AOS-CX 10.17 Monitoring Guide

8100, 83xx, 8360, 93xx, 10000 Switch Series

Hewlett Packard Enterprise

Published: November 2025
Version: 1

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

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 8100 Switch Series (R9W94A, R9W95A, R9W96A, R9W97A)
- HPE Aruba Networking 8320 Switch Series (JL479A, JL579A, JL581A)
- HPE Aruba Networking 8325 Switch Series (JL624A, JL625A, JL626A, JL627A)
- HPE Aruba Networking 8325H Switch Series (S4B20A, S4B21A, S4B22A, S4B23A, S2T42A, S2T46A, S2T47A, S2T48A)
- HPE Aruba Networking 8325P Switch Series (S0G12A, S4A48A)
- HPE Aruba Networking 8360 Switch Series (JL700A, JL701A, JL702A, JL703A, JL706A, JL707A, JL708A, JL709A, JL710A, JL711A, JL700C, JL701C, JL702C, JL703C, JL706C, JL707C, JL708C, JL709C, JL710C, JL711C, JL704C, JL705C, JL719C, JL718C, JL717C, JL720C, JL722C, JL721C)
- HPE Aruba Networking 9300 Switch Series (R9A29A, R9A30A, R8Z96A, S0F81A, S0F82A, S0F83A)
- HPE Aruba Networking 10000 Switch Series (R8P13A, R8P14A)

Latest version available online

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

Command syntax notation conventions

Convention	Usage
`example-text`	Identifies commands and their options and operands, code examples, filenames, pathnames, and output displayed in a command window. Items that appear like the example text in the previous column are to be entered exactly as shown and are required unless enclosed in brackets (**[ ]**).
example-text	In code and screen examples, indicates text entered by a user.
Any of the following: ■ *`<example-text>`* ■ `<example-text>` ■ *`example-text`* ■ *example-text*	Identifies a placeholder—such as a parameter or a variable—that you must substitute with an actual value in a command or in code: ■ For output formats where italic text cannot be displayed, variables are enclosed in angle brackets (**< >**). Substitute the text—including the enclosing angle brackets—with an actual value.

Convention	Usage
	■ For output formats where italic text can be displayed, variables might or might not be enclosed in angle brackets. Substitute the text including the enclosing angle brackets, if any, with an actual value.
|	Vertical bar. A logical **OR** that separates multiple items from which you can choose only one. Any spaces that are on either side of the vertical bar are included for readability and are not a required part of the command syntax.
{ }	Braces. Indicates that at least one of the enclosed items is required.
[ ]	Brackets. Indicates that the enclosed item or items are optional.
... 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 8xxx, 93xx, and 10xxx Switch Series

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

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

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

Chapter 2

Monitoring hardware through visual

Diagnosing with the LEDs

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.

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*.

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 provided 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
packetDeltaCount
protocolIdentifier
tcpControlBits
sourceTransportPort

Private Information Elements
sourceVpcUUID (ID 769 ,1886 Pensando Systems, Inc.)
sourceSubnetUUID (ID 770,1886 Pensando Systems, Inc.)
flowLastTimeStamp (ID771,1886 Pensando Systems, Inc.)

Standard Information Elements
sourceIPv4Address
ingressInterface
destinationTransportPort
destinationIPv4Address
sourceIPv6Address
destinationIPv6Address
egressinterface
sourceIPv6Address
destinationIPv6Address
icmpTypeCodeIPv4
sourceMacAddress
vlanId
ipVersion
flowDirection
destinationMacAddress
octetTotalCount
packetTotalCount
forwardingStatus
droppedOctetDeltaCount
droppedPacketDeltaCount
droppedOctetTotalCount
droppedPacketTotalCount
flowEndReason
flowStartMilliseconds
flowEndMilliseconds
flowStartMicroseconds
flowEndMicroseconds

Private Information Elements
tcpReTransmitCount (ID 772,1886 Pensando Systems, Inc.)
destVpcUUID (ID 773,1886 Pensando Systems, Inc.)
destSubnetUUID (ID774,1886 Pensando Systems, Inc.)
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)
cngGlobalViewId (ID 495, 14823 Aruba)
cngGlobalBufferUsage (ID 496, 14823 Aruba)
cngSourceIP (ID 497, 14823 Aruba)
cngDestinationIP (ID 498, 14823 Aruba)

Standard Information Elements
udpSourcePort
udpDestinationPort
tcpSourcePort
tcpDestinationPort
tcpSequenceNumber
tcpAcknowledgementNumber
tcpOptions
dot1qVlanId
dot1qPriority
ingressInterface
ingressPhysicalInterface
egressPhysicalInterface
ethernetType

Private Information Elements
cngQueueDroppedOctetCount (ID 499, 14823 Aruba)
cngQueueRxOctetCount (ID 500, 14823 Aruba)
cngQueueBufferUsage (ID 501, 14823 Aruba)
cngQueueId (ID 502, 14823 Aruba)
cngQueueDroppedPacketCount (ID 503, 14823 Aruba)
cngQueueRxPacketCount (ID 504, 14823 Aruba)

Private information elements 495-504 are exclulsively reported by a flow congestion monitor.

Data generated by a flow congestion monitor is exported simultaneously in two types of IPFIX data records. Each record type contains a subset of information elements, listed in the tables below.

All Information Elements related to flow congestion-monitor are only available on the 8325 W/S/P.

