the queue statistics history data in list format for interface **1/3/1** on all the active Ingress slots.

```
switch# show interface 1/3/1 queue-monitor list
Interface Queue Congestion History
Time Range: 2024-12-06 02:23:59 to 2024-12-06 10:23:59 (UTC)
Sort: Time-ascending, Interface, Ingress-Slot, Queue
Min-Depth Filter: 1 KB
```

Time	Interface	Ingress-Slot	Queue	Depth (B)
2024-12-06 10:23:40	1/3/1	3	7	1154
2024-12-06 10:23:40	1/3/1	4	7	1086
2024-12-06 10:23:40	1/3/1	5	7	1103
2024-12-06 10:23:50	1/3/1	3	2	1207
2024-12-06 10:23:50	1/3/1	4	2	1238
2024-12-06 10:23:50	1/3/1	5	2	1420

On a 6400v2 Switch series, display the queue statistics history data in list format for interface **1/3/1** on Ingress slot **3**.

```
6400# show interface 1/3/1 queue-monitor list slot 3
Interface Queue Congestion History
Time Range: 2024-12-06 02:23:59 to 2024-12-06 10:23:59 (UTC)
Sort: Time-ascending, Interface, Ingress-Slot, Queue
Min-Depth Filter: 1 KB
```

Time	Interface	Ingress-Slot	Queue	Depth (B)
2024-12-06 10:23:40	1/3/1	3	7	1154
2024-12-06 10:23:50	1/3/1	3	2	1207

On a 6400v2 Switch series, display the queue statistics history data in tabular format for interface **1/4/3**.

```
6400(config)# show interface 1/4/3 queue-monitor table
Ingress Slot ID: 4
1/4/3 Congestion History
Collection Interval: 10 seconds
Time Range: 2024-10-30 10:31:10 to 2024-10-30 10:36:10 (UTC)
Min-Depth Filter: 1 KB
```

Queue Depth - bytes								
Time (s)	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
-300	-	-	-	-	-	-	-	-
-290	-	-	-	-	-	-	-	-
-280	-	-	-	-	-	-	-	-
-270	-	-	-	-	-	-	-	-
-260	-	-	-	-	-	-	-	-
-250	-	-	-	-	-	-	-	-
-240	-	-	-	-	-	-	-	-
-230	-	-	-	-	-	-	-	-
-220	-	-	-	-	-	-	-	-
-210	-	-	-	-	-	-	-	-
-200	-	27596	-	-	-	-	-	-
-190	-	212172	-	-	-	-	-	-
-180	-	458598	-	-	-	-	-	-
-170	-	704870	-	-	-	-	-	-
-160	-	951500	-	-	-	-	-	-
-150	-	1197926	-	-	-	-	-	-
-140	-	1444198	-	-	-	-	-	-
-130	-	1690828	-	-	-	-	-	-
-120	-	1937254	-	-	-	-	-	-
-110	-	2183526	-	-	-	-	-	-
-100	-	2429798	-	-	-	-	-	-
-90	-	2676428	-	-	-	-	-	-
-80	-	2922496	-	-	-	-	-	-
-70	-	3168972	-	-	-	-	-	-
-60	-	3415756	-	-	-	-	-	-
-50	-	3662182	-	-	-	-	-	-
-40	-	3908454	-	-	-	-	-	-
-30	-	4154726	-	-	-	-	-	-
-20	-	4400998	-	-	-	-	-	-
-10	-	-	-	-	-	-	-	-

Ingress Slot ID: 5
 1/4/3 Congestion History
 Collection Interval: 10 seconds
 Time Range: 2024-10-30 10:31:10 to 2024-10-30 10:36:10 (UTC)
 Min-Depth Filter: 1 KB

Queue Depth - bytes								
Time (s)	Q0	Q1	Q2	Q3	Q4	Q5	Q6	Q7
-300	-	-	-	-	-	-	-	-
-290	-	-	-	-	-	-	-	-
-280	-	-	-	-	-	-	-	-
-270	-	-	-	-	-	-	-	-
-260	-	-	-	-	-	-	-	-
-250	-	-	-	-	-	-	-	-
-240	-	-	-	-	-	-	-	-
-230	-	-	-	-	-	-	-	-
-220	-	-	-	-	-	-	-	-
-210	-	-	-	-	-	-	-	-
-200	-	27596	-	-	-	-	-	-
-190	-	212172	-	-	-	-	-	-
-180	-	458598	-	-	-	-	-	-
-170	-	704870	-	-	-	-	-	-
-160	-	951500	-	-	-	-	-	-
-150	-	1197926	-	-	-	-	-	-
-140	-	1444198	-	-	-	-	-	-
-130	-	1690828	-	-	-	-	-	-
-120	-	1937254	-	-	-	-	-	-
-110	-	2183526	-	-	-	-	-	-
-100	-	2429798	-	-	-	-	-	-

```

-90          - 2676428      -      -      -      -      -      -
-80          - 2922496      -      -      -      -      -      -
-70          - 3168972      -      -      -      -      -      -
-60          - 3415756      -      -      -      -      -      -
-50          - 3662182      -      -      -      -      -      -
-40          - 3908454      -      -      -      -      -      -
-30          - 4154726      -      -      -      -      -      -
-20          - 4400998      -      -      -      -      -      -
-10          -      -      -      -      -      -      -      -
Ingress Slot ID: 6
1/4/3 Congestion History
Collection Interval: 10 seconds
Time Range: 2024-10-30 10:31:10 to 2024-10-30 10:36:10 (UTC)
Min-Depth Filter: 1 KB
Queue Depth - bytes
Time (s)      Q0      Q1      Q2      Q3      Q4      Q5      Q6      Q7
-----
-300          -      -      -      -      -      -      -      -
-290          -      -      -      -      -      -      -      -
-280          -      -      -      -      -      -      -      -
-270          -      -      -      -      -      -      -      -
-260          -      -      -      -      -      -      -      -
-250          -      -      -      -      -      -      -      -
-240          -      -      -      -      -      -      -      -
-230          -      -      -      -      -      -      -      -
-220          -      -      -      -      -      -      -      -
-210          -      -      -      -      -      -      -      -
-200          -      27596      -      -      -      -      -      -
-190          -      212172      -      -      -      -      -      -
-180          -      458598      -      -      -      -      -      -
-170          -      704870      -      -      -      -      -      -
-160          -      951500      -      -      -      -      -      -
-150          -      1197926      -      -      -      -      -      -
-140          -      1444198      -      -      -      -      -      -
-130          -      1690828      -      -      -      -      -      -
-120          -      1937254      -      -      -      -      -      -
-110          -      2183526      -      -      -      -      -      -
-100          -      2429798      -      -      -      -      -      -
-90           -      2676428      -      -      -      -      -      -
-80           -      2922496      -      -      -      -      -      -
-70           -      3168972      -      -      -      -      -      -
-60           -      3415756      -      -      -      -      -      -
-50           -      3662182      -      -      -      -      -      -
-40           -      3908454      -      -      -      -      -      -
-30           -      4154726      -      -      -      -      -      -
-20           -      4400998      -      -      -      -      -      -
-10           -      -      -      -      -      -      -      -

```

On a 6400v2 Switch series, display the queue statistics history data in tabular format for interface **1/4/3** on Ingress slot **4**.

```

6400(config)# show interface 1/4/3 queue-monitor table slot 4
Ingress Slot ID: 4
1/4/3 Congestion History
Collection Interval: 10 seconds
Time Range: 2024-10-30 10:31:10 to 2024-10-30 10:36:10 (UTC)
Min-Depth Filter: 1 KB
Queue Depth - bytes
Time (s)      Q0      Q1      Q2      Q3      Q4      Q5      Q6      Q7
-----
-300          -      -      -      -      -      -      -      -

```

-290	-	-	-	-	-	-	-	-
-280	-	-	-	-	-	-	-	-
-270	-	-	-	-	-	-	-	-
-260	-	-	-	-	-	-	-	-
-250	-	-	-	-	-	-	-	-
-240	-	-	-	-	-	-	-	-
-230	-	-	-	-	-	-	-	-
-220	-	-	-	-	-	-	-	-
-210	-	-	-	-	-	-	-	-
-200	-	27596	-	-	-	-	-	-
-190	-	212172	-	-	-	-	-	-
-180	-	458598	-	-	-	-	-	-
-170	-	704870	-	-	-	-	-	-
-160	-	951500	-	-	-	-	-	-
-150	-	1197926	-	-	-	-	-	-
-140	-	1444198	-	-	-	-	-	-
-130	-	1690828	-	-	-	-	-	-
-120	-	1937254	-	-	-	-	-	-
-110	-	2183526	-	-	-	-	-	-
-100	-	2429798	-	-	-	-	-	-
-90	-	2676428	-	-	-	-	-	-
-80	-	2922496	-	-	-	-	-	-
-70	-	3168972	-	-	-	-	-	-
-60	-	3415756	-	-	-	-	-	-
-50	-	3662182	-	-	-	-	-	-
-40	-	3908454	-	-	-	-	-	-
-30	-	4154726	-	-	-	-	-	-
-20	-	4400998	-	-	-	-	-	-
-10	-	-	-	-	-	-	-	-

On a 6400v2 Switch series, display the queue statistics history data for interface **1/3/1** in histogram format on specified slot.

```
6400# show interface 1/3/1 queue-monitor histogram slot 3
Ingress Slot ID: 3
1/3/1 Congestion History
Time Range: 2024-12-06 02:59:19 to 2024-12-06 10:59:19 (UTC)
Samples per queue: 750
Count at Depth
Depth (KB)      Q0      Q1      Q2      Q3      Q4      Q5      Q6      Q7
-----
0                250     250     250     250     250     250     250     249
1-64             0        0        0        0        0        0        0        1
65-128           0        0        0        0        0        0        0        0
129-256          0        0        0        0        0        0        0        0
257-512          0        0        0        0        0        0        0        0
513-1024         0        0        0        0        0        0        0        0
1025-2048        0        0        0        0        0        0        0        0
2049-4096        0        0        0        0        0        0        0        0
4097-8192        0        0        0        0        0        0        0        0
8193-16384       0        0        0        0        0        0        0        0
16385-32768     0        0        0        0        0        0        0        0
```

On a 6400v2 Switch series, display the queue statistics history data for interface **1/3/1** in histogram format for all slots.

```
6400# show interface 1/3/1 queue-monitor histogram
Ingress Slot ID: 3
1/3/1 Congestion History
```

Time Range: 2024-12-06 02:59:19 to 2024-12-06 10:59:19 (UTC)

Samples per queue: 250

Depth (KB)	Count at Depth							Q7
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	
0	250	250	250	250	250	250	250	249
1-64	0	0	0	0	0	0	0	1
65-128	0	0	0	0	0	0	0	0
129-256	0	0	0	0	0	0	0	0
257-512	0	0	0	0	0	0	0	0
513-1024	0	0	0	0	0	0	0	0
1025-2048	0	0	0	0	0	0	0	0
2049-4096	0	0	0	0	0	0	0	0
4097-8192	0	0	0	0	0	0	0	0
8193-16384	0	0	0	0	0	0	0	0
16385-32768	0	0	0	0	0	0	0	0
32769-65536	0	0	0	0	0	0	0	0

Ingress Slot ID: 4

1/3/1 Congestion History

Time Range: 2024-12-06 02:59:19 to 2024-12-06 10:59:19 (UTC)

Samples per queue: 750

Depth (KB)	Count at Depth							Q7
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	
0	250	250	250	250	250	250	250	249
1-64	0	0	0	0	0	0	0	1
65-128	0	0	0	0	0	0	0	0
129-256	0	0	0	0	0	0	0	0
257-512	0	0	0	0	0	0	0	0
513-1024	0	0	0	0	0	0	0	0
1025-2048	0	0	0	0	0	0	0	0
2049-4096	0	0	0	0	0	0	0	0
4097-8192	0	0	0	0	0	0	0	0
8193-16384	0	0	0	0	0	0	0	0
16385-32768	0	0	0	0	0	0	0	0
32769-65536	0	0	0	0	0	0	0	0

Ingress Slot ID: 5

1/3/1 Congestion History

Time Range: 2024-12-06 02:59:19 to 2024-12-06 10:59:19 (UTC)

Samples per queue: 750

Depth (KB)	Count at Depth							Q7
	Q0	Q1	Q2	Q3	Q4	Q5	Q6	
0	250	250	250	250	250	250	250	249
1-64	0	0	0	0	0	0	0	1
65-128	0	0	0	0	0	0	0	0
129-256	0	0	0	0	0	0	0	0
257-512	0	0	0	0	0	0	0	0
513-1024	0	0	0	0	0	0	0	0
1025-2048	0	0	0	0	0	0	0	0
2049-4096	0	0	0	0	0	0	0	0
4097-8192	0	0	0	0	0	0	0	0
8193-16384	0	0	0	0	0	0	0	0
16385-32768	0	0	0	0	0	0	0	0
32769-65536	0	0	0	0	0	0	0	0

Related Commands

Command	Description
queue-monitor	This command enables the queue monitoring feature, allowing the switch to collect queue statistics at 10-second time intervals. These aggregate statistics can be used to monitor network health, troubleshoot network issues, and identify normal network behavior patterns and performance anomalies. This feature is disabled by default, but can be enabled on a maximum of 52 interfaces per switch.

Command History

Release	Modification
10.17	Command introduced on 5420, 6200, 6300, and 6400v2 Switch Series

Command Information

Platforms	Command context	Authority
6300M/F (except for JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL669A, JL670A, JL762A, JL761A, JL760A) 6400v2	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 queue-monitor status

```
show interface {<IFNAME> | <IFRANGE>} queue-monitor status [vsx-peer]
```

Description

Displays the queue monitoring status for a specific interface or range of interfaces. If no interface is specified, the command will display the queue monitoring status for all interfaces configured.