Global IE Number	Name	Type/Unit	Description
154	flowStartMicroseconds	Timestamp	Displays the timestamp in microseconds, in which the view was collected.
163	observedFlowTotalCount	Number	Displays the total number of flows observed in this view.
495*	cngGlobalViewId	Number	Associates Flow Views to Global View.
496*	cngGlobalBufferUsage	Percentage	Displays the global buffer pool utilization.

Flow View Record	Name	Type/Unit	Description
60	ipVersion	Number	Specifies flow IP version (4 or 6).

Flow View Record	Name	Type/Unit	Description
497*	cngSourceIP	IP Address	Specifies flow source IP address.
498*	cngDestinationIP	IP Address	Specifies flow destination IP address.
4	protocolIdentifier	Number	Specifies flow IP protocol number (e.g. 6=TCP, 17=UDP).
7	sourceTransportPort	Port	Specifies flow source transport port.
11	destinationTransportPort	Port	Specifies flow destination transport port.
252	ingressPhysicalInterface	HW Intf. ID	Specifies flow ingress physical interface.
253	egressPhysicalInterface	HW Intf. ID	Specifies flow egress physical interface.
1	octetDeltaCount	Bytes	Specifies flow byte counter delta during this view's time window.
2	packetDeltaCount	Packets	Specifies flow packet counter delta during this view's time window.
502*	cngQueueId	Queue	Specifies queue ID used by this flow.
500*	cngQueueRxOctetCount	Bytes	Specifies queue Rx byte counter delta during this view's time window.
504*	cngQueueRxPacketCount	Packets	Specifies queue Rx packet counter delta during this view's time window.
499*	cngQueueDroppedOctetCount	Bytes	Specifies queue byte drop count delta during this view's time window.
503*	cngQueueDroppedPacketCount	Packets	Specifies queue packet drop count delta during this view's time window.
501*	cngQueueBufferUsage	Cells (256B)	Specifies queue buffer pool utilitization.
495*	cngGlobalViewId	Number	Associates Flow Views to Global Views.

Information elements marked with an asterisk (*) use private enterprise number 14823.

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: {CSRFTOKEN}"
    -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 8100, 8325, 8325H, 8325P, 9300 and 10000 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, 8325P, 8325H, 8360, 10000 Switch series
Martian Address	Packet has invalid source or destination IP address, for example, **0.0.0.0** or **127.0.0.0**.	8325, 8325P, 8325H, 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

Drop reason	Example Cause	Supported Platforms
Vlan Mbr Chk Fail	The packet has a VLAN that is not part of the member VLANs for the ingress port.	9300 Switch series
L3 Header Error	The packet has one or more malformed/invalid L3 header fields.	9300 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 as 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 9300 Switch series:

If there is no data in the **drop group name** or **drop reason name** fields, there have been less than two drop reasons for the flow. Every time the drop reason changes, the next set of drop reason IEs will be populated. If more than two 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 8325, 8325P, 8325H, and 10000 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 8100 and 8360 Switch Series:

If the **drop end** timestamp has the same value as the **drop star**t 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. 8100 and 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. 5420, 8325, 8325H, 8325P, 10000, 9300, and 9300S switch series support one flow monitor and only one flow exporter can be applied to the flow monitor.

Flow monitor is mutually exclusive with other device features due to shared hardware resource use. During device initialization, a feature is enabled based on the precedence list below.

On the 8325 and 8325H switch series:

- Flow monitor
- Flow congestion-monitor

On the 8325P switch series:

- Precision Time Protocol (PTP)
- Flow monitor
- Flow congestion-monitor

The first feature in the precedence list, which is fully configured, will be enabled upon device initialization. If no feature is configured, the first feature in the list is enabled by default. If flow monitor is not enabled, applying a flow-monitor on an interface will result in a monitor status of Disabled (Blocked by a higher precedence feature) on the interface. Refer to the FAQs and Troubleshooting section for remediation steps.

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. On the 8325, 8325H, 8325P, 9300, 9300S and 10000 switch series, a maximum of one flow record 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

For 8100 and 8360 Switch series:

- 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. 8100 and 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. 5420, 8325, 8325H, 8325P, 10000, 9300 and 9300S switch series support one flow monitor and only one flow exporter can be applied to the flow monitor.

Flow Congestion Monitor Resource Use

A flow congestion monitor on 8325, 8325H and 8325P switch series is applied to an interface to perform network traffic monitoring at a high frequency. Flows are monitored as they egress configured queues of one or more ports. This feature helps detect microbursts on monitored queues and associate those bursts with specific flows. A maximum of one flow congestion monitor can be created.

Flow congestion monitor cannot operate while flow monitor is active.

Flow congestion-monitor is mutually exclusive with other device features due to shared hardware resource use. During device initialization, a feature is enabled based on the precedence list below.

On the 8325 and 8325H switch series:

- Flow monitor
- Flow congestion-monitor

On the 8325P switch series:

- Precision Time Protocol (PTP)
- Flow monitor
- Flow congestion-monitor

The first feature in the precedence list, which is fully configured, will be enabled upon device initialization. If no feature is configured, the first feature in the list is enabled by default. If flow congestion-monitor is not enabled, applying a flow congestion-monitor on an interface will result in a monitor status of Disabled (Blocked by a higher precedence feature) on the interface. Refer to the FAQs and Troubleshooting section for remediation steps.

IPFIX records

The flow congestion monitor scans all learned flows for updated counters every 1.6 milliseconds. If any counter delta for one or more learned flows is nonzero, the flow congestion monitor creates and exports a new **view**. A view represents a snapshot of global, per-queue, and per-flow data collected over the last 1.6 milliseconds. A view always consists of two things:

1. 1 Global View Record
2. 1-1000 Flow View Records (one per active flow)

All flow records in a single export cycle represent 1 view, and contain the same **view ID**. It is possible to match a Global View Record to one or more Flow View Records by joining the records on **view ID**. For example, if a collector receives a **Flow View Record** indicating a packet count of 100 packets and **view ID 43**, the collector may determine a precise timestamp for those 100 packets by locating the **Global View Record** with **view ID 43** and inspecting the **Global View Record's flowStartMicroseconds** field.

Flow statistics reflect the single flow described by this Flow View Record. **Queue** statistics reflect the sum of all flows using the same egress physical interface and queue as the flow described by this Flow View Record (including flows not monitored by this flow congestion monitor). **Global** statistics reflect the state of the entire device.

Table 1: *Global View Record Data*

IE#	Name	Type/Unit	Description

Table 2: *Flow View Record Data*

IE#	Name	Type/Unit	Description

IE#	Name	Type/Unit	Description

* Information elements with an asterisk (*) use private enterprise number 14823.

Destinations

The destination specifies where flow reports are sent. There is one possible destination for a flow exporter within the 8100, 8325, 8325H, 8325P, 8360, and 9300S switch series:

1. Hostname or IP address of a device with an optional (or default on the 8325, 8325H, and 8325P switch series) 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 supported on 5420, 6200, 6300, 6400, 8100, 8360, and 9300S Switch series. 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.

The following is the only possible type of destination for a flow exporter within the 10000 switch series:

- Traffic Insight instance

A destination type is required for the flow exporter configuration to be complete.

Configuring IP Flow Information Export on 8100, 8360, 8325, 8325H, 8325P, 9300, and 9300S 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 for an 8360 Switch series.

```
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 timestamp absolute first
switch(config-flow-record)# collect timestamp absolute last
```

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)# 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
```

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
```

A collecting process is not configured with IPFIX.

An intermediate collecting process is configured with IPFIX by:

1. Creating a flow collector.
2. Adding any additional information elements to be appended in software to the flow collector with the append keyword.
3. Adding the flow collector to a flow monitor which is applied on a port.

Switch software can act as an intermediate collecting process for flow reports from the hardware to append certain additional IPFIX information elements to the flow reports. When the switch software is configured, acts as an intermediate exporting process to export the augmented flow reports to any flow exporters that are configured.

Natively, the hardware supporting flow monitoring does not generate reports on flows that are dropped. The hardware only produces reports on forwarded flows.

The collecting process enables the switch to provide flow monitoring reports in the IPFIX format for flows that are dropped and provides the reason for this drop. For more information on how to configure monitoring for dropped flows, refer to configuration information in the **configure-flow-collector-append-fields** command regarding forwarding-status.

Configuring IP Flow Information Export on 10000 Switch Series

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

- **Metering Process:** This process monitors an observation point within the network for flows and generates flow reports for the flows. An observation point is a location in the network where packets can be observed. The metering process passes the flow reports to the exporting process. A metering process is configured with IPFIX by:

    1. Creating a flow record.

2. Creating a flow monitor.
3. Assigning the flow record to the flow monitor

- **Exporting Process:** This process sends flow reports generated by the metering process to a collecting process.
- **Collecting Process:** This process receives flow reports from the exporting process and stores it or further processes it.

Figure 1 *IPFIX to Traffic Insight Workflow*

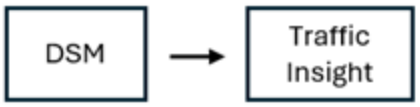

Metering and exporting processes are configured through AMD Pensando Policy and Services Manager (PSM), a cloud native application, that is connected to the device. IPFIX must be enabled under the DSM context and IP source-interface must be specified for IPFIX via the CLI or REST interface. The information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

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

Figure 2 *IPFIX Flow Collector Workflow*

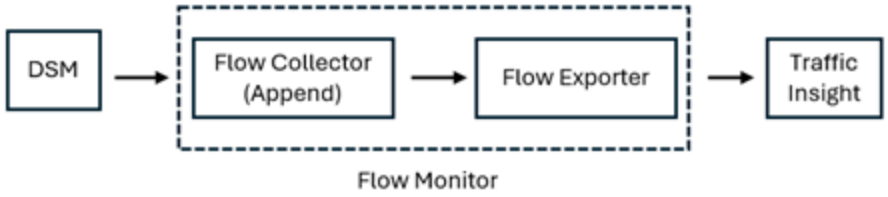

Hardware supporting flow monitoring does not generate reports on flows that are dropped, and instead only produces reports on forwarded flows. The collecting process enables the switch to provide flow monitoring reports in the IPFIX format for flows that are dropped along with the reason for this drop.

Refer to the **Configure flow collector append fields** section in this guide for more information on how to configure monitoring for dropped flows and forwarding status.

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

- Create a flow collector.

Configuring PSM with an export policy destined to an IP address on the device, for example, the IP of a local loopback interface.

- Add any additional information elements to be appended to the flow collector with the **append** keyword.

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

- Create a flow exporter.
- Configure the flow exporter with a destination of a traffic insight instance.