Parameter	Description
{<IFNAME> <IFRANGE>}	Use the queue-monitor status parameters to show whether the feature is enabled or disabled, and information about memory consumption and monitored interfaces. Include

Parameter	Description
	the optional <IFNAME> parameter to filter the output of this command to show data for just the specified interface, or include both the <IFNAME> parameter and the additional <IFRANGE> parameter also to view data for a range of interfaces. If no interface is specified, the command will display the queue monitor status for all interfaces that have queue monitoring configured.
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Examples

Showing queue monitoring status for interface **1/1/1** when queue monitoring is not configured on that interface.

```
switch# show interface 1/1/1 queue-monitor status
Interface 1/1/1 does not have queue monitoring enabled.
```

Showing the queue monitoring status for interface **1/1/1** when queue monitoring is enabled on that interface.

```
switch# show interface 1/1/1 queue-monitor status
Interface 1/1/1

Status                : Enabled
Monitored Statistics  : queue-depth
```

Showing the queue monitoring status for a range of interfaces when queue monitoring is enabled on those interfaces.

```
switch# show interface 1/1/1-1/1/2 queue-monitor status
Interface 1/1/1

Status                : Enabled
Monitored Statistics  : queue-depth

Interface 1/1/2

Status                : Enabled
Monitored Statistics  : queue-depth
```

Showing the queue monitoring status for all interfaces that have queue monitoring enabled, but the interface does not support any statistics.

```
switch# show interface queue-monitor status
Interface 1/1/1
```

```
Status                : Blocked (No queue statistics supported for monitoring)

Interface 1/1/2

Status                : Blocked (No queue statistics supported for monitoring)

Interface 1/1/3 does not have queue monitoring enabled.
```

Showing the queue monitoring status for all interfaces where queue monitoring is enabled on interface **1/1/1**.

```
switch# show interface queue-monitor status
Interface 1/1/1

Status                : Enabled
Monitored Statistics  : queue-depth

Interface 1/1/2 does not have queue monitoring enabled.

Interface 1/1/3 does not have queue monitoring enabled.
```

Showing the queue monitoring status for all interfaces when the feature pack is invalid and queue monitoring is configured only on interface **1/1/1**.

```
switch# show interface queue-monitor status
Interface 1/1/1

Status                : Blocked (Invalid feature pack)

Interface 1/1/2 does not have queue monitoring enabled.

Interface 1/1/3 does not have queue monitoring enabled.
```

Showing the queue monitoring status for interface **1/1/1** when the feature pack is invalid and queue monitoring is configured on that interface.

```
switch# show interface 1/1/1 queue-monitor status
Interface 1/1/1

Status                : Blocked (Invalid feature pack)
```

Command History

Release	Modification
10.17	Command introduced

Command Information

Platforms	Command context	Authority
6300M/F (except for JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL669A, JL670A, JL762A, JL761A, JL760A) 6400v2	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 queue-monitor status

show queue-monitor status

Description

Displays the global state of the queue monitoring feature. The output includes the global enable/disable status, the memory consumed by the feature, the statistics being collected, and the interfaces currently enabled for monitoring.

Examples

Showing queue monitoring status when queue monitoring is not enabled on any interfaces.

```
switch# show queue-monitor status
Feature Status           : Running
Memory Consumption       : 0 kbytes
Monitored Interfaces/Max : 0/52
```

Showing the queue monitoring status when queue monitoring is enabled on interface 1/1/1.

```
switch# show queue-monitor status
Feature Status           : Running
Memory Consumption       : 184 kbytes
Polling Interval         : 10 seconds
Data Retention Duration  : 8 hours
Monitored Interfaces/Max : 1/52
Monitored Interfaces     : 1/1/1
```

Related Commands

Command	Description
queue-monitor	This command enables the queue monitoring feature, allowing the switch to collect queue statistics at 10-second time intervals. These aggregate statistics can be used to monitor network health, troubleshoot network issues, and identify normal network behavior patterns and performance anomalies. This feature is disabled by default, but can be enabled on a maximum of 52 interfaces per switch.

Command History

Release	Modification
10.17	Command introduced on 5420, 6200, 6300, and 6400v2 Switch Series

Command Information

Platforms	Command context	Authority
6300M/F (except for JL658A, JL659A, JL660A, JL661A, JL662A, JL663A, JL664A, JL665A, JL666A, JL667A, JL668A, JL669A, JL670A, JL762A, JL761A, JL760A) 6400v2	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.

IP Flow Path Trace

IP Flow Path Trace (flowtraced) feature helps to trace the complete path taken by an Application Flow while traversing from Source device to Destination device. There is a new dedicated daemon (flowtraced) created to support the feature. A new proprietary UDP protocol is designed to handle the feature requirements. These CPU generated packets traverse through the network and collect all the necessary telemetry information from each of the HPE ANW OS-CX devices traversed where flow path trace is supported. All the information collected from each of the device traversed is stored in the payload as TLVs.



IP Flow Path Trace is supported on 5420, 6200, 6300, 6400, 8360 and 8100 Switch Series.

Limitations, Conflicts or Exclusions

- Platform should support Forwarding Info infra to fetch the nexthop and interface details.
- Traffic Insight feature needs to be running on all devices in the path to match on the Policy/ACL drop matches. If Traffic Insight is not supported on the devices, then the Flow drop will not be detected.
- If Traffic Insight feature is not supported or enabled along with Application recognition, Flow Tracking, IPFix, TI with AppFlow monitor on the first-hop device, then the flow trace query should have the complete 5 tuple information. Partial query having only the Client IP and destination (AppName, DomainName) will result in lookup failure.
- Client-Insight ARPtoGW needs to be enabled on the first Access switch if acting as an L2 device to make sure that the MAC-IP association of the GW are available.
- Client-IP-Tracker needs to be enabled on the first Access switch if acting as an L2 device to make sure that the MAC-IP association of the Client/Source are available.
- If the Clients have static IPs, the **gateway-ip** needs to be provided in the flow trace query.
- The feature sends out proprietary packets using UDP Source L4 Port set to 55000 and Destination L4 Port set to 55001, which are private ports as per IANA.
- In case of Flow Trace capability unaware devices present on the path, the pathtrace post that device could result in an incorrect path.
- If the last-hop device IP is not provided along with the 5-tuple query, then packet would get leaked to the destination device if the last hop switch is not Flow Trace capable.
- The MAC and Routing tables may be constantly changing. The egress port received is valid for that snapshot in time. It may not remain the same in the future.
- This feature is not applicable for broadcast, multicast or unknown unicast packets.
- IPv6 support is not provided.
- Below drops are not captured in current release:
 - GBP Policy drops
 - QoS drops
 - Per-Hop latency
 - Path Latency
 - GBP relay issues
 - PBR drops
- Support for UBT, IPSEC or GRE tunnels is not present.
- Periodic probes for the 5-tuple query is not supported.
- Number of non-flowtrace capable devices in the path (applicable for L3 only).
- Automatic Reverse Flow Trace is not supported.
- If Application provided is not active in TI for the provided source IP, then the flow trace query will fail.
- If destination details are not provided, then the flow trace query will fail.
- Uses IP exception CoPP class which can be updated while executing the **flowtrace** query.
- If the ARP table in case of L3 or the CIPT/CIARPGW table in case of L2 does not have the MAC-IP binding on the first-hop device, then the query is not initiated and failed.
- LLDP neighbor information is used to figure out if the adjacent device is a Flow Trace capable or not. So all the devices in the network need to have LLDP enabled to make sure that the device classification is correct.
- If the TLV packet payload increases above the MTU of the egress interface, the flow trace will be stopped on that device and a partial trace information will be sent back to the initiating device. Each node will add multiple TLV information depending on the attributes getting collected.

- It is advisable to have the MTU in the network to be jumbo capable to make sure the flow trace is complete especially when there are more than 7-8 nodes between the source and destination device.

IP Flow Path Trace commands

IP Flow Path Trace

```
flowtrace {source-ip-address <IP-ADDR>} {destination-type <ipv4|domainname|appname>
destination <string>}
[source-l4-port <L4-PORT> destination-l4-port <L4-PORT>] [transport-protocol <PROTOCOL>]
[vrf <VRF-NAME>] [last-hop-device-ip <IP-ADDR>] [gateway-ip <IP-ADDR>]
```

Description

This command is used to get the complete flow path trace of the 5 tuple query which is initiated from the CLI. The Source IP and the Destination details are mandatory for the command to initiate the query. The other fields are optional parameters. If **last-hop-device-ip** is not provided, then the proprietary packets will traverse until the last hop, HPE ANW OS-CX, where flow path trace is supported, or until the destination device, in case the device before destination does not support flow path trace. The **gateway-ip** is also optional. This is useful especially when the first-hop device is an L2 device to get the next-hop MAC details needed for forwarding info lookup.



The flow path trace command does not validate the transport protocol of the flows and provides output using the user-specified protocol number instead of verifying the actual transport protocol.

Parameter	Description
IP-ADDR	Specify the IPv4 address (A.B.C.D). <i>Required</i>
L4-PORT	Specify the L4 Port (1-65535). <i>Optional</i>
PROTOCOL	Specify the transport protocol number (1-255) <i>Optional</i>
VRF-NAME	Specify the VRF. <i>Optional</i>
DOMAIN-NAME	Specify the Domain Name. <i>Optional</i>
APP-NAME	Specify the Application Name. <i>Optional</i>

Examples

This example shows the complete Flow Path trace successful until the last-hop-device:

```
switch# flowtrace source-ip-address 50.0.0.5 destination-type ipv4
destination 80.0.0.2 transport-protocol 6 source-l4-port 23456
destination-l4-port 34567 vrf default last-hop-device-ip 60.0.0.1
```

IP Flow Trace Summary:

src_ip : 50.0.0.5
src_port : 23456
protocol : 6

dst_ip : 80.0.0.2
dst_port : 34567

PathNode 1 :

Device Type : FlowTrace_Capable
System MAC : 90:aa:aa:aa:aa:80
Ingress Interface : 1/1/43
Ingress VLAN : 10
Flow Status : Forwarded

Device ID : SGXXXXXXAL
Egress Subnet : NA
Egress Interface : 1/1/45
Egress VLAN : 10

PathNode 2 :

Device Type : FlowTrace_Capable
System MAC : f8:bb:bb:bb:bb:00
70.0.0.0/24
Ingress Interface : 1/3/47
Ingress VLAN : 10
Flow Status : Forwarded

Device ID : SGXXXXXXBL
Egress Subnet : ECMP
Egress Interface : 1/5/47
Egress VLAN : 20

PathNode 3 :

Device Type : FlowTrace_Capable
System MAC : f8:cc:cc:cc:cc:00
Ingress Interface : 1/6/3
Ingress VLAN : 20
Flow Status : Forwarded

Device ID : SGXXXXXXCL
Egress Subnet : 60.0.0.0/24
Egress Interface : lag20->1/3/2
Egress VLAN : 30

PathNode 4 :

Device Type : Last_Hop_Device
System MAC : 90:dd:dd:dd:dd:00
Ingress Interface : 1/2/2
Ingress VLAN : 30
Flow Status : Forwarded

Device ID : SGXXXXXXDL
Egress Subnet : 80.0.0.0/24
Egress Interface : 1/1/1
Egress VLAN : 40

In this example the flow Path trace failed because of a Policy drop in an intermediate device:

switch# **flowtrace source-ip-address 50.0.0.5 destination-type ipv4
destination 80.0.0.2 transport-protocol 6 source-l4-port 23456
destination-l4-port 34567 vrf default last-hop-device-ip 60.0.0.1**

IP Flow Trace Summary:

src_ip : 50.0.0.5
src_port : 23456
protocol : 6

dst_ip : 80.0.0.2
dst_port : 34567

PathNode 1 :

Device Type : FlowTrace_Capable
System MAC : 90:aa:aa:aa:aa:80
Ingress Interface : 1/1/43
Ingress VLAN : 10
Flow Status : Forwarded

Device ID : SGXXXXXXAL
Egress Subnet : NA
Egress Interface : 1/1/45
Egress VLAN : 10

PathNode 2 :

Device Type : FlowTrace_Capable
System MAC : f8:bb:bb:bb:bb:00
70.0.0.0/24
Ingress Interface : 1/3/47
Ingress VLAN : 10
Flow Status : Forwarded

Device ID : SGXXXXXXBL
Egress Subnet : ECMP
Egress Interface : 1/5/47
Egress VLAN : 20

PathNode 3 :

Device Type : FlowTrace_Capable
System MAC : f8:cc:cc:cc:cc:00
70.0.0.0/24
Ingress Interface : 1/3/47

Device ID : SGXXXXXXCL
Egress Subnet : ECMP
Egress Interface : 1/5/47

```

Ingress VLAN      : 10                               Egress VLAN      : 20
Flow Status       : Dropped
Drop Reason:
Policy : testPolicy1
Class  : testClass1
RuleID : 30

```

This example shows Flow Path trace via an intermediate Non-Flow trace capable device.

```

switch# flowtrace source-ip-address 50.0.0.5 destination-type ipv4
destination 80.0.0.2 transport-protocol 6 source-l4-port 23456
destination-l4-port 34567 vrf default last-hop-device-ip 60.0.0.1
IP Flow Trace Summary:
-----
src_ip      : 50.0.0.5                               dst_ip      : 80.0.0.2
src_port    : 23456                                   dst_port    : 34567
protocol    : 6
-----
PathNode 1 :
Device Type      : FlowTrace_Capable                 Device ID    : SGXXXXXXAL
System MAC       : 90:aa:aa:aa:aa:80                 Egress Subnet : NA
Ingress Interface : 1/1/43                           Egress Interface : 1/1/45
Ingress VLAN     : 10                               Egress VLAN   : 10
Flow Status      : Forwarded
PathNode 2 :
Device Type      : FlowTrace_Capable                 Device ID    : SGXXXXXXBL
System MAC       : f8:bb:bb:bb:bb:00                 Egress Subnet : ECMP
70.0.0.0/24
Ingress Interface : 1/3/47                           Egress Interface : 1/5/47
Ingress VLAN     : 10                               Egress VLAN   : 20
Flow Status      : Forwarded
PathNode 3 :
Device Type      : NonFlowTrace_Capable              Device ID    : SGXXXXXXCL
System MAC       : f8:cc:cc:cc:cc:00
Flow Status      : Forwarded
PathNode 4 :
Device Type      : Last_Hop_Device                   Device ID    : SGXXXXXXDL
System MAC       : 90:dd:dd:dd:dd:00                 Egress Subnet : 80.0.0.0/24
Ingress Interface : 1/2/2                             Egress Interface : 1/1/1
Ingress VLAN     : 30                               Egress VLAN   : 40
Flow Status      : Forwarded

```

This example shows a successful complete Flow Path trace for an Application Based query:

```

switch# flowtrace source-ip-address 10.0.0.10 destination-type app_name
destination ssh
IP Flow Trace Summary:
-----
src_ip      : 10.0.0.10                               dst_ip      : 60.0.0.2
src_port    : 23456                                   dst_port    : 22
protocol    : 6                                       app_name    : ssh
-----
PathNode 1 :
Device Type      : FlowTrace_Capable                 Device ID    : SGXXXXXXAL
System MAC       : 90:ea:ea:ee:aa:50                 Egress Subnet : NA
Ingress Interface : 1/1/1                             Egress Interface : 1/1/48
Ingress VLAN     : 10                               Egress VLAN   : 10
Flow Status      : Forwarded
PathNode 2 :