- Assign the flow exporter to the previously-created flow monitor.
- Applying the flow monitor to an observation point for a particular traffic type and direction. IPFIX flow monitoring supports traffic type IPv4, and traffic direction ingress. Only physical interfaces and LAG interfaces can be monitored.
- An exporting process is configured with IPFIX by:
  - Creating a flow exporter.
  - Assigning the flow exporter to the flow monitor. IPFIX flow exporting supports IPv4 transport.

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

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

It is possible for multiple export policies to be applied through Policy and Services Manager (PSM. In a configuration where multiple export policies are exporting to IP addresses local to the switch, duplicate reports for the same flow will be combined, and packet and byte counts could be counted multiple times in the combined report.

Compatibility with Traffic Insight

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

FAQs and Troubleshooting

- On 8325, 8325P, and 9300S Switch series, IPFIX does not monitor unresolved IP unicast traffic. Any traffic still being ARP or neighbor-resolved that is received on an interface with IPFIX monitoring will not be learned and flow reports will not be generated for that traffic. Once the traffic has sbeen resolved, IPFIX will start reporting those flows. On the 8325 switch series, any ICMP traffic currently being received on an interface where IPFIX monitoring has been applied will not be learned and flow reports will not be generated. For the 9300 switch series, for ICMP flows, IPFIX only matches on the IPv4 version, IPv4 source address, IPv4 destination address, and IPv4 protocol. The ICMP type and code are not included as key fields. Flow reports for ICMP flows have transport destination port and transport source port set to 0.
- To enable flow congestion-monitor:
  1. Disable PTP in the "ptp" configuration context.
  2. Remove "flow monitor" configurations from all interfaces.
  3. Apply a flow congestion-monitor on one or more interfaces.
  4. Reboot the device.
- To enable flow monitor:
  - Disable PTP in the "ptp" configuration context.
  - Apply a flow monitor on one or more interface(s).
  - Reboot the device.

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

Flow monitoring commands

diag-dump ipfix basic

```
diag-dump ipfix basic
```

Description

Displays diagnostic information for IPFIX.

Examples

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

....

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

```
1_3217499901_1_14                0xffff91186e80   1

....

------------------------------------------------------------------------
[End] Daemon ops-switchd
------------------------------------------------------------------------
========================================================================
[End] Feature ipfix
========================================================================
Diagnostic-dump captured for feature ipfix
```

Command History

Release	Modification
10.16	Command introduced on 8325P Switch series.
10.15	Command introduced on 9300S and 6200 Switch series
10.14	Command introduced on 8325 and 10000 Switch series.
10.11	Command introduced on 6300, 6400, 8100 and 8360 Switch series.

Command Information

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 10000	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
- flow collector in the **config-flow-collector** context

The no form deletes a flow congestion monitor.

Parameter	Description
<DESCRIPTION>	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 collector, **flow-collector-1**:

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

```
switch(config-flow-exporter)# description Exports flows to Traffic Insight
```

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
8100	`config-flow-record`	Administrators or local user group members

Platforms	Command context	Authority

8325
8325P
8325H
8360
9300
10000
	`config-flow-monitor`
	`config-flow-congestion-monitor`
	`config-flow-exporter`
	`config-flow-collector`
		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.

A maximum of two flow exporters can be applied to each flow monitor, and one flow exporter can be applied to each flow congestion monitor.

A maximum of one flow exporter can be applied to each flow monitor, and one flow exporter can be applied to each flow congestion 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
8100 8325 8325P	`config-flow-monitor` `config-flow-congestion-monitor`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325H 8360 9300 10000		

flow collector

```
[no] flow collector <name>
   [no] listen <IPv4 address> [vrf <VRF-name>]
   [no] append egress interface
   [no] append egress queue
   [no] append forwarding-status
```

Description

Creates or modifies a flow collector.

A flow collector allows switch software to act as an intermediate collecting process for flows provided by hardware, appending additional information elements in these flows and then acting as an intermediate exporting process to export the augmented flow reports to any configured flow exporters. The **no** form of the command deletes the flow collector.

The AOS-CX 10000 series switches contain embedded Distributed Services Modules (DSMs) which can be configured via PSM to export flows. A flow collector allows the switch to act as an intermediate collecting process for flows, append additional information elements in these flows, and then act as an intermediate exporting process to export the augmented flow reports to any configured flow exporters.

A maximum of one flow collector can be configured.

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 collector, up to 64 characters. The special characters allowed in the name are **.**, **_**, and **-**.
`listen <IPv4 address>`	This parameter is deprecated as of AOS-CX 10.17.
`vrf [VRF-name]`	Name of the VRF. If a VRF is not specified, then the default VRF will be used.
`append egress interface` `append egress queue`	Configure fields which will be added to the collected flow reports. **NOTE:** Only one append field can be specified per line in a configuration.
`forwarding-status`	Configure flow forwarding-status to be appended in software.

Examples

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

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

The following example displays an error message when more than one flow collector is configured.

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

The following example configures an interface to listen for flows.

```
switch(config)# flow collector collector-1
switch(config-flow-collector)# listen 1.2.3.4 vrf vrf2
```

The following example adds egress interface to **collector-1** as an append field

```
switch(config)# flow collector collector-1
switch(config-flow-collector)# append egress interface
```

The following example displays an error message when more than one append field is configured.

```
switch(config)# flow collector collector-1
switch(config-flow-collector)# append egress interface
switch(config-flow-collector)# append egress queue
Remove an egress interface append field from flow collector **collector-1**
```

The append forwarding-status option is only available on the 8325 Switch Series.