```

```

Device Type      : FlowTrace_Capable      Device ID       : SGXXXXXXBXL
System MAC       : f8:bb:bb:bb:bb:00      Egress Subnet   : 70.0.0.0/24
Ingress Interface : 1/1/47                Egress Interface : 1/5/47
Ingress VLAN     : 10                    Egress VLAN     : 20
Flow Status      : Forwarded

PathNode 3 :
Device Type      : Last_Hop_Device        Device ID       : SGXXXXXXCL
System MAC       : f8:cc:cc:cc:cc:00      Egress Subnet   : 60.0.0.0/24
Ingress Interface : 1/6/3                Egress Interface : 1/6/6
Ingress VLAN     : 20                    Egress VLAN     : 30
Flow Status      : Forwarded

```

This example shows a successful complete Flow Path trace for a VxLAN based topology:

```

switch# flowtrace source-ip-address 10.0.0.10 destination-type app_name
destination ssh
IP Flow Trace summary:
-----
src_ip      : 10.0.0.10                dst_ip      : 60.0.0.2
src_port    : 23456                    dst_port    : 22
protocol    : 6
-----

PathNode 1 :
Device Type      : FlowTrace_Capable      Device ID       : SGXXXXXXAL
System MAC       : 90:aa:aa:aa:aa:80      Egress Subnet   : NA
Ingress Interface : 1/1/43                Egress Interface : 1/1/45
Ingress VLAN     : 10                    Egress VLAN     : 10
Flow Status      : Forwarded

PathNode 2:
Device Type      : Tunnel_Entry_Device    Device ID       : SGXXXXXXBXL
System MAC       : BB:A4:7D:29:C4:50      Egress IP       : 30.0.0.1
Ingress Interface : 1/1/10                Egress Interface : 1/1/12
Egress VLAN     : 30                    Flow Status     : Forwarded

PathNode 3:
Device Type      : Flow_Trace Capable     Device ID       : SGXXXXXXCL
System MAC       : CC:A4:7D:29:C4:50      Egress IP       : 40.0.0.1
Ingress Interface : 1/1/13                Egress Interface : 1/1/15
Egress VLAN     : 40                    Flow Status     : Forwarded

PathNode 4:
Device Type      : Tunnel_End_Device      Device ID       : SGXXXXXXDL
System MAC       : DD:A4:7D:29:C4:50      Egress IP       : 50.0.0.1
Ingress Interface : 1/1/4                Egress Interface : 1/1/5
Egress VLAN     : 30                    Flow Status     : Forwarded

PathNode 5:
Device Type      : Last_Hop_Device        Device ID       : SGXXXXXXEL
System MAC       : EE:A4:7D:29:C4:60      Egress IP       : 60.0.0.1
Ingress Interface : 1/1/20                Egress Interface : 1/1/21
Egress VLAN     : 50                    Flow Status     : Forwarded

```

This example shows an error due to a missing Source IP:

```

switch# flowtrace destination-type ipv4 destination 50.0.0.50
transport-protocol 6 source-l4-port 23456 destination-l4-port 34567
vrf default last-hop-device-ip 50.0.0.1

Missing Required Parameter Source IP

```

This example shows an error since there is no active Application:

```
switch# flowtrace source-ip-address 100.0.0.10 destination-type appname
destination ssh
```

```
No active session to the 'ssh' application from source 100.0.0.10
```

This example shows an error since Forwarding Info infra is not successful

```
switch# flowtrace source-ip-address 10.0.0.10 destination-type ipv4
destination 50.0.0.50 transport-protocol 6 source-l4-port 23456
destination-l4-port 34567 vrf default last-hop-device-ip 50.0.0.1
```

```
Forwarding-info lookup was not successful
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

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

gRPC network management interface

gRPC is a RPC framework developed by Google to create distributed systems. This protocol allows a client running on one system to call the services defined in a remote server as a local server.

The gRPC Network Management Interface (gNMI) feature provides secure, high-performance, standards-based network management and gRPC-based access to network device configuration and operational data. Using industry-standard OpenConfig models, gNMI enables the building of modern network automation and monitoring solutions.

gNMI is implemented as an HTTP/2 server supporting encrypted communication and role-based access control. It allows network management systems to interact with the device for configuration and telemetry using gRPC, supporting VRF-aware access and role-based control.

gNMI supports access over SVIs, loopback, and data ports.

gNMI allows for the following:

- Streaming of real-time telemetry data from switches
- Monitoring of interface statistics, system performance, and hardware health
- Integration with modern network management platforms
- Implementation of event-driven network automation

Refer to [gNMI commands](#) for information on enabling and configuring gNMI.

gNMI prerequisites and setup

Before implementing gNMI, certain network, authentication, and software requirements must be met.

Network requirements:

- Network connectivity between the gNMI client and AOS-CX switches
- IPv4 or IPv6 connectivity (both are supported)
- Access to TCP port 9339 on system switches

Authentication requirements:

- Valid user credentials (local users or RADIUS/TACACS+ authentication)
- TLS 1.2 or 1.3 support on client systems
- Certificate management for secure connections

Client software:

- gNMI client software (such as gNMIC, which is used in this guide)
- JSON parsing capabilities in monitoring or automation tools

Before implementing gNMI, it is recommended to configure gNMI certificates using **crypto pki application gnmi certificate <certificate-name>**. This command assigns a specific certificate for gNMI to use instead of the default local certificate. The certificate must be instanced prior to using this command. Refer to [the Security Guide](#) for more information on certificates. An example is as follows:

```
switch(config)# crypto pki certificate my-gnmi-cert  
switch(config-cert-my-gnmi-cert)# crypto pki application gnmi certificate my-gnmi-cert
```

gNMI capabilities, limitations, and best practices

gNMI has the following capabilities:

- Monitor up to 30 data paths per subscription request
- Maintain up to 9 concurrent data streams per switch
- Enable gNMI on multiple VRFs simultaneously
- Use read-only access to device data
- Choose from multiple subscription modes (on-change, sample, once)
- Set sample intervals from 10 seconds to 30 minutes
- Monitor interface statistics, system performance, and hardware status
- Supports access over SVIs, loopback, and data ports

gNMI has the following limitations:

- Read-only access (configuration changes are not supported via gNMI)
- Subinterfaces and VSF links are not supported
- XPath wildcards (* or ...) are not supported
- Sample mode requires the suppress-redundant flag
- gNMI streaming connections have a 24-hour session timeout. After this period, the session is automatically terminated and clients must reconnect. This timeout is not user-configurable.

The following best practices are suggested for secure gNMI implementation:

- Use proper certificates. Configure CA-signed certificates instead of self-signed certificates.
- Avoid bypassing validation. Do not use certificate validation bypass flags in production.
- Dedicated management networks. Use dedicated management VRFs when possible.
- Proper authentication. Implement strong authentication mechanisms.
- Network segmentation. Restrict gNMI access to authorized management networks.

The following best practices are suggested for attribute selection:

- Use local authentication for initial setup and testing
- Implement RADIUS/TACACS+ authentication for production environments
- Ensure proper user permissions for gNMI access
- Start with general paths for discovery and initial testing
- Use specific paths in production to minimize bandwidth and processing overhead
- Combine multiple specific paths (up to 30) for efficient data collection
- Always use secure connections in production environments
- Use proper CA-signed certificates instead of self-signed certificates
- Configure hostname-based connections for better security validation
- Choose appropriate subscription modes based on data update requirements
- Use sample intervals that balance timeliness with system performance
- Monitor concurrent stream usage (maximum 9 streams per switch)

Getting started with gNMI

Use the following steps to configure gNMI.

1. Enable gNMI on the switch using the following configuration:

```
switch(config)# gnmi vrf default
```

Verify the configuration using the following:

```
switch(config)# show gnmi
gNMI Configuration
-----
-
VRF : default
Access mode : read-only
Global stream limit : 9
```

2. Install a gNMI client that supports the gNMI specification. The example below uses **gnmic**, which can be downloaded from the official repository.

```
bash -c "$(curl -sL https://get-gnmic.openconfig.net)" -- --use-pkg
```

This command automatically downloads and installs the latest **gnmic** package for the operating system.

3. Test basic connectivity.

For IPv4 connections:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-password> --skip-verify --format json capabilities --encoding json_ietf
```

For IPv6 connections:

```
gnmic -a [<your-switch-ipv6>]:9339 --username <your-username> --password <your-password> --skip-verify --format json capabilities --encoding json_ietf
```



The **--skip-verify** flag bypasses certificate validation and should only be used for testing. In production, configure proper certificates and remove this flag.

Discover available data using the following query:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-password> --skip-verify --format json capabilities --encoding json_ietf
```

This command returns information about what data models your switch supports and confirms that your connection is working properly. The following is an example response excerpt:

```
{
  "version": "0.10.0",
  "supported-models": [
    {
      "name": "openconfig-interfaces",
      "organization": "OpenConfig",
      "version": "2024-04-04"
    },
    {
      "name": "openconfig-system",
      "organization": "OpenConfig",
      "version": "2024-09-24"
    }
  ],
  "encodings": ["JSON_IETF"]
}
```

gNMI connection types and common attributes

Understanding the common attributes and connection types is essential for effectively using gNMI with AOS-CX switches.

Connection types

An unsecure connection is a connection without certificate validation between device and client. This should be used for testing purposes only, such as in laboratory and testing environments, at initial setup and troubleshooting, or for proof-of-concept implementations.



Never use unsecure connections in production environments due to security risks.

A secure connection is a connection with proper certificate validation between device and client and is recommended for production environments. Use a secure connection in production environments, in any network with security requirements, or in long-term monitoring implementations.

Common attributes

The following are basic connection parameters for gNMI:

- Device IP/Hostname: Communication IP address (IPv4 or IPv6) or configured hostname from the device to the client
- Port: 9339 (standard gNMI port, not configurable)
- Username: User credentials for authentication
 - Local user: A user defined directly on the switch
 - Remote user: A user defined in a remote authentication server (RADIUS/TACACS+)
- Password: The password assigned to the specified user

The following are data path parameters for gNMI:

- Paths: OpenConfig module paths (up to 30 paths per subscription request)
 - General Paths: Broad data collection
 - openconfig:/system - Complete system information
 - openconfig:/interfaces - All interface data
 - openconfig:/components - All hardware component data
 - Specific Paths: Targeted data collection (there are many more specific paths than specified below)
 - openconfig:/system/state/hostname - System hostname only
 - openconfig:/interfaces/interface[name="<interface-id>"] - Specific interface
 - openconfig:/components/component[name="<component-name>"] - Specific component

The following are subscription mode parameters for gNMI:

- Mode: Subscription behavior
 - once - Single data retrieval, then close connection
 - stream - Continuous streaming connection
- Stream Mode: (when using stream mode)
 - on-change - Immediate updates when values change
 - sample - Periodic updates at specified intervals
 - target-defined - Same behavior as on-change mode
- Sample Interval: (for sample mode only)
 - Minimum: 10 seconds
 - Maximum: 1800 seconds (30 minutes)
- Suppress Redundant: Mandatory flag for streaming mode sample.

The following are security parameters. These apply only to secure connections.

- TLS Certificate Authority (--tls-ca): Path to CA certificate file for validating server certificates
- TLS Certificate (--tls-cert): Path to client certificate file for mutual authentication
- TLS Version (--tls-version): Specify TLS protocol version
 - Recommended: 1.2 or 1.3
- Hostname Verification: Enabled by default with secure connections

The following are output and debugging parameters:

- Format: Output format (json recommended for readability)
- Encoding: Data encoding (json_ietf recommended)
- Debug (-d): Enable debug output for troubleshooting
- Skip Verify (--skip-verify): Bypass certificate validation (unsecure connections only)

Supported OpenConfig models

AOS-CX switches support the OpenConfig models listed in [Primary model support](#) for gNMI access. Use this information to understand what data can be monitored and collected from AOS-CX switches.

Table 1: *Primary model support*

Model Category	OpenConfig Module	Version	Description
Interfaces	openconfig-interfaces	3.7.1	Interface configuration and state
System	openconfig-system	2.3.0	System information and resources
Platform	openconfig-platform	0.30.0	Hardware components and sensors

AOS-CX switches support specific interface, system, hardware platform, and supporting OpenConfig models. For a full list of supported leaves, refer to the [HPE ANW github repo](#).



Use the capabilities query shown in [Getting started with gNMI](#) for the most up-to-date list of supported data paths within these models.

Interface-related models:

- openconfig-interfaces (2024-04-04) - Core interface management
- openconfig-if-ethernet (2024-09-17) - Ethernet-specific features
- openconfig-if-aggregate (2022-06-28) - Link aggregation
- openconfig-if-poe (2018-11-21) - Power over Ethernet
- openconfig-if-ethernet-ext (2018-11-21) - Extended ethernet features

System models:

- openconfig-system (2024-09-24) - System configuration and state
- openconfig-types (2024-01-31) - Common data types
- openconfig-yang-types (2024-05-30) - YANG data types
- openconfig-inet-types (2024-01-05) - Internet address types

Hardware platform models:

- openconfig-platform (2024-10-13) - Hardware platform information
- openconfig-platform-fan (2018-11-21) - Fan monitoring
- openconfig-platform-cpu (2018-11-21) - CPU monitoring
- openconfig-platform-psu (2018-11-21) - Power supply monitoring
- openconfig-platform-types (2024-11-04) - Platform type definitions

Supporting models:

- openconfig-extensions (2024-09-19) - OpenConfig extensions
- openconfig-alarm-types (2018-11-21) - Alarm type definitions
- openconfig-transport-types (2024-11-21) - Transport layer types
- ietf-interfaces (2018-02-20) - IETF interface standard
- iana-if-type (2017-01-19) - IANA interface types

gNMI subscription modes

gNMI offers different subscription modes to fit various monitoring needs. Choose the right mode based on how frequently data is needed and whether to collect all data or just change detection.

Sample mode

Sample mode is used to collect data at regular intervals with change detection. Use the following configuration to start a stream of sampled data

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert  
<tls-certificate-file-name> --path openconfig:/system --mode sample --sample-  
interval 60s -suppress redundant --encoding json_ietf --tls-version 1.2 -d
```

Use sample mode when regular updates are needed at specific intervals, to reduce bandwidth by only receiving data when values change, or when building time-series databases or trend analysis.

Sample intervals can be set from 10 seconds to 1800 seconds (30 minutes). The **--suppress-redundant** flag is required and ensures updates only when values change. It is recommended to start at 60.

On-change mode

On-change mode is used to receive immediate updates when monitored values change. Use the following configuration to start a stream where changes are immediately reported:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert  
<tls-certificate-file-name> --path openconfig:/system --mode on-change --encoding  
json_ietf --tls-version 1.2 -d
```

Use on-change mode to get immediate notification of changes, to implement event-driven automation, to minimize bandwidth usage, and to monitor for configuration changes or alarm conditions.

Once mode

Once mode is used for one-time data collection. Use the following configuration to enter into once mode:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert  
<tls-certificate-file-name> --path openconfig:/system --mode once --encoding json_  
ietf --tls-version 1.2 -d
```

Use once mode when building polling-based monitoring systems, to get data snapshots for reports, to control exactly when data is collected, and when implementing custom scheduling logic.

gNMI secure connectivity

gNMI over HTTP/2 supports secure connections over TLS using server certificates at the AOS-CX device. The default certificate is **local-cert**. To assign a certificate, use **crypto pki application gnmi-certificate <CERT-NAME>**. If a specific certificate is not set, the default is used. It is recommended to replace with a CA-signed certificate where possible as the default may not be valid for all deployments.



mTLS is not supported.