The following example displays how to add forwarding-status to flow collector **flow-collector-1** as an append field

```
switch(config)# flow collector flow-collector-1
switch(config-flow-collector)# append forwarding-status
```

The following example displays how to remove a forwarding-status append field from flow collector **flow-collector-1**

```
switch(config)# flow collector flow-collector-1
switch(config-flow-collector)# no append forwarding-status
```

Command History

Release	Modification
10.17	In AOS-CX 10.16 and previous releases, a flow collector required a **listen** address configuration. That configuration has been deprecated in AOS-CX 10.17, and a flow collector is valid without a listen address.

Release	Modification
10.16	Command introduced on 8325P Switch series.
10.14	Command introduced.

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 10000	`config` `config-flow-collector`	Administrators or local user group members with execution rights for this command.

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.

Parameter	Description
`<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.
`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
8100 8360	16	Two flow exporters can be applied to a single flow monitor.
8325 8325H 8325P	1	One flow exporter can be applied to the flow monitor.
9300	1	One flow exporter can be applied to the flow monitor.
10000	1	One flow exporter can be applied to the flow monitor.

On the 8325, 8325H, 8325P, 9300 Switch series, the exporter can be configured only with IPv4 address destinations.

Examples

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

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

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 attempts to create more than the maximum number of allowed flow exporters

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

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
10.16	Command introduced on 8325P Switch series.
10.15	Command introduced on 9300, 9300S and 6200 Switch series
10.14	Command introduced on 10000 and 8325 Switch series.
10.11	Command introduced on 6300, 6400, 8100 and 8360 Switch series.

Command Information

Platforms	Command context	Authority
8100 8325 8325H	`config` `config-flow-exporter`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 9300 10000		

flow monitor

```
flow monitor <name>
    exporter <name>
    collector <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.

On an HPE Aruba Networking 10000 Switch series, a flow monitor is applied to the system to define an intermediate collecting and exporting process for flow reports generated by the Distributed Services Module (DSM). A flow monitor consists of a flow collector and flow exporter. Flow reports are collected from the flow collector's listen address, appended with fields defined in the flow collector, and exported by the flow exporter assigned to the flow monitor.

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. HPE Aruba Networking 8325, 8325H 8325P, 9300S,10000 Switch series support one flow monitor and only one flow exporter can be applied to the 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. **NOTE**: This parameter is not supported on the HPE Aruba Networking 10000 Switch series.

Parameter	Description
`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 8325, 8325H, 8325P and 9300 Switch Series, the supported inactive timeout range is 30-120 seconds, and the default is 30 seconds. 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. **NOTE**: This parameter is not supported on the HPE Aruba Networking 10000 Switch series.
`description`	A description up to 256 characters long, including spaces.
`exporter <name>`	Assign a flow exporter to a flow monitor.
`collector <name>`	For HPE Aruba Networking 8325, 8325H, 8325P, and 10000 Switch series, assign a flow collector to a flow monitor. This command will override any configuration of "**traffic-insight flow-collector**" on the associated interface. Only one flow collector can be applied to each flow monitor. **NOTE:** A flow collector can be assigned to be a flow monitor only when a valid **listen address** is configured.
`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 example assigns **collector-1** as the collector associated with this monitor.

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

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.
```

Create more than the maximum number of allowed flow monitors

```
switch(config)# flow monitor monitor-1
switch(config)# flow monitor monitor-2
No more than 1 flow monitor 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
switch(config-flow-exporter)# description Exports flows to Traffic Insight
```

Command History

Release	Modification
10.16	Command introduced on 8325P Switch series.
10.15	Command introduced on 9300S and 6200 Switch series
10.14	Command introduced on 8325 and 10000 Switch series.
10.11	Command introduced on 6400, 6400, 8200 and 8360 Switch series.

Command Information

Platforms	Command context	Authority
8100 8325 8325H	`config` `config-flow-monitor`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325P 8360 10000		

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
      forwarding-status
      drop ingress-exceptions
      egress interface
      egress queue
      counter {packets|bytes}
      ingress interface
      datalink mac {source | destination} address in
      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 ■ i**p 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 **NOTE:** The HPE Aruba Networking 8325, 8325H, 8325P and 9300 Switch series support source, destination, and protocol matching on IPv4 networks only.
`description`	A description for the flow record up to 256 characters long, including spaces.
`collect`	Configures data fields to be included 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. For 8100 and 8300 switches, Byte count represents the number of incoming bytes since the previous report. For 8325 switches, byte count represents the total number of incoming bytes since the flow started. ■ **counter packets**: Collect counter data for packets in the flow. For 8100 and 8300 switches, Packet count represents the number of incoming packets since the previous report. For 8325 and 8325P switches, packet count represents the total number of incoming packets since the flow started. ■ **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. ■ **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. ■ **ingress interface:** Add ingress interface as a collect field to flow record ■ **datalink mac {source | destination} address in**: Configure a MAC source or destinatio address to monitor inbound traffic in a flow.

Examples

Add ingress packet destination MAC address as a collect field to flow record **flow-record-1.**

```
switch(config)# flow record flow-record-1
switch(config-flow-record)# collect mac destination address input
```

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

```
switch(config)# flow record flow-record-1
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.
```

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
10.17	The **collect drop ingress-exceptions** parameter is introduced on 5420, 6200, 6200, 8325H, 9300 and 10000 Switch series. Support for **egress vlan** and **egress interface** parameters is added on 8100 and 8360 Switch series.
10.16	Command introduced on 8325P and 9300 Switch series.
10.16	The **drop ingress-exceptions** parameter was added for 6300 and 8325 Switch series.
10.15	Command introduced on 6200 and 9300S Switch series
10.14	The **ipv4** parameter is deprecated and replaced with **ip**. Command introduced on the 8325 Switch series.
10.11	Command introduced on 6300, 6300, 8100 and 8360 Switch Series.

Command Information

Platforms	Command context	Authority
8100 8325 8325P 8360 9300	`config` `config-flow-record`	Administrators or local user group members with execution rights for this command.

flow congestion-monitor

```
flow congestion-monitor <MONITOR-NAME>
no flow congestion-monitor <MONITOR-NAME>
```

Description

Creates or modifies a flow congestion monitor.

The **no** form of this command deletes a flow congestion monitor.

Parameter	Description
<MONITOR-NAME>	Specifies the name of the flow congestion monitor. The name can be comprised of up to 64 alphanumeric, underscore, hyphen and period characters.

A maximum of one flow congestion monitor can be created.

Examples

Creating a flow congestion monitor named **congestion-monitor-1**:

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

Deleting the flow congestion monitor named **congestion-monitor-1**:

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

An error message is displayed if attempting to create more than the maximum allowed number of flow congestion monitors:

```
switch(config)# flow congestion-monitor monitor-1
switch(config)# flow congestion-monitor monitor-2
Unable to add another flow congestion-monitor; the limit is 1.
```

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 **congestion-monitor-1**:

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

Command History

Release	Modification
10.16	Command introduced on 8325 and 8325P

Command Information

Platforms	Command context	Authority
8325 8325P	config	Administrators or local user group members with execution rights for this command.

ingress-interface

```
ingress interface <IFRANGE>
no ingress interface <IFRANGE>
```

Description

Adds or removes monitored ingress interfaces when in the **config-flow-congestion-monitor** context.

The **no** form of this command stops monitoring one or more ingress interfaces.

Parameter	Description
<IFRANGE>	One or more comma - or hyphen-separated interfaces.

Usage

A flow congestion monitor only exports flows destined for monitored queues received on a monitored ingress (source) interface. If a flow does not use an ingress interface specified in the congestion monitor, or if the flow egresses on a queue that is not monitored, then the flow will not be exported.

A maximum of 52 ingress interfaces may be added to a flow congestion monitor.

Examples

Adding a single interface to the monitored list of ingress interfaces:

```
switch(config-flow-congestion-monitor)# ingress-interface 1/1/5
```

Adding multiple interfaces to the monitored list:

```
switch(config-flow-congestion-monitor)# ingress-interface 1/1/5,1/1/7-1/1/10
```

Removing a single interface from the monitored list:

```
switch(config-flow-congestion-monitor)# no ingress-interface 1/1/7
```

Removing multiple interfaces from the monitored list:

```
switch(config-flow-congestion-monitor)# no ingress-interface 1/1/7,1/1/9
```

Attempting to add more than the maximum allowed number of ingress interfaces:

```
switch(config-flow-congestion-monitor)# ingress-interface 1/1/1
A maximum of 52 ingress interfaces may be added to a flow congestion monitor
```

Command History

Release	Modification
10.16	Command introduced on 8325 and 8325P

Command Information

Platforms	Command context	Authority
8325 8325P	`config-flow-congestion-monitor`	Administrators or local user group members with execution rights for this command.

ip-all flow congestion-monitor

```
ip-all flow congestion-monitor <MONITOR-NAME> out queue <QUEUES>
no ip-all flow congestion-monitor <MONITOR-NAME> out queue <QUEUES>
```

Description

Enables flow monitoring on an interface when in the **config-if** or**config-lag-if** contexts. Flow congestion monitors are uniquely applied in the "out" direction to monitor outbound (egress) traffic on queues on an interface. A congestion monitor may be applied to one or more interface queues.

The **no** form disables the flow monitoring on an 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. The monitor must be applied on the LAG itself. If an interface is part of a LAG while attempting to apply a flow monitor, an error message will be displayed.

Parameter	Description
<MONITOR-NAME>	The name of the flow monitor.
out	Specifies monitoring of outbound (egress) traffic.
queue	Specifies that specific queues on this interface should be monitored.
<QUEUES>	One or more egress queues to monitor as a range or comma-separated list.

Usage

A maximum of 160 egress queues may be monitored by a congestion monitor across all interfaces in the system.