To verify the certificate, use **show crypto pki application**. The output of this command indicates which certificate is active for gNMI.

gNMI use case 1: interface monitoring

This section covers monitoring network interfaces on a switch. Interface monitoring is the most common use case for gNMI.

Monitoring a complete interface

To get comprehensive information about a specific interface, including configuration, operational state, and statistics, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path  
openconfig:/interfaces/interface[name=1/1/1] --encoding json_ietf
```

Monitoring a complete interface is ideal for complete visibility into an interface's status, including PoE information, speed negotiation, and traffic counters. The response includes interface configuration, operational status, ethernet-specific information, PoE data (if applicable), and all traffic counters.

Monitoring all interface counters

To collect traffic statistics from all interfaces simultaneously, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path  
openconfig:/interfaces/interface/state/counters --encoding json_ietf
```

Monitoring all interface counters is ideal for network-wide traffic monitoring, creating dashboards that show traffic across the entire switch, or feeding data into network monitoring systems. The response includes traffic counters for every interface on the switch, including packet counts, byte counts, error counts, and status change counters.

Monitoring specific interface metrics

To monitor just one specific metric, such as incoming packet count on a particular interface, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path  
openconfig:/interfaces/interface[name=1/1/1]/state/counters/in-pkts --encoding  
json_ietf
```

Monitoring specific interface metrics is ideal for focused monitoring of specific metrics for alerting, trending, or troubleshooting specific interface issues. The response includes only the specific counter you requested, reducing bandwidth and processing overhead.

gNMI use case 2: system monitoring

This section covers monitoring a switch's system resources and operational status.

Monitoring complete system status

To get comprehensive system information including CPU, memory, software version, and configuration, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path openconfig:/system --  
encoding json_ietf
```

Monitoring complete system status is ideal for system health dashboards, compliance reporting, or comprehensive device status checks. The response includes hostname, software version, memory usage, CPU utilization, timezone, and last configuration change timestamp.

Monitoring memory usage

To track memory utilization for capacity planning, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path  
openconfig:/system/memory/state/used --encoding json_ietf
```

Monitoring memory usage is ideal for capacity planning, alerting on high memory usage, or trending memory consumption over time.

Monitoring CPU utilization

To monitor real-time CPU performance, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path  
openconfig:/system/cpus/cpu/state/total/instant --encoding json_ietf
```

Monitoring CPU utilization is ideal for performance monitoring, identifying high CPU utilization periods, or capacity planning.

gNMI use case 3: platform and hardware monitoring

This section covers monitoring the physical health of switches.

Monitoring complete hardware inventory

To get detailed information about all hardware components, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path openconfig:/components --  
encoding json_ietf
```

Monitoring complete hardware inventory is ideal for asset management, hardware health monitoring, environmental monitoring, or generating hardware inventory reports. The response includes information about chassis, power supplies, fans, temperature sensors, management modules, and line cards.

Monitoring hardware component discovery

To get a list of all hardware component names, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path  
openconfig:/components/component/name --encoding json_ietf
```

Monitoring hardware component discovery is ideal for automated discovery of hardware components or building inventory management systems.

Monitoring fan speed

To monitor cooling system performance, use the following:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path  
openconfig:/components/component/fan/state/speed --encoding json_ietf
```

Monitoring fan speed is ideal for environmental monitoring, predictive maintenance, or ensuring proper cooling performance.

gNMI troubleshooting

gNMI client cannot connect

Use the following steps if the gNMI client cannot connect and there is no response:

1. Ensure gNMI is enabled on the correct VRF.

```
switch# show gnmi  
switch(config)# gnmi vrf <vrf-name>
```

2. Check network connectivity, credentials, and logs.

```
telnet <switch-ip> 9339
```

```
ping <switch-ip>
```

3. Verify username and password.

```
switch# show user <username>  
switch# show aaa authentication
```

4. Check log for errors.

```
switch# show logging -d nginx-configurator  
switch# show logging -d yang-resolverd
```

Use the following best practices for certificate and TLS issues:

- If certificate or TLS verification is not skipped, ensure the gNMI server is configured with a valid certificate signed by a trusted authority.
- The client must trust the server's certificate, otherwise the connection may be refused or fail its TLS handshake.
- Check that the certificate's Common Name (CN) or Subject Alternative Name (SAN) matches the server address used by the client.
- If using self-signed certificates, import the server certificate into the client's trusted certificate store.
- Review client-side logs for TLS or certificate errors, such as **certificate not trusted** or **hostname mismatch**.
- On the switch, use **show crypto pki application** to verify the assigned certificate for gNMI.
- Replace the default local-cert with a valid certificate if needed.

Certificate validation issues

Use the following steps if the gNMI client has certificate validation errors, if there is a TLS handshake failure, or a hostname mismatch error:

1. Verify certificate configuration.

```
switch# show crypto pki application
```

2. Check the certificate Subject Alternative Names (SANs) include client connection IP/hostname.
3. Ensure client certificate store includes the CA that signed the server certificate.
4. For self-signed certificates, import the server certificate into client's trusted store.

Service recovery

Use the following steps for service recovery if gNMI appears non-responsive:

1. Restart gNMI services by disabling and re-enabling.

```
switch(config)# no gnmi vrf <vrf-name>  
switch(config)# gnmi vrf <vrf-name>
```

2. Check service status using diagnostic commands.

```
switch# show tech gnmi
switch(config)# diag-dump gnmi basic
```

Troubleshooting caveats

Consider the following caveats when troubleshooting:

- gNMI requests are rate-limited with a 1-second throttling interval between subscription requests. Wait at least 1 second between consecutive subscription requests to avoid "Request throttled. Try again later" errors. Multiple paths within a single request are processed together.
- The **--suppress-redundant** flag is mandatory for all sample mode subscriptions. Sample mode will only send notifications when values change, not periodic updates, making it functionally similar to on-change mode.
- On Core-Aggregation platforms (8xxx, 9xxx, and 10k switch series), former LAG member interfaces may not generate gNMI notifications after LAG deletion due to default shutdown state. Configure the interface to restore gNMI visibility.
- When moving interfaces between VRFs, gNMI streams may remain active on the previous VRF for up to 15 minutes, counting against the 9 concurrent stream limit. Disable gNMI on the previous VRF to immediately close streams.

gNMI diagnostics and monitoring

Event logging

AOS-CX provides comprehensive logging for gNMI operations. Key events include the following:

- Connection events
 - gNMI connection open and closed
 - Connection rejections and reasons
 - Authentication failures
- Configuration events
 - VRF enable and disable changes
 - Certificate configuration changes

To view gNMI-related logs, using the following:

```
# View nginx-configurator events
switch# show logging -d nginx-configurator
switch# show events -d nginx-configurator
# View yang-resolverd events
switch# show logging -d yang-resolverd
switch# show events -d yang-resolverd
# View authentication events
switch# show logging -d hpe-restd
```

Diagnostic commands

Advanced troubleshooting must be done by an administrator.

For comprehensive gNMI status, use the following:

```
switch# show tech gnmi
```

For service status dump, use the following:

```
switch(config)# diag-dump gnmi basic
```

The following is an example of a diagnostic output:

```
[Start] Feature gnmi Time : Thu Jun 26 14:30:13 2025
=====
gNMI nginx VRF service status dump:
-----
| VRF          | LoadState   | ActiveState | SubState     |
-----
| VRF_2        | loaded      | active     | running      |
-----
Diagnostic-dump captured for feature gnmi
```

gNMI event reference

The following events are logged for gNMI operations:

Event ID	Description	Severity
17101	gNMI connection opened	Informational
17102	gNMI connection closed	Informational
17103	gNMI connection rejected - maximum connections reached	Warning
17104	gNMI enabled on VRF	Informational
17105	gNMI disabled on VRF	Informational
17106	gNMI VRF reconfigured	Informational
17107	gNMI maximum paths per subscription limit reached	Warning

Use the following best practices for logging:

- Check logs periodically for connection issues or configuration warnings.
- Use specific daemon filters to focus on gNMI-related events.
- Watch for events 17103 and 17107 which indicate resource limits.
- Review authentication failures in hpe-restd logs

Certificate verification

Use the following to check the certificate assignment:

```
switch# show crypto pki application
```

The following is an example output:

Associated Applications	Certificate Name	Cert Status
gnmi	not configured, using local-cert	

REST API certification

Use the following to configure gNMI using the REST API:

- URI: /rest/v10.17/system/vrfs/{VRF_NAME}
- Method: PATCH
- Body: {"gnmi": {"enable": true}}

gNMI command syntax templates

There are two command syntax templates for gNMI.

Unsecure connection template

```
gnmic -a <device-ip-or-hostname>:9339 \
  --username <username> \
  --password <password> \
  --format json \
  --skip-verify \
  subscribe \
  --path <openconfig-path> \
  --mode <once|stream> \
  [--stream-mode <on-change|sample|target-defined>] \
  [--sample-interval <10s-1800s>] \
  [--suppress-redundant] \
  --encoding json_ietf
```

Secure connection template

```
gnmic -a <device-hostname>:9339 \
  --username <username> \
  --password <password> \
  --format json \
  --tls-ca <ca-certificate-file> \
  --tls-cert <client-certificate-file> \
  --tls-version <1.2|1.3> \
  -d \
  subscribe \
  --path <openconfig-path> \
  --mode <once|stream> \
  [--stream-mode <on-change|sample|target-defined>] \
  [--sample-interval <10s-1800s>] \
  [--suppress-redundant] \
  --encoding json_ietf
```

gNMI implementation examples

The following examples demonstrate practical gNMI usage with actual commands and response outputs.

These examples represent the different stream modes and modes.

Device capabilities discovery

Device capabilities discovery using an unsecure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-password> --skip-verify --format json capabilities --encoding json_ietf
```

Device capabilities discovery using a secure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert <tls-certificate-file-name> --tls-version 1.2 -d capabilities --encoding json_ietf
```

Example output:

```
{
  "version": "0.10.0",
  "supported-models": [
    {
      "name": "openconfig-if-ethernet",
      "organization": "OpenConfig",
      "version": "2024-09-17"
    },
    {
      "name": "iana-if-type",
      "organization": "IANA",
      "version": "2017-01-19"
    },
    {
      "name": "openconfig-system",
      "organization": "OpenConfig",
      "version": "2024-09-24"
    },
    {
      "name": "openconfig-interfaces",
      "organization": "OpenConfig",
      "version": "2024-04-04"
    },
    {
      "name": "openconfig-platform",
      "organization": "OpenConfig",
      "version": "2024-10-13"
    }
  ],
  "encodings": [
    "JSON_IETF"
  ]
}
```

Monitor specific interface

Monitor a specific interface using an unsecure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-password> --skip-verify --format json subscribe --path openconfig:/interfaces/interface[name=1/1/1] --encoding json_ietf
```

Monitor a specific interface using a secure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert  
<tls-certificate-file-name> --path openconfig:/interfaces/interface[name=1/1/1] --  
mode once --encoding json_ietf --tls-version 1.2 -d
```

Example output:

```
{  
  "source": "<your-switch-ip>:9339",  
  "subscription-name": "default-1759857080",  
  "timestamp": 1759858552012475920,  
  "time": "2025-10-07T11:35:52.01247592-06:00",  
  "updates": [  
    {  
      "Path": "openconfig:interfaces/interface[name=1/1/1]",  
      "values": {  
        "interfaces/interface": [  
          {  
            "openconfig-if-ethernet:ethernet": {  
              "openconfig-if-poe:poe": {  
                "state": {  
                  "enabled": true,  
                  "power-class": 6,  
                  "power-used": "11.055804"  
                }  
              },  
              "state": {  
                "auto-negotiate": true,  
                "counters": {  
                  "openconfig-if-ethernet-ext:in-distribution": {  
                    "in-frames-128-255-octets": "17266",  
                    "in-frames-512-1023-octets": "55246",  
                    "in-frames-65-127-octets": "7863"  
                  }  
                },  
                "enable-flow-control": false,  
                "fec-mode": "FEC_DISABLED",  
                "hw-mac-address": "64:e8:81:d4:2a:66",  
                "mac-address": "64:e8:81:d4:2a:66",  
                "negotiated-duplex-mode": "FULL",  
                "negotiated-port-speed": "SPEED_5GB"  
              }  
            },  
            "openconfig-interfaces:config": {  
              "description": "AP1",  
              "enabled": true,  
              "mtu": 1500,  
              "name": "1/1/1",  
              "type": "iana-if-type:ethernetCsmacd"  
            },  
            "openconfig-interfaces:name": "1/1/1",  
            "openconfig-interfaces:state": {  
              "admin-status": "UP",  
              "counters": {  
                "carrier-transitions": "1025",  
                "in-broadcast-pkts": "55246",  
                "in-multicast-pkts": "25129",  
                "in-octets": "34049144",  
                "in-pkts": "80375",
```



```

{
  "Path": "interface[name=1/1/1]/state/counters",
  "values": {
    "interface/state/counters": {
      "openconfig-interfaces:carrier-transitions": "1027",
      "openconfig-interfaces:in-broadcast-pkts": "55344",
      "openconfig-interfaces:in-multicast-pkts": "25175",
      "openconfig-interfaces:in-octets": "34109737",
      "openconfig-interfaces:in-pkts": "80519",
      "openconfig-interfaces:out-broadcast-pkts": "8934",
      "openconfig-interfaces:out-multicast-pkts": "256167",
      "openconfig-interfaces:out-octets": "22319007",
      "openconfig-interfaces:out-pkts": "265101",
      "openconfig-interfaces:resets": "1027"
    }
  }
},
{
  "Path": "interface[name=1/1/26]/state/counters",
  "values": {
    "interface/state/counters": {
      "openconfig-interfaces:carrier-transitions": "1",
      "openconfig-interfaces:out-multicast-pkts": "496486",
      "openconfig-interfaces:out-octets": "37623886",
      "openconfig-interfaces:out-pkts": "496486",
      "openconfig-interfaces:resets": "1"
    }
  }
}
]
}

```

Monitor specific interface metric

Monitor a specific interface metric using an unsecure connection:

```

gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --skip-verify --format json subscribe --path
openconfig:/interfaces/interface[name=1/1/1]/state/counters/in-pkts --encoding
json_ietf

```

Monitor a specific interface metric using a secure connection:

```

gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert
<tls-certificate-file-name> --path openconfig:/interfaces/interface
[name=1/1/1]/state/counters/in-pkts --mode once --encoding json_ietf --tls-version
1.2 -d

```

Example output:

```

{
  "source": "<your-switch-ip>:9339",
  "subscription-name": "default-1759857122",
  "timestamp": 1759858593881890640,
  "time": "2025-10-07T11:36:33.88189064-06:00",
  "updates": [
    {

```

```

        "Path": "openconfig:interfaces/interface[name=1/1/1]/state/counters/in-
pkts",
        "values": {
            "interfaces/interface/state/counters/in-pkts": "80382"
        }
    }
]
}

```

Complete system information

Get complete system information using an unsecure connection:

```

gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --skip-verify --format json subscribe --path openconfig:/system --
encoding json_ietf

```

Get complete system information using a secure connection:

```

gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert
<tls-certificate-file-name> --path openconfig:/system --mode once --encoding json_
ietf --tls-version 1.2 -d

```

Example output:

```

{
  "source": "<your-switch-ip>:9339",
  "subscription-name": "default-1759857501",
  "timestamp": 1759858972325622400,
  "time": "2025-10-07T11:42:52.3256224-06:00",
  "updates": [
    {
      "Path": "openconfig:system",
      "values": {
        "system": {
          "openconfig-system:clock": {
            "state": {
              "timezone-name": "UTC"
            }
          },
          "openconfig-system:cpus": {
            "cpu": [
              {
                "index": "ALL",
                "state": {
                  "index": "ALL",
                  "total": {
                    "avg": 16,
                    "instant": 9,
                    "interval": "60000000000"
                  }
                }
              }
            ]
          },
          "openconfig-system:memory": {
            "state": {

```

```

        "counters": {
            "correctable-ecc-errors": "0"
        },
        "free": "3189243904",
        "physical": "7447310336",
        "reserved": "4099072000",
        "used": "4258066432"
    }
},
"openconfig-system:state": {
    "domain-name": "apps.com",
    "hostname": "6300-VSF2-U41",
    "last-configuration-timestamp": "1759352701000000000",
    "login-banner": "",
    "software-version": "FL.10.17.1000C-209-g08df8e207061"
}
}
}
}
]
}

```

Monitor memory usage

Monitor memory usage using an unsecure connection:

```

gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --skip-verify --format json subscribe --path
openconfig:/system/memory/state/used --encoding json_ietf

```

Monitor memory usage using a secure connection:

```

gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert
<tls-certificate-file-name> --path openconfig:/system/memory/state/used --mode
once --encoding json_ietf --tls-version 1.2 -d

```

Example output:

```

{
    "source": "<your-switch-ip>:9339",
    "subscription-name": "default-1759858357",
    "timestamp": 1759859828204425880,
    "time": "2025-10-07T11:57:08.20442588-06:00",
    "updates": [
        {
            "Path": "openconfig:system/memory/state/used",
            "values": {
                "system/memory/state/used": "4261863424"
            }
        }
    ]
}

```

Monitor CPU usage

Monitor CPU usage using an unsecure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path  
openconfig:/system/cpus/cpu/state/total/instant --encoding json_ietf
```

Monitor CPU usage using a secure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert  
<tls-certificate-file-name> --path openconfig:/system/cpus/cpu/state/total/instant  
--mode once --encoding json_ietf --tls-version 1.2 -d
```

Example output:

```
{  
  "source": "<your-switch-ip>:9339",  
  "subscription-name": "default-1759858445",  
  "timestamp": 1759859916724282200,  
  "time": "2025-10-07T11:58:36.7242822-06:00",  
  "updates": [  
    {  
      "Path": "openconfig:system/cpus/cpu[index=ALL]/state/total/instant",  
      "values": {  
        "system/cpus/cpu/state/total/instant": 21  
      }  
    }  
  ]  
}
```

Complete platform information

Access complete platform information using an unsecure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --skip-verify --format json subscribe --path openconfig:/components --  
encoding json_ietf
```

Access complete platform information using a secure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-  
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert  
<tls-certificate-file-name> --path openconfig:/components --mode once --encoding  
json_ietf --tls-version 1.2 -d
```

Example output (abbreviated):

```
{  
  "source": "<your-switch-ip>:9339",  
  "subscription-name": "default-1759859341",  
  "timestamp": 1759860812583004440,  
  "time": "2025-10-07T12:13:32.58300444-06:00",  
  "updates": [  
    {  
      "Path": "openconfig:components",  
      "values": {
```

```

"components": {
  "openconfig-platform:component": [
    {
      "config": {
        "name": "chassis - 1"
      },
      "name": "chassis - 1",
      "state": {
        "base-mac-address": "64:e8:81:d4:2a:40",
        "description": "6300M 24-port HPE Smart Rate 1/2.5/5GbE Class 6
          PoE and 4-port SFP56 Switch",
        "hardware-version": "3",
        "mfg-name": "HPE ANW",
        "model-name": "6300M 24SR5 CL6 PoE 4SFP56 Swch",
        "oper-status": "openconfig-platform-types:ACTIVE",
        "part-no": "JL660A",
        "removable": false,
        "serial-no": "SG08KMZ065",
        "type": "openconfig-platform-types:CHASSIS",
        "used-power": 114
      }
    },
    {
      "config": {
        "name": "chassis - 1 : PSU - 1/1"
      },
      "name": "chassis - 1 : PSU - 1/1",
      "power-supply": {
        "state": {
          "openconfig-platform-psu:capacity": "AAAqRA==",
          "openconfig-platform-psu:input-voltage": "AADuQg==",
          "openconfig-platform-psu:output-power": "AADKQg==",
          "openconfig-platform-psu:output-voltage": "mpnvQg=="
        }
      },
      "state": {
        "model-name": "JL086A",
        "oper-status": "openconfig-platform-types:ACTIVE",
        "removable": true,
        "type": "openconfig-platform-types:POWER_SUPPLY"
      }
    }
  ]
}
}
]
}

```

List all components

List all components using an unsecure connection:

```

gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --skip-verify --format json subscribe --path
openconfig:/components/component/name --encoding json_ietf

```

List all components using a secure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert
<tls-certificate-file-name> --path openconfig:/components/component/name --mode
once --encoding json_ietf --tls-version 1.2 -d
```

Example output:

```
{
  "source": "<your-switch-ip>:9339",
  "subscription-name": "default-1759859634",
  "timestamp": 1759861105370409160,
  "time": "2025-10-07T12:18:25.37040916-06:00",
  "prefix": "openconfig:components",
  "updates": [
    {
      "Path": "component[name=fan_tray - 1/1]/name",
      "values": {
        "component/name": "fan_tray - 1/1"
      }
    },
    {
      "Path": "component[name=fan_tray - 1/1 : Fan - Tray-1/1/2]/name",
      "values": {
        "component/name": "fan_tray - 1/1 : Fan - Tray-1/1/2"
      }
    },
    {
      "Path": "component[name=management_module - 1/1]/name",
      "values": {
        "component/name": "management_module - 1/1"
      }
    },
    {
      "Path": "component[name=chassis - 1]/name",
      "values": {
        "component/name": "chassis - 1"
      }
    }
  ]
}
```

Monitor fan speed

Monitor fan speed using an unsecure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --skip-verify --format json subscribe --path
openconfig:/components/component/fan/state/speed --encoding json_ietf
```

Monitor fan speed using a secure connection:

```
gnmic -a <your-switch-ip>:9339 --username <your-username> --password <your-
password> --format json --tls-ca <tls-certificate-authority-file-name> --tls-cert
<tls-certificate-file-name> --path
openconfig:/components/component/fan/state/speed --mode once --encoding json_ietf
--tls-version 1.2 -d
```

Example output:

```
{
  "source": "<your-switch-ip>:9339",
  "subscription-name": "default-1759860011",
  "timestamp": 1759861482023930600,
  "time": "2025-10-07T12:24:42.0239306-06:00",
  "prefix": "openconfig:components",
  "updates": [
    {
      "Path": "component[name=fan_tray - 1/1 : Fan - Tray-1/1/2]/fan/state/speed",
      "values": {
        "component/fan/state/speed": 5917
      }
    },
    {
      "Path": "component[name=fan_tray - 1/1 : Fan - Tray-1/1/1]/fan/state/speed",
      "values": {
        "component/fan/state/speed": 5882
      }
    }
  ]
}
```

gNMI commands

crypto pki application gnmi certificate

crypto pki application gnmi certificate <CERT-NAME>

Description

Configures a certificate for the gNMI server. **local-cert** is used by default. For more details, refer to the [Public Key Infrastructure guide] (./Functionality_Guide_PKI.md)

Parameter	Description
<CERT-NAME>	Specifies the certificate name.

Examples

Configure sign-cert for the gNMI server:

```
switch(config)# crypto pki application gnmi certificate sign-cert
```

Command History

Release	Modification
10.17	Command introduced

Command Information

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

gnmi vrf

```
gnmi vrf <VRF NAME>
no gnmi vrf <VRF NAME>
```

Description

Enables the gNMI server on a given VRF. gNMI can be enabled on multiple VRFs simultaneously.



There is a maximum of 30 gNMI subscriptions per request and a maximum of 9 concurrent streams per switch. These are system-wide limits shared among all authenticated users regardless of connection method (local or remote).

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

Disabling the gNMI server on a VRF immediately closes any active streams.

Parameter	Description
<VRF NAME>	Specifies the VRF name.

Examples

Enable the gNMI server on VRF mgmt, this allows access to the gNMI server from the OOBM port in the "management VRF":

```
switch(config)# gnmi vrf mgmt
```

Removing the configuration of gNMI server on mgmt:

```
switch(config)# no gnmi vrf mgmt
```

Command History

Release	Modification
10.17	Command introduced

Command Information

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

show gnmi

```
show gnmi
```

Description

Displays the current gNMI configuration.

Examples

Display the current gNMI configuration:

```
switch(config)# show gmni
gNMI Configuration
-----
VRF                : mgmt, default
Access mode        : read-only
Global stream limit : 9
```

Command History

Release	Modification
10.17	Command introduced

Command Information

Platforms	Command context	Authority
6300 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.

boot fabric-module

boot fabric-module <SLOT-ID>

Description

Reboots the specified fabric module.

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

Usage

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

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

This command is valid for fabric modules only.

Examples

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

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

switch#
```

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

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

switch#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
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. Changes to this setting may impact other features, such as the job scheduler feature configured via the **reload at** or **reload after** commands.

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

Example

Selecting the primary image as the default boot image:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

boot system

boot system [primary | secondary | serviceos]

Description

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

Parameter	Description
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.

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-<VOLUME-NAME>	Administrators or local user group members with execution rights for this command.

directory

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

Description

Selects an existing directory on the external storage volume.

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

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

Examples

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

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

Clearing the directory /home:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
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	<code>config-external-storage-<VOLUME-NAME></code>	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


```
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.
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-<VOLUME-NAME>	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
<USER-NAME>	Specifies the username.

Examples

Creating a volume named logfiles with the user name nassuser:

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

Clearing the user name nasuser from accessing the logfiles volume:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 6400	config-external-storage-<VOLUME-NAME>	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
<VRF-NAME>	Specifies the VRF name.

Examples

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

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

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
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 500.
- 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 explicitly configured, the system automatically uses the first DNS server available in its default configuration.
- 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.
- The configuration of history results is limited to a maximum of 8 IP-SLA sessions. This means that you can enable and store historical performance data, such as response times and availability, for up to 8 individual IP SLA sessions at any given time.
- 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 maximum value when UDP jitter SLA is enabled. Otherwise, 100% packet loss can be seen and UDP jitter SLA probes will result in failure:

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

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> [history-interval <HISTORY-INTERVAL>] [cache disable] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}] [probe-interval <PROBE-INTERVAL>] [proxy <PROXY-URL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|IFNAME}] [source-port <SOURCE-PORT-NUM>] [version <VERSION-NUMBER>] [http-raw-request <RAW-PAYLOAD>]
```

```
no http {get | raw} <URL> [history-interval <HISTORY-INTERVAL>] [cache disable] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}] [probe-interval <PROBE-INTERVAL>] [proxy <PROXY-URL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|IFNAME}] [source-port <SOURCE-PORT-NUM>] [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).

The **no** version of this command disables HTTP as the IP-SLA test mechanism.

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>.
history-interval <HISTORY-INTERVAL>	Configures the history interval for the IP-SLA. Set the history

Parameter	Description
	interval to minimum of two times the probe-interval for the SLA. Range: 60 to 7200.
cache disable	Selects cache option for the HTTP server. By default the option is enabled.
name-server {<IPV4-ADDR-DNS-SERVER> <IPV6-ADDR-DNS-SERVER>}	Specifies the DNS server for destination hostname resolution.
probe-interval <PROBE-INTERVAL>	Specifies the probe interval in seconds. Range: 30 to 604800.
proxy <PROXY-URL>	Specifies the probe interval in seconds. Range: 30 to 604800.
source {<IPV4-ADDR> <IPV6-ADDR> IFNAME}	Selects the source, either an IPv4 address, an IPv6 address, or hostname for SLA probes.
source-port <SOURCE-PORT-NUM>	Specifies the value of the source port for the IP-SLA probes.
version <VERSION-NUMBER>	Specifies the source interface to use for sending IP-SLA probes.
http-raw-request <RAW-PAYLOAD>	Specifies the HTTP raw request. String.

Examples

Configuring HTTP get with parameters, including history interval:

```
switch(config)# ip-sla 1
switch(config-ipsla-1)# http get http://device.arubanetworks.com/root/home.html
history-interval 120 cache disable name-server 10.10.10.2 probe-interval 30
```

Configuring HTTP raw with parameters:

```
switch(config-ipsla-1)# http raw http://2.2.2.2 http-raw-request "GET
/en/US/hmpgs/index.html HTTP/1.0\r\n\r\n"
```

Disabling HTTP get with parameters:

```
switch(config-ipsla-1)# no http get http://device.example.com/root/home.html name-
server 10.10.10.2 history-interval 120
```

Disabling HTTP raw with parameters:

```
switch(config-ipsla-1)# no http raw http://device.example.com/root/home.html http-raw-request "GET /en/US/hmpgs/index.html HTTP/1.0\r\n\r\n"
```

Command History

Release	Modification
10.16.1000	Added new parameter history-interval . Also, IPv6 addresses can be used.
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> [history-interval <HISTORY-INTERVAL>] [cache disable] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}] [probe-interval <PROBE-INTERVAL>] [proxy <PROXY-URL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|IFNAME}] [source-port <SOURCE-PORT-NUM>] [version <VERSION-NUMBER>] [http-raw-request <RAW-PAYLOAD>]
```

```
no https {get | raw} <URL> [history-interval <HISTORY-INTERVAL>] [cache disable] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}] [probe-interval <PROBE-INTERVAL>] [proxy <PROXY-URL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|IFNAME}] [source-port <SOURCE-PORT-NUM>] [version <VERSION-NUMBER>] [http-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>.
history-interval <HISTORY-INTERVAL>	Configures the history interval

Parameter	Description
	for the IP-SLA. Set the history interval to minimum of two times the probe-interval for the SLA. Range: 60 to 7200.
cache disable	Selects cache option for the HTTPS server. By default the option is enabled.
name-server {<IPv4-ADDR-DNS-SERVER> <IPv6-ADDR-DNS-SERVER>}	Specifies the IPv4 address of DNS server.
probe-interval <PROBE-INTERVAL>	Specifies the probe interval in seconds. Range: 30 to 604800.
proxy <PROXY-URL>	Specifies the probe interval in seconds. Range: 30 to 604800.
source {<IPv4-ADDR> <IPv6-ADDR> IFNAME}	Selects the source, either an IPv4 address, an IPv6 address, or hostname for SLA probes.
source-port <SOURCE-PORT-NUM>	Specifies the value of the source port for the IP-SLA probes.
version <VERSION-NUMBER>	Specifies the source interface to use for sending IP-SLA probes.
https-raw-request <RAW-PAYLOAD>	Specifies the HTTPS raw request. String.

Examples

Configuring HTTPS get with parameters:

```
switch(config-ipsla-1)# https get https://device.arubanetworks.com/root/home.html
```

Configuring HTTPS raw with parameters:

```
switch(config-ipsla-1)# https raw https://device.arubanetworks.com/root/home.html
raw-request "GET /en/US/hmpgs/index.html"
```

Removing the HTTPS raw:

```
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
10.16.1000	Added new parameters history-interval . Also, IPv6 addresses can be used.
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>|<DEST-IPV6-ADDR>|<HOSTNAME>} [history-interval <HISTORY-INTERVAL>] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>} [payload-size <PAYLOAD-SIZE>] [probe-interval <PROBE-INTERVAL>] [source {<SOURCE-IPV4-ADDR>|<SOURCE-IPV6-ADDR>|<IFNAME>}]] [timeout <TIMEOUT>] [tos <TYPE-OF-SERVICE>]
```

```
no icmp-echo {<DEST-IPV4-ADDR>|<DEST-IPV6-ADDR>|<HOSTNAME>} [history-interval <HISTORY-INTERVAL>] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>} [payload-size <PAYLOAD-SIZE>] [probe-interval <PROBE-INTERVAL>] [source {<SOURCE-IPV4-ADDR>|<SOURCE-IPV6-ADDR>|<IFNAME>}]] [timeout <TIMEOUT>] [tos <TYPE-OF-SERVICE>]
```

Description

Configures ICMP echo as the IP-SLA test mechanism. Requires destination address for the IP-SLA test. The **no** form of this command disables the ICMP echo as the IP-SLA test mechanism.

Parameter	Description
{<DEST-IPV4-ADDR> <DEST-IPV6-ADDR> <HOSTNAME>}	Selects the destination, either an IPv4 address, an IPv6 address, or hostname, for the IP-SLA.
history-interval <HISTORY-INTERVAL>	Specifies the history interval for the IP-SLA. Set the history interval to minimum of two times the probe-interval for the SLA. Range: 10 to 7200.
name-server {<IPV4-ADDR-DNS-SERVER> <IPV6-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.
probe-interval <PROBE-INTERVAL>	Specifies the probe interval in seconds. Range: 5 to 604800.
source {<SOURCE-IPV4-ADDR> <SOURCE-IPV6-ADDR> <IFNAME>}	Selects the source IPv4 or IPv6 address for SLA probes or the

Parameter	Description
	source interface to use for sending IP-SLA probes.
<code>timeout <TIMEOUT></code>	Specifies the interval before a probe is timed out. Range: 5 to 604800.
<code>tos <TYPE-OF-SERVICE></code>	Specifies the type of service value to be used in probe packets. Range: 0 to 255.

Examples

Configuring icmp-echo:

```
switch(config)# ip-sla test
switch(config-ip-sla-test)# icmp-echo 2.2.2.2 name-server 4.4.4.4 source 3.3.3.3
```

Configuring ICMP echo with several parameters, including history interval and timeout:

```
switch(config)# ip-sla test
switch(config-ip-sla-test)# icmp-echo 2.2.2.2 history-interval 160 name-server
4.4.4.4 payload-size 400 probe-interval 80 source 3.3.3.3 timeout 20 tos 255
```

Command History

Release	Modification
10.16.1000	Added two new parameters history-interval and timeout . Also, IPv6 addresses can be used.
10.07 or earlier	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400	<code>config-ip-sla-<IP-SLA-NAME></code>	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
<code><IP-SLA-NAME></code>	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>|<SOURCE-IPV6-ADDR>|<IFNAME>}] [vrf <VRF-NAME>] [ipv6]
```

```
no ip-sla responder <SLA-NAME> (udp-echo | tcp-connect | udp-jitter-voip) [<PORT-NUM>]
[source {<SOURCE-IPV4-ADDR>|<SOURCE-IPV6-ADDR>|<IFNAME>}] [vrf <VRF-NAME>] [ipv6]
```

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.

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

Parameter	Description
<code><SLA-NAME></code>	Specifies the IP-SLA responder name. Length: 1 to 64 characters.
udp-echo	Enables responder for udp-echo probes.
tcp-connect	Selects TCP connect as the IP-SLA

Parameter	Description
	test mechanism.
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> <SOURCE-IPV6-ADDR> <IFNAME>}	Selects the source IPv4 or IPv6 address for SLA probes or the source interface to use for sending IP-SLA probes.
vrf <VRF-NAME>	Specifies the name of the VRF to use.
ipv6	Configures IPv6 responder. This keyword is required if an IPv6 address is being used by the source interface or VRF. By default, it will be considered as an IPv4 address.

Usage

The IPv6 keyword is required if an IPv6 address is being used by the source interface or VRF. Otherwise, by default, it will be considered as an IPv4 address.

Examples

Configuring IP-SLA responder for udp-echo:

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

Configuring IP-SLA responder with IPv6:

```
switch(config)# ip-sla responder SLA1 udp-echo 8000 source 1/1/1 ipv6
```

Configuring IP-SLA responder for udp-jitter-voip:

```
switch(config)# ip-sla responder SLA1 udp-jitter-voip 1025 vrf <VRF>
```

Disabling IP-SLA responder:

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

Command History

Release	Modification
10.15	Added ipv6 parameter.
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 all

show ip-sla all

Description

Shows all ip-sla source configuration and status.

Examples

Showing results for ip-sla all:

```
switch# show ip-sla all
SLA Name           : 2 (non-persistent)
Status             : running
SLA Type           : udp-echo
VRF                : default
Source IP          :
Source Interface    :
Domain Name Server :
Payload Size       : 28
TOS                 : 0
Probe Interval(seconds) : 60
History Interval(seconds) : 0
Timeout Interval(seconds) : 45

IP-SLA session status
IP-SLA Name         : 2 (non-persistent)
IP-SLA Type         : udp-echo
Destination Host Name/IP Address : 10.1.1.2
Destination Port     : 8888
Source IP Address/IFName :
Source Port          :
Status              : running

IP-SLA Session Cumulative Counters
Total Probes Transmitted : 10
Probes Timed-out         : 10
Bind Error               : 0
Destination Address Unreachable : 0
DNS Resolution Failures  : 0
Reception Error          : 0
Transmission Error       : 0
Operational Status       : down

IP-SLA Latest Probe Results
```

```
Last Probe Time           :
Packets Sent              : 1
Packets Received          : 0
Packet Loss in Test       : 100%
```

```
Minimum RTT(ms)          :
Maximum RTT(ms)          :
Average RTT(ms)          :
DNS RTT(ms)              :
```

```
-----
SLA Name                  : echo-udp-sess2 (non-persistent)
Status                    : running
SLA Type                  : udp-echo
VRF                        : default
Source IP                 :
Source Interface          :
Domain Name Server        :
Payload Size              : 28
TOS                       : 0
Probe Interval(seconds)   : 60
History Interval(seconds) : 0
Timeout Interval(seconds) : 45
```

```
IP-SLA session status
IP-SLA Name               : echo-udp-sess2 (non-persistent)
IP-SLA Type               : udp-echo
Destination Host Name/IP Address : 100.1.1.2
Destination Port          : 8888
Source IP Address/IFName  :
Source Port               :
Status                    : running
```

```
IP-SLA Session Cumulative Counters
Total Probes Transmitted   : 10
Probes Timed-out          : 0
Bind Error                 : 0
Destination Address Unreachable : 0
DNS Resolution Failures    : 4
Reception Error           : 0
Transmission Error        : 0
Operational Status        : Up
```

```
IP-SLA Latest Probe Results
Last Probe Time           :
Packets Sent              : 1
Packets Received          : 1
Packet Loss in Test       : 0.0000%
```

```
Minimum RTT(ms)          :
Maximum RTT(ms)          :
Average RTT(ms)          :
DNS RTT(ms)              :
```

Showing results for non-configured ip-sla all:

```
switch# show ip-sla all
IPSLA source is not configured
```

Command History

Release	Modification
10.16.1000	Updated to display history interval .
10.12.1000	Updated to display https as an IP-SLA type.
10.07 or earlier	Command introduced.

Command Information

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

show ip-sla responder

```
show ip-sla responder <SLA-NAME> [initiator {<SOURCE-IPV4-ADDR>|<SOURCE-IPV6-ADDR>}]  
[<SOURCE-PORT-NUM>] [results]
```

Description

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

Parameter	Description
<SLA-NAME>	Specifies the SLA name.
initiator {<SOURCE-IPV4-ADDR> <SOURCE-IPV6-ADDR>}	Selects the source IPv4 or IPv6 address for SLA probes to use.
<SOURCE-PORT-NUM>	Configures the source port for the IP-SLA test. Range: 1 to 65535.
results	Displays the statistics for a given source IP and port.

Examples

Showing IP-SLA responder configuration:

```
switch# 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
```

Showing IP-SLA responder with initiator and results parameters:

```
switch# show ip-sla responder SLA1 initiator 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
10.16.1000	Added new parameters: initiator and results .
10.07 or earlier	Command introduced.

Command Information

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

show ip-sla

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

Description

Shows the given IP-SLA source configuration and status.

Parameter	Description
<SLA-NAME>	Specifies the SLA name.
results	Displays the statistics calculated for an SLA type.
history-results	Displays the history statistics calculated for the SLA ID.

Examples

Showing results for ip-sla:

```
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

```

Showing history results for ip-sla:

```

switch# sh ip-sla abcd history-results
IP-SLA Name: abcd

Session Details
=====
IP-SLA Type          : tcp-connect      Status              : running

Probe Interval(seconds) : 30          History Interval(seconds) : 600

Source Port          : 3000          Destination Port      : 4000

Source IP Address     : 10.0.0.1
Dest Host Name/IP Address : 10.0.0.2

History Probe Results
=====
Packet Stats
-----
Probes Transmitted    : 4          Packets Sent         : 4
Packets Received      : 4          Loss Percentage       : 0.0000%

Error Stats
-----
Transmission Errors   : 0          Reception Errors     : 0
Bind Errors           : 0          Dest. Unreachable    : 0
Probes Timed-Out      : 0          DNS Resolution Failures : 0

Probe RTT Stats
-----
Min RTT(ms)          : 0          Max RTT(ms)          : 0
Avg RTT(ms)           : 0

DNS RTT Stats

```

```
-----  
Min  RTT (ms)           :           Max  RTT (ms)           :  
Avg  RTT (ms)           :           :
```

Command History

Release	Modification
10.16.1000	Added new parameter history-results . Updated to display history interval .
10.12.1000	Updated to display https as an IP-SLA type.
10.07 or earlier	Command introduced.

Command Information

Platforms	Command context	Authority
6300 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 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 6400	config-ip-sla- <i><IP-SLA-NAME></i>	Administrators or local user group members with execution rights for this command.

tcp-connect

```
tcp-connect {<DEST-IPV4-ADDR>|<DEST-IPV6-ADDR>|<HOSTNAME>} <DEST-PORT-NUM> [history-
interval <HISTORY-INTERVAL>] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}
[probe-interval <PROBE-INTERVAL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|<IFNAME>}] [source-port
<PORT-NUM>]
```

```
no tcp-connect {<DEST-IPV4-ADDR>|<DEST-IPV6-ADDR>|<HOSTNAME>} <DEST-PORT-NUM> [history-
interval <HISTORY-INTERVAL>] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}
[probe-interval <PROBE-INTERVAL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|<IFNAME>}] [source-port
<PORT-NUM>]
```

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.

The **no** form of this command removes the the TCP connection.

Parameter	Description
{<DEST-IPV4-ADDR> <DEST-IPV6-ADDR> <HOSTNAME>}	Selects the destination, either an IPv4 address, an IPv6 address, or hostname, for the IP-SLA.
<DEST-PORT-NUM>	Destination port for the IP-SLA. Range: 1 to 65535.
history-interval <HISTORY-INTERVAL>	Configures the history interval for the IP-SLA. Set the history interval to minimum of two times the probe-interval for the SLA. Range: 60 to 7200.

Parameter	Description
name-server {<IPV4-ADDR-DNS-SERVER> <IPV6-ADDR-DNS-SERVER>}	Specifies the DNS server for destination hostname resolution.
probe-interval <PROBE-INTERVAL>	Probe interval in seconds. Range: 30 to 604800.
source {<IPV4-ADDR> <IPV6-ADDR> IFNAME>}	Selects the source, either an IPv4 address, an IPv6 address, or hostname for SLA probes.
source-port <PORT-NUM>	Specifies the port for the IP-SLA test.

Examples

Configuring TCP connect echo with parameters, including history interval:

```
switch(config)# ip-sla tcp
switch(config-ip-sla-tcp)# tcp-connect https://device.example.com 8080 name-server
10.10.10.2 history-interval 180 source 1/1/1 source-port 6000
```

Disabling the TCP connect:

```
switch(config-ip-sla-tcp)# no tcp-connect 10:1::1:1 8080 source 1/1/1 source-port
6000
```

Command History

Release	Modification
10.16.1000	Added new parameters history-interval . Also, IPv6 addresses can be used.
10.07 or earlier	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.

udp-echo

```
udp-echo {<DEST-IPV4-ADDR>|<DEST-IPV6-ADDR>|<HOSTNAME>} <DEST-PORT-NUM> [history-interval
<HISTORY-INTERVAL>] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}
[payload-size <PAYLOAD-SIZE>] [probe-interval <PROBE-INTERVAL>] [source {<IPV4-
ADDR>|<IPV6-ADDR>|IFNAME>}] [source-port <SOURCE-PORT-NUM>] [timeout <TIMEOUT>] [tos
<TYPE-OF-SERVICE>]
```

```
no udp-echo {<DEST-IPV4-ADDR>|<DEST-IPV6-ADDR>|<HOSTNAME>} <DEST-PORT-NUM> [history-
interval <HISTORY-INTERVAL>] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}]
```

```
[payload-size <PAYLOAD-SIZE>] [probe-interval <PROBE-INTERVAL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|IFNAME}] [source-port <SOURCE-PORT-NUM>] [timeout <TIMEOUT>] [tos <TYPE-OF-SERVICE>]
```

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.

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

Parameter	Description
{<DEST-IPV4-ADDR> <DEST-IPV6-ADDR> <HOSTNAME>}	Selects the destination, either an IPv4 address, an IPv6 address, or hostname, for the IP-SLA.
<DEST-PORT-NUM>	Specifies the destination port for the IP-SLA. Range: 1 to 65535.
history-interval <HISTORY-INTERVAL>	Configures the history interval for the IP-SLA. Set the history interval to minimum of two times the probe-interval for the SLA. Range: 10 to 7200.
name-server {<IPV4-ADDR-DNS-SERVER> <IPV6-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.
probe-interval <PROBE-INTERVAL>	Specifies the probe interval in seconds. Range: 5 to 604800.
source {<IPV4-ADDR> <IPV6-ADDR> IFNAME}	Selects the source, either an IPv4 address, an IPv6 address, or hostname for SLA probes.
source-port <SOURCE-PORT-NUM>	Configures the source port for the IP-SLA test. Range: 1 to 65535.
timeout <TIMEOUT>	Specifies the interval before a probe is timed out. Range: 5 to 604800.
tos <TYPE-OF-SERVICE>	Specifies the type of service. Range: 0 to 255.

Examples

Configuring UDP echo with parameters, including history interval and timeout:

```
switch(config-ipsla-1)# udp-echo https://device.example.com 4000 history-interval 180 name-server 2.2.2.2 payload-size 100 probe-interval 90 source 4.4.4.4 source-
```

```
port 8000 timeout 20
```

Removing UDP echo configuration:

```
switch(config-ipsla-1)# no udp-echo https://device.example.com 8080 history-  
interval 160 name-server 10.10.10.2 payload-size 50 source 2.2.2.1 timeout 20
```

Command History

Release	Modification
10.16.1000	Added two new parameters history-interval and timeout . Also, IPv6 addresses can be used.
10.07 or earlier	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.

udp-jitter-voip

```
udp-jitter-voip {<DEST-IPV4-ADDR>|<DEST-IPV6-ADDR>|<HOSTNAME>} <DEST-PORT-NUM>  
[advantage-factor <ADVANTAGE-FACTOR>] [codec-type <CODEC-TYPE>] [history-interval  
<HISTORY-INTERVAL>] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}] [probe-  
interval <PROBE-INTERVAL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|IFNAME}] [source-port  
<SOURCE-PORT-NUM>] [tos <TYPE-OF-SERVICE>]
```

```
no udp-jitter-voip {<DEST-IPV4-ADDR>|<DEST-IPV6-ADDR>|<HOSTNAME>} <DEST-PORT-NUM>  
[advantage-factor <ADVANTAGE-FACTOR>] [codec-type <CODEC-TYPE>] [history-interval  
<HISTORY-INTERVAL>] [name-server {<IPV4-ADDR-DNS-SERVER>|<IPV6-ADDR-DNS-SERVER>}] [probe-  
interval <PROBE-INTERVAL>] [source {<IPV4-ADDR>|<IPV6-ADDR>|IFNAME}] [source-port  
<SOURCE-PORT-NUM>] [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.

The **no** form of this command removes the UDP jitter VoIP configuration.

Parameter	Description
{<DEST-IPV4-ADDR> <DEST-IPV6-ADDR> <HOSTNAME>}	Selects the destination, either an IPv4 address, an IPv6 address, or hostname, for the IP-SLA.
<DEST-PORT-NUM>	Selects the port number for the IP-SLA. Range: 1 to 65535.

Parameter	Description
<code>advantage-factor <ADVANTAGE-FACTOR></code>	Selects the value for the advantage factor. Default value is 0. Range: 0 to 20.
<code>codec-type <CODEC-TYPE></code>	Selects the codec-type for the Voip IP-SLA test.
<code>history-interval <HISTORY-INTERVAL></code>	Configures the history interval for the IP-SLA. Set the history interval to minimum of two times the probe-interval for the SLA. Range: 240 to 7200.
<code>name-server {<IPV4-ADDR-DNS-SERVER> <IPV6-ADDR-DNS-SERVER>}</code>	Specifies the DNS server for destination hostname resolution.
<code>probe-interval <PROBE-INTERVAL></code>	Specifies the probe interval in seconds. Range: 120 to 604800.
<code>source {<IPV4-ADDR> <IPV6-ADDR> IFNAME}</code>	Selects the source, either an IPv4 address, an IPv6 address, or hostname for SLA probes.
<code>source-port <SOURCE-PORT-NUM></code>	Specifies the value of source port for the IP-SLA probes.
<code>tos <TYPE-OF-SERVICE></code>	Specifies the type of service. Range: 0 to 255.

Examples

Configuring udp-jitter-voip with optional parameters, including history interval:

```
switch(config-ipsla-1)# udp-jitter-voip https://device.arubanetworks.com 8080
advantage-factor 10 codec-type g711a history-interval 240 name-server 10.10.10.2
probe-interval 120 source 10.1.1.1 source-port 8888 tos 10
```

Configuring udp-jitter-voip with optional parameters:

```
switch(config-ipsla-1)# udp-jitter-voip 2.2.2.2 8080 advantage-factor 10 codec-
type g711a
```

Disabling udp-jitter-voip:

```
switch(config-ipsla-1)# no udp-jitter-voip https://device.example.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
10.16.1000	Added new parameters history-interval . Also, IPv6 addresses can be used.
10.07 or earlier	Command introduced.

Command Information

Platforms	Command context	Authority
6300 6400	config-ip-sla- <i><IP-SLA-NAME></i>	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
<i><VRF-NAME></i>	Specifies a VRF name. Length: Default: default. If no VRF name is specified, then the default VRF name will be taken.

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 6400	config-ip-sla- <i><IP-SLA-NAME></i>	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 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	Shows additional statistics, including the tx filtered and rx filtered counters. ▪ Rx filter packets are protocol packets received when the protocol is disabled on the switch and there is only one port in the VLAN. Protocols include OSPF, PIM, RIP, LACP, and LLDP. ▪ An example of a Tx filtered packet would be a multicast packet being filtered from going out of the ingress port.

Parameter	Description
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		Status	Config	Port Status	Config	Status	Config
Status	Config	(Watts)	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                      RX                      TX                      Total (RX+TX)
-----
Bits / sec                3M                      3M                      6M
Pkts / sec                316                    316                    633
Unicast                   319                    319                    638
Multicast                  0                      0                      0
Broadcast                  0                      0                      0
Utilization %             < 1                    < 1                    < 1
Statistic                 RX                      TX                      Total
-----
Packets                   577K                    577K                    1M
Unicast                   577K                    577K                    1M
Multicast                  0                      51                     51
Broadcast                  0                      15                     15
Bytes                     744M                    745M                    1G
Jumbos                     0                      0                      0
Dropped                    0                      0                      0
Filtered                   0                      0                      0
Pause Frames               0                      0                      0
Errors                     0                      0                      0
CRC/FCS                    0                      n/a                    0
Collision                  n/a                    0                      0
Runts                      0                      n/a                    0
Giants                     0                      n/a                    0

```

Showing information about extended counters:



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

```
switch(config-if)# show interface 1/1/17 extended
-----
Interface 1/1/17
-----
Statistics                                     Value
-----
Dot1d Tp Port In Frames                      547
Dot1d Tp Port Out Frames                     608
Dot3 In Pause Frames                         0
Dot3 Out Pause Frames                       0
Ethernet Stats Broadcast Packets             19
Ethernet Stats Bytes                        40162
Ethernet Stats Packets                      342
...
-----
Error-Statistics                             Value
-----
Dot1d Base Port MTU Exceeded Discards        0
Dot3 Control In Unknown Opcodes              0
Dot3 Stats Alignment Errors                  0
Dot3 Stats FCS Errors                       0
Dot3 Stats Frame Too Longs                  0
Dot3 Stats Internal Mac Transmit Errors      0
Ethernet RX Oversize Packets                 0
...
```

Showing interface link-status:

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

Port          Type          Physical Link      Last
              Type          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
              Type          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	--

Showing state information when interface is blocked:

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

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

Rate	RX	TX	Total (RX+TX)
Mbits / sec	0.00	0.00	0.00
KPkts / sec	0.00	0.00	0.00
Unicast	0.00	0.00	0.00
Multicast	0.00	0.00	0.00
Broadcast	0.00	0.00	0.00
Utilization %	0.00	0.00	0.00
Statistic	RX	TX	Total
Packets	15	15	30

Unicast	12	12	24
Multicast	3	3	6
Broadcast	0	0	0
Bytes	1350	1350	2700
Jumbos	0	0	0
Dropped	0	0	0
Pause Frames	0	0	0
Errors	0	0	0
CRC/FCS	0	n/a	0
Collision	n/a	0	0
Runts	0	n/a	0
Giants	0	n/a	0

Command History

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

Command Information

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

show interface statistics

```

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

Parameter	Description
<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 P
1/1/1		2727136	1975	0	17796	195	0	82	1788	96	54	0	
1/1/10		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	
1/1/12		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	
1/1/31	- lag2	2360	25	50	2732119	2040	0	0	0	178	1839	0	
1/1/32	- lag2	0	0	0	11373	93	0	0	0	0	51	0	
vlan1		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 <code>moitor</code> parameter.
10.10	Added <code>human-readable</code> parameter.
10.07 or earlier	--

Command Information

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

show interface downshift-enable

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

Description

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

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

Examples

Showing automatic downshift information:

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

Showing automatic downshift information on per interface:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 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 [<IFNNAME>|<IFRANGE>]
show running-config interface [lag | loopback | tunnel | vlan ] [<ID>]
```

Description

Displays active configurations of various switch interfaces.

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

Examples

Showing 1/1/2 interface configuration:

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

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

Showing loopback interfaces configured:

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

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

Showing loopback interfaces not configured:

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

No loopback interfaces configured.
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

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

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

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



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

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

- Support for a VLAN as source for a mirror session
- Ability for a given Mirror source in one session to act as source in another Mirror Session
- Support for a Layer 2/bridged Link Aggregation Group (LAG) as Session destination
- Support for a Layer 3/Route-only Link Aggregation Group (LAG) as Session destination

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



- Management
 - Persona
 - Loopback
 - VxLAN
 - Subinterfaces
-

Hit Count Behavior for Policy Mirroring Actions

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

Mirroring and CoPP

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

Control plane packets that are already destined for the switch (for example, routing protocol packets and management traffic) will be processed by their respective CoPP classes. These packets will still be

mirrored to the CPU as part of the Mirror-to-CPU session, but they will not be counted under the Mirror-to-CPU CoPP class. The Mirror-to-CPU CoPP class is specifically designed to manage dataplane traffic that is mirrored to the CPU. Any control plane traffic mirrored to the CPU will bypass this class and be handled by the CoPP class corresponding to its original purpose (e.g., OSPF, BGP, or ARP). Administrators should account for this distinction when analyzing traffic mirrored to the CPU.

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

Mirror statistics

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

Classifier policies and mirroring sessions

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

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

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

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

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

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

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

Mirroring on a VSX configuration

On 6400 switch series with a VSX configuration, north to south traffic may enter one switch, then cross the ISL to the other switch to reach its destination. If both 6400 switches have the following mirror configuration and the analyzer is connected to 1/1/10 on both 6400s, then the packet is mirrored by each 6400 switch, and two copies are sent to the analyzer.

```
mirror session 1
source vlan 1 both
```

```
source vlan 2 both
...
destination 1/1/10
```

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

Example

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

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 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 6400	config-mirror- <i><SESSION-ID></i>	Administrators or local user group members with execution rights for this command.

destination interface

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

Description

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

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

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

Usage

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

Examples

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

Examples

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

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

Replacing the existing tunnel destination:

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

Replacing the existing destination with a different DSCP value:

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

Removing the destination:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 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.

Parameter	Description
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
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 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 HPE Aruba Networking 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-<SESSION-ID></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><SESSION-ID></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)#
```



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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	<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 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
<SESSION-ID>	Specifies the session identifier. Range: 1 to 4
vsx-peer	Shows the output from the VSX peer switch. If the switches do not have the VSX configuration or the ISL is down, the output from the VSX peer switch is not displayed. This parameter is available on switches that support VSX.

Usage

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

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

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

- **dest_doesnt_exist**: The configured destination interface is not found in the system. The mirroring session cannot be enabled.
- **destination_shutdown**: **The mirroring session is enabled, but the destination interface is shut down. No traffic** can be monitored.
- **disabled**: The mirroring session is disabled and is not in an error condition.
- **enabled**: The mirroring session is enabled.
- **external/driver_error**: An internal ASIC hardware error occurred.
- **hit_active_sessions_capacity**: The mirroring session could not be enabled because the maximum number of supported mirroring sessions are already enabled.
- **internal_error**: An invalid parameter was passed to the ASIC software layer.

- **no_dest_configured:** The mirroring session does not have a destination interface configured.
- **no_name_configured:** A software error occurred. The mirroring session does not have a session ID in its configuration.
- **null_mirror:** A software error occurred. The session object reference is invalid.
- **out_of_memory:** The system is out of memory, reboot recommended.
- **tunnel_route_resolution_not_populated:** If the destination tunnel IP address is not reachable.
- **unknown_error:** An unexpected error occurred.

Examples

On the HPE Aruba Networking 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
switch#
```

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

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

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

```
switch: show mirror 1
Admin Status: enable
Operation Status: enabled
Source: interface 1/1/1 rx
Source: interface lag1 rx
```

```
Destination: interface 1/1/2
Output Packets: 0
Output Bytes: 0
```

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

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

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

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

show mirror endpoint

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

Description

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

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

Examples

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

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

Showing the details of enabled mirror endpoint test:

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



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

Command History

Release	Modification
10.07 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.

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 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 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 member/slot/port (for example, 1/3/1).
<LAG-NAME>	Specifies the identifier for the LAG (link aggregation group).
<DIRECTION>	Selects the direction of traffic to be mirrored from this source interface. There is no default for this parameter. Valid values are the following: - both: Mirror both transmitted and received packets. - rx: Mirror only received packets. - tx: Mirror only transmitted packets.

Usage

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

Switch	Source interface limit 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- <i><SESSION-ID></i>	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. tx (transmit), rx (receive), or both .

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:
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 6400	config	Administrators or local user group members with execution rights for this command.

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.