Examples

Enabling a flow congestion monitor on a physical interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip-all flow congestion-monitor congestion-monitor-1 out queue 3
```

Enabling a flow congestion monitor on all queues on a physical interface:

```
switch(config)# interface 1/1/1
switch(config-if)# ip-all flow congestion-monitor congestion-monitor-1 out queue
0-7
```

Attempting to enable a flow congestion monitor on more queues than are allowed across the system:

```
switch(config)# interface 1/1/1-1/1/23
switch(config-if)# ip-all flow congestion-monitor congestion-monitor-1 out queue
0-7
[1/1/21] A maximum of 160 egress queues can be monitored.
[1/1/22] A maximum of 160 egress queues can be monitored.
[1/1/23] A maximum of 160 egress queues can be monitored.
switch(config)# interface 1/1/30
switch(config-if)# ip-all flow congestion-monitor congestion-monitor-1 out queue 0
A maximum of 160 egress queues can be monitored.
```

Command History

Release	Modification
10.16	Command introduced on 8325 and 8325P switch series

Command Information

Platforms	Command context	Authority
8325 8325P	`config` `config-flow-exporter`	Administrators or local user group members with execution rights for this command.

ip-all flow monitor

```
[no] ip-all flow monitor <name>
```

Description

Enables flow monitoring for all flows passing through Distributed Services Modules (DSMs) in the DSM configuration context. The flow reports arriving from DSM on the monitor's associated collector will be processed by the monitor and the appended fields specified in the monitor's collector will be added. Finally, it will be exported by the monitor's exporter.

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

Examples

The following example enables a flow monitor.

```
switch(config)# dsm
switch(config-dsm)# ip-all flow monitor flow-monitor-1
```

Command History

Release	Modification
10.14	Command introduced.

Command Information

Platforms	Command context	Authority
10000	`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.

An IPv6 flow monitor is not supported on 8325, 8325H, 8325P, and 9300S Switch Series.

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
10.16	Command introduced on 8325P Switch series.
10.15	Command introduced on 9300S and 6200 Switch Series
10.11	Command introduced on 6300, 6400 and 8325 Swtich Series

Command Information

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300	`config` `config-flow-monitor`	Administrators or local user group members with execution rights for this command.

show interface flow-monitor

`show interface [<IFRANGE>] flow-monitor`

Description

Displays flow-telemetry monitor status as enabled on one or all interfaces that have flow-telemetry monitor configured.

Parameter	Description
<IFRANGE>	Specifies the port identifier range.
`flow-monitor`	Displays IPFIX flow monitor information.

Examples

The following example shows no configured monitor:

```
switch# show interface 1/1/1 flow-monitor
switch#
```

The following example shows no monitor configured in any port in the system:

```
switch# show interface flow-monitor
switch#
```

The following example shows a configured flow monitor:

```
switch# show interface 1/1/2 flow-monitor
Interface 1/1/2
IPv4 Flow Monitor (ingress): monitor-1
     Status: Enabled
```

The following example shows a configured flow monitor with no specified port:

```
switch# show interface flow-monitor
Interface 1/1/2
IPv4 Flow Monitor (ingress): monitor-1
     Status: Enabled
```

The following example shows a configured flow monitor that failed due to a mutually exclusive enabled feature:

```
switch# show interface flow-monitor
Interface 1/1/2
IPv4 Flow Monitor (ingress): monitor-1
Status: Disabled (Blocked by a higher precedence feature)
```

The following example shows a configured flow monitor that failed due to invalid configuration:

```
switch# show interface flow-monitor
Interface 1/1/2
   IPv4 Flow Monitor (ingress): monitor-1
Status: Disabled (Invalid configuration)
```

The following example shows a previously accepted configuration and enabled flow monitor occurence when the configuration causes an applied flow monitor to go invalid:

```
switch# show interface flow-monitor
 Interface 1/1/2
   IPv4 Flow Monitor (ingress): monitor-1
Status: Enabled (Using previously accepted configuration)
```

Command History

Release	Modification
10.16	Command introduced.

Command Information

Platforms	Command context	Authority
8325H 8325P 8325S 8325W	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 flow-congestion-monitor

`show interface [<IFRANGE>] flow-congestion-monitor`

Description

Displays flow congestion monitor status as enabled on one or all interfaces with a flow congestion monitor configured.

Parameter	Description
<IFRANGE>	Specifies the port identifier range.
`flow-congestion-monitor`	Displays IPFIX flow congestion monitor information.

Examples

The following example shows no configured monitor:

```
switch# show interface 1/1/1 flow-congestion-monitor
switch#
```

The following example shows no monitor configured in any port in the system:

```
switch# show interface flow-congestion-monitor
switch#
```

The following example shows a configured flow monitor:

```
switch# show interface 1/1/2 flow-congestion-monitor
Interface 1/1/2
    Flow Congestion Monitor (egress queue 0-7): fqm-1
     Status: Enabled
```

The following example shows a configured flow monitor with no specified port:

```
switch# show interface flow-congestion-monitor
Interface 1/1/2
    Flow Congestion Monitor (egress queue 0-7): fqm-1
     Status: Enabled
Interface 1/1/5
    Flow Congestion Monitor (egress queue 0-3,5): fqm-1
     Status: Enabled
```

The following example shows a configured flow monitor that failed due to a mutually exclusive enabled feature:

```
switch# show interface flow-congestion-monitor
Interface 1/1/2
    Flow Congestion Monitor (egress queue 0-7): fqm-1
     Status: Disabled (Blocked by a higher precedence feature)
Interface 1/1/5
    Flow Congestion Monitor (egress queue 0-3,5): fqm-1
     Status: Disabled (Blocked by a higher precedence feature)
```

The following example shows a configured flow monitor that failed due to rejected configuration:

```
switch# show interface flow-congestion-monitor
Interface 1/1/2
   Flow Congestion Monitor (egress queue 0-7): fqm-1
    Status: Disabled (Flow congestion monitor is rejected)
Interface 1/1/5
   Flow Congestion Monitor (egress queue 0-3,5): fqm-1
    Status: Disabled (Flow congestion monitor is rejected)
```

Command History

Release	Modification
10.16	Command introduced.

Command Information

Platforms	Command context	Authority
8325H 8325P 8325S 8325W	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 flow collector

show flow collector <name>

Description

Displays flow collector configuration and status. If no collector name is specified, the output of this command displays information for all flow collectors.

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

- Accepted
- Rejected (Internal error: failed to process collector)
- Rejected (The configured listen address is missing or invalid)
- Rejected (The configured listen address not associated with a local interface)

Parameter	Description
<name>	Name of the flow collector.

Examples

The following example displays the configuration of a flow collector named **collector-1**.