Packet Capture

The Packet Capture is a solution which enables the Central administrator to capture or view packets from clients connected to Halon switches. The solution serves as a tool which an admin can leverage to debug clients connected to managed devices across branches in an enterprise. The switch requires packet exchanges to be copied to the CPU (both ingress *and* egress packets) per client session, so that they can be streamed to Central where the admin can view the packet exchanges for further analysis.

This feature is supported on the following platforms:

- Rosewood (6200, 6300, 6400, 8360)
- Matrix (4100i, 6000, 6100)

packet-capture commands

copy packet-capture

```
copy packet-capture-pcap <FILE_NAME> <REMOTE_URL>
```

Description

This command copies a capture file created using the [packet-capture](#) to a remote URL.

Parameter	Description
<FILE_NAME>	Specifies the name of the <i>tcpdump</i> capture file to be copied.
<REMOTE_URL>	Specifies a URL to copy the command output. Valid URL syntaxes: ▪ <code>sftp://USER@{IP HOST}[:PORT]/FILE</code> ▪ <code>tftp://{IP HOST}[:PORT][;blocksize=VAL]/FILE</code>

Example

```
8400x# copy packet-capture-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
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
5420	config	Administrators or local user group members with execution rights for this command.
6200		
6300		
6400		
8100		
8325		
8360		
9300		
10000		

packet-capture delete-pcap

packet-capture delete-pcap <FILE_NAME>

Description

This command deletes a capture file created with the [packet-capture](#) command.

Parameter	Description
<FILE_NAME>	Specifies the name of the <i>tcpdump</i> capture file to be deleted.

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
5420	config	Administrators or local user group members with execution rights for this command.
6200		

Platforms	Command context	Authority
6300		
6400		
8100		
8325		
8360		
9300		
10000		

packet-capture enable

```
packet-capture <SESSION_NAME> [enable]
no...
```

Description

Enables the packet capture session.

The **no** form of this command disables the packet capture for the specified session.

Parameter	Description
<SESSION_NAME>	Specifies the session name.
enable	Enables packet capture for the specified session.



Note: A session name must be specified first before issuing the *enable* command or its *no* form. If the session hasn't been configured as a [packet-capture](#) session prior to running these commands, the error message *Packet Capture session does not exist* is displayed.

Examples

Enable packet capture for *session1*:

```
switch# packet-capture session1 enable
```

Disable packet capture for *session1*:

```
switch# no packet-capture session1 enable
```



Note: If *session1* has been configured as a [packet-capture](#) session but isn't enabled, the error message *Packet Capture session not running* is displayed when issuing the *no* form of the command.

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
5420 6200 6300 6400 8100 8325 8360 9300 10000	config	Administrators or local user group members with execution rights for this command.

packet-capture

```
packet-capture <SESSION_NAME> <PORT-ID> ( mac { <SMAC> <DMAC> <ETHERTYPE> } | ipv4
{<PROTOCOL> <SIP> eq <SPORT> <DIP> eq <DPORT> } | ipv6 {<PROTOCOL> <SIP> eq <SPORT> <DIP>
eq <DPORT> }) file-name <PCAP_NAME> [timeout <TIMEOUT>] [buffer-size <BUFFER_SIZE>] [sw-
filter <SW_FILTERS>] [vlan <VLAN>]
```

Description

Configures a new packet capture session.

Parameter	Description
<SESSION_NAME>	Configures the session name with a maximum of 32 characters.
<PORT-ID>	Configures the port ID. Applies to physical ports only.
<SMAC>	Configures the source MAC address.
<DMAC>	Configures the destination MAC address.
<ETHERTYPE>	Configures either a hexadecimal <i>EtherType</i> value (<H:600–FFFF>) or a supported protocol name (aarp, any, appletalk, arp, fcoe, fcoe-init, ip, ipv6, ipx-arpa, ipx-non-arpa, is-is, lldp, mpls-multicast, mpls-unicast, q-in-q, rbridge, trill, wake-on-lan).
<PROTOCOL>	Configures either a numeric protocol value (0–255) or a supported protocol name (ah, any, esp, gre, icmp, igmp, ip, ospf, pim, sctp, tcp, udp).
<SIP>	Configures the source IP address.

Parameter	Description
<DIP>	Configures the destination IP address.
<SPORT>	Configures the source port.
<DPORT>	Configures the destination port.
<TIMEOUT>	Configures the maximum session time for packet capture in seconds. Default: 300.
<BUFFER_SIZE>	Configures the buffer size in MB. Range: 1-100. Default: 100.
<PCAP_NAME>	Configures the packet capture file name. Length: 5-128 characters.
<SW_FILTERS>	Configures the software filter with a maximum of 512 characters. Standard <i>tcpdump</i> filters are supported for this parameter.
<VLAN>	Configures the VLAN identifier. Range: 1-4094.

Examples

Packet capture using MAC:

```
switch# packet-capture session1 1/1/1 mac 0050.5696.7cef 0050.5696.7cea 0800
```

Packet capture using IPV4:

```
switch# packet-capture session1 1/1/1 ipv4 tcp any eq any any eq any
```

Packet capture using IPV6:

```
switch# packet-capture session1 1/1/1 ipv6 any 2001:db8:abcd:12::2/ffff:ffff::0000
eq any any eq any
```

General Limitations

- Only 3 *pcap* files are allowed at one time. Creating a fourth *pcap* file will replace the oldest one.
- Packet capture is only supported on physical ports.
- Only one `packet-capture` session can be active a time.
- A `packet-capture` session cannot be active at the same time as a mirror session with destination CPU.
- Packets captured by `packet-capture` sessions will not include the Ethernet *frame check* sequence in the resulting capture.

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
5420	config	Administrators or local user group members with execution rights for this command.
6200		
6300		
6400		
8100		
8325		
8360		
9300		
10000		

show packet-capture pcaps

`show packet-capture pcaps`

Description

This command shows capture files created with the [packet-capture](#) command.

Parameter	Description
pcaps	Lists the saved <i>tcpdump</i> capture files.

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
5420	config	Administrators or local user group members with execution rights for this command.
6200		
6300		
6400		
8100		
8325		
8360		

Platforms	Command context	Authority
9300		
10000		

show packet-capture session

show packet-capture [<SESSION_NAME>]

Description

Displays information about packet capture sessions configured on the switch. If a *SESSION_NAME* is specified, it shows details for that session only.

Parameter	Description
<SESSION_NAME>	Specifies the name of the packet capture session.

Examples

Packet capture using MAC:

```
switch# show packet-capture
Name                : orange0
Admin State         : Disabled
Operation Status    : Inactive
Reason code         : max-timeout-reached
Port                : 1/1/1
Hardware filter     : ipv4
  Ethertype         : -
  Source MAC        : -
  Dest MAC          : -
  Source IP         : -
  Destination IP    : -
  Protocol          : -
  Source port min   : 21
  Source port max   : 21
  Dest port min     : 1001
  Dest port max     : 1001
  Vlan              : -
Software Filter     : -
Buffer size (in MB) : 100
Timeout (in second) : 300
Pcap file name      : session1.pcap
Capture direction   : Ingress and egress
```



On 5420, 6200, 6300, 6400, 8100 and 8360 Switch Series, both *ingress* and *egress* are supported capture directions.

On 8325, 9300 and 10000 Switch Series, only *ingress* is a supported capture direction.

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
5420	config	Administrators or local user group members with execution rights for this command.
6200		
6300		
6400		
8100		
8325		
8360		
9300		
10000		

Chapter 11

Power-over-Ethernet

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

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



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

PoE commands

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

lldp dot3 poe

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

Description

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

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

Examples

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 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 HPE Aruba Networking 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 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 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 HPE Aruba Networking 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 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.

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

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

Examples

On the HPE Aruba Networking 6400 Switch Series, interface identification differs.

Enabling pre-standard device detection:

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

Disabling pre-standard device detection:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 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 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 quick-poe <MODULE-ID>
```

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 HPE Aruba Networking 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 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 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 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: member/slot/port

Examples

On the HPE Aruba Networking 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 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: member/slot/port

Examples

On the HPE Aruba Networking 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 HPE_ANW IAP
Neighbor Chassis-ID        : 84:d4:7e:ce:5d:68
Neighbor Management-Address : 169.254.41.250
Chassis Capabilities Available : Bridge, WLAN
Chassis Capabilities Enabled :
Neighbor Port-ID           : 84:d4:7e:ce:5d:68
Neighbor Port-Desc         : eth0
TTL                         : 120
Neighbor Port VLAN ID      :
Neighbor PoEplus information : DOT3
Neighbor Device Type       : TYPE2 PD
Neighbor Power Priority     : Unknown
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 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

HPE Aruba Networking 6300 Switch Series:

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

HPE Aruba Networking 6400 Switch Series:

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

HPE Aruba Networking 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
      Total Power
PS      (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
      Total Power
PS      (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 show power-over-ethernet <IFRANGE> output.
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
6300 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
Configuration management	Device topology
	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
<code>{<IPv4-ADDR> <IPv6-ADDR> <HOSTNAME>}</code>	Selects the IPv4 address, IPv6 address, or host name of the remote syslog server. Required.
<code>[udp [<PORT-NUM>] tcp [<PORT-NUM>]]</code>	Specifies the UDP port or TCP port of the remote syslog server to receive the forwarded syslog messages.
<code>udp [<PORT-NUM>]</code>	Range: 1 to 65535. Default: 514
<code>tcp [<PORT-NUM>]</code>	Range: 1 to 65535. Default: 1470
<code>tls [<PORT-NUM>]</code>	Range: 1 to 65535. Default: 6514
<code>include-auditable-events</code>	Specifies that auditable messages are also logged to the remote syslog server.
<code>severity <LEVEL></code>	Specifies the severity of the syslog messages: ▪ <code>alert</code>: Forwards syslog messages with the severity of <code>alert</code> (6) and <code>emergency</code> (7). ▪ <code>crit</code>: Forwards syslog messages with the severity of <code>critical</code> (5) and above. ▪ <code>debug</code>: Forwards syslog messages with the severity of <code>debug</code> (0) and above. ▪ <code>emerg</code>: Forwards syslog messages with the severity of <code>emergency</code> (7) only. ▪ <code>err</code>: Forwards syslog messages with the severity of <code>err</code> (4) and above ▪ <code>info</code>: Forwards syslog messages with the severity of <code>info</code> (1) and above. Default. ▪ <code>notice</code>: Forwards syslog messages with the severity of <code>notice</code> (2) and above. ▪ <code>warning</code>: Forwards syslog messages with the severity of <code>warning</code> (3) and above.
<code>auth-mode</code>	Specifies the TLS authentication mode used to validate the certificate. ▪ <code>certificate</code>: Validates the peer using trust anchor certificate based authentication. Default. ▪ <code>subject-name</code>: Validates the peer using trust anchor certificates as well as subject-name based authentication.
<code>legacy-tls-renegotiation</code>	Enables the TLS connection with a remote syslog server supporting legacy renegotiation.
<code>vrf <VRF-NAME></code>	Specifies the VRF used to connect to the syslog server. Optional. Default: <code>default</code>

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.
trap version v3	Specifies the trap notification type for SNMPv3.

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

Example

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

snmpv3 context

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

Description

Creates an SNMPv3 context on the specified VRF.

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

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

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

Examples

Creating an SNMPv3 context named **newContext**:

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

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><NAME></code>	Specifies the SNMPv3 username. Range 1 - 32 printable ASCII characters, excluding space and question mark.
<code>auth <AUTH-PROTOCOL></code>	Specifies the authentication protocol used to validate user logins. Available options are: <code>md5</code> or <code>sha</code> .
<code>auth-pass {plaintext ciphertext} <AUTH-PWORD></code>	Specifies the SNMPv3 user password. Range for <code>plaintext</code> is 8 - 32 printable ASCII characters, excluding space and question mark. Range for <code>ciphertext</code> is 1 - 120 printable ASCII characters. This option is only used when copying user configuration settings between switches. It enables you to duplicate a user's configuration on another switch without having to know their password.
<code>priv <PRIV-PROTOCOL></code>	Specifies the SNMPv3 security protocol (encryption method). Available options are: <code>aes</code> or <code>des</code> .
<code>priv-pass {plaintext ciphertext} <PRIV-PWORD></code>	Specifies the SNMPv3 user privacy passphrase. Range for <code>plaintext</code> is 8 - 32 printable ASCII characters, excluding space and question mark. Range for <code>ciphertext</code> is 1 - 120 printable ASCII characters. This option is only used when copying user configuration settings between switches. It enables you to duplicate a user's configuration on another switch without having to know their password.

Examples

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

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

Removing an SNMPv3 user named **Admin**:

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

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

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

Copying an SNMP user from switch 1 to switch 2.

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

```
switch1(config)# snmpv3 user Admin auth sha auth-pass plaintext mypassword
priv des priv-pass plaintext myprivpass
switch1(config)# exit
switch1# show running-config
Current configuration:
!
!Version AOS-CX TL.10.00.0003-8017-gdeb0606~dirty
!
!
!
snmpv3 user Admin auth sha auth-pass ciphertext
AQBapZHf2d20GYr/xcGUzYzm0zjNf/4VKHtSqbNImqtFYbJYCgAAALkGFJVcSp3nZ3o=
priv des priv-pass ciphertext
AQBapb0H2poBQKXPoVsC9L9qzZyfJQnzR7hmTr7LGsOsI7K3CgAAAKP98Rq2jfTrFwQ=
ssh server vrf mgmt
!
!
!
!
interface mgmt
    no shutdown
    ip dhcp
    vlan 1
```

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

```
switch1(config)# snmpv3 user Admin auth sha auth-pass
ciphertextAQBapZHf2d20GYr/xcGUzYzm0zjNf/4VKHtSqbNImqtFYbJYCgAAALkGFJVcSp3nZ3o=priv
des priv-pass ciphertext
AQBapb0H2poBQKXPoVsC9L9qzZyfJQnzR7hmTr7LGsOsI7K3CgAAAKP98Rq2jfTrFwQ=
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

Accessing HPE Aruba Networking Support

HPE Aruba Networking Support Services	https://www.hpe.com/us/en/networking/hpe-aruba-networking-support-services.html
AOS-CX Switch Software Documentation Portal	https://arubanetworking.hpe.com/techdocs/AOS-CX/help-portal/Content/home.htm
HPE Aruba Networking Support Portal	https://networkingsupport.hpe.com/home
North America telephone	1-800-943-4526 (US & Canada Toll-Free Number) +1-650-750-0350 (Backup—Toll Number)
International telephone	https://www.hpe.com/psnow/doc/a50011948enw

Be sure to collect the following information before contacting Support:

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

Other useful sites

Other websites that can be used to find information:

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

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

Accessing Updates

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

Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

To subscribe to eNewsletters and alerts:

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

Warranty Information

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

Regulatory Information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at <https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts>

Additional regulatory information

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

Documentation Feedback

HPE Aruba Networking is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback

(docsfeedback-switching@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.