```
switch# show flow collector collector-1
-------------------------------------------------------------------------------
Flow collector 'collector-1'
-------------------------------------------------------------------------------
Description             : Collects flows from DSM
Status                  : Accepted
Listen Address          : 1.1.1.1
Append Fields
   egress interface
   egress queue
```

The following example displays the configuration of all flow collectors:

```
switch# show flow collector
-------------------------------------------------------------------------------
Flow collector 'collector-1'
-------------------------------------------------------------------------------
Description             : Collects flows from hardware
Status                  : Accepted
Listen Address          : 1.1.1.1
Append Fields
forwarding-status
```

The following example displays the configuration of a flow collector with no listen address configured:

```
switch# show flow collector collector-1
-------------------------------------------------------------------------------
Flow collector 'collector-1'
-------------------------------------------------------------------------------
Description             : Collects flows from hardware
Status                  : Rejected (The configured listen address is missing or
invalid)
Listen Address          :
Append Fields
```

```
forwarding-status
```

Release	Modification
10.16	Command introduced on 8325P switch series
10.14	Command introduced.

Command Information

Platforms	Command context	Authority
8325 8325P 10000	Manager (#)	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.

Displays flow exporter configuration and status.

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)
- Pending (Route Resolution for destination in progress)
- Pending (Destination MAC Resolution in progress)
- Pending (Source interface Resolution in progress)

- Pending (Source Vlan resolution in progress)
- Rejected (Internal error: route resolution for destination failed)

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 the configuration of all flow exporters:

```
switch# show flow exporter
------------------------------------------------------------------------------
Flow exporter 'exporter-1'
 -----------------------------------------------------------------------------
Reports sent            : 0
```

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
10.16	Command introduced on 8325P switch series
10.15	Command introduced on 9300S and 6200 Switch series
10.14	Command supported on 8325 and 10000 Switch Series.
10.11	Command introduced on 6300, 6400, 8100 and 8630 Switch Series.

Command Information

Platforms	Command context	Authority
8100 8325 8325P 8360 10000	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)
- Possible status types for 8325, 8325H, 8325P and 10000 Switch series:
  - Rejected (The state of the assigned flow collector is rejected)

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	8100 8325 8325P 8325H 8360

Statistics name	Meaning	Switches that support this statistic
		9300
Flows Added	Total number of flows added to the flow cache for this flow monitor since it was created	8100 8325 8325P 8325H 8360 9300
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	8100 8325 8325P 8325H 8360 9300
Flows Aged	Number of flows removed from the flow cache for this flow monitor since it was created due to active or inactive timeout	8100 8325 8325P 8325H 8360 9300
Inactive Timeout	Number of flows removed from the flow cache for this flow monitor since it was created due to inactive cache timeout.	8100 8325 8325P 8325H 8360 9300
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.	8100 8360
Current Dropped Entries	Current number of dropped flows in the flow cache for this flow monitor	8325 8325P
Current Forwarded Entries	Current number of forwarded flows in the flow cache for this flow monitor	8325 8325P

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 for a flow monitor on a 5420, 9300, 9300S 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
Cache Configuration
Inactive Timeout    : 1800
Active Timeout      : 300
```

Display information for a flow monitor on a 8325, 8325P, or 8325H Switch series:

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

Display information for a flow monitor on an 10000 Switch series:

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

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
```

Display information and statistics for a flow monitor on a 9300S Switch series

```
show flow monitor monitor-1 statistics
--------------------------------------------------------------------------------
Flow monitor 'monitor-1'
--------------------------------------------------------------------------------
Current Entries          : 2
Flows Added              : 6
Total Flows Terminated   : 4
  Flows Aged             : 2
    Active Timeout       :
    Inactive Timeout     : 1
  Forced End             :
```

Display information and statistics for a flow monitors on a 9300 Switch series:

```
switch# show flow monitor monitor-1 statistics
--------------------------------------------------------------------------------
Flow monitor 'monitor-1'
--------------------------------------------------------------------------------
Current Entries          : 2
Flows Added              : 6
Total Flows Terminated   : 4
Flows Aged               : 4
Inactive Timeout         : 4
```

Display information for a flow monitor on a 8325 and 8325P Switch series:

```
switch# show flow monitor monitor-1 statistics
--------------------------------------------------------------------------------
Flow monitor 'monitor-1'
--------------------------------------------------------------------------------
Current Entries             : 2
Flows Added                 : 6
Total Flows Terminated      : 4
Flows Aged                  : 4
Inactive Timeout            : 4
Current Dropped Entries     : 4
Current Forwarded Entries   : 0
```

Display information and statistics for a flow monitor on a 8325H Switch series:

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

```
Flow monitor 'monitor-1'
-------------------------------------------------------------------------------
Current Entries           : 2
Flows Added               : 6
Total Flows Terminated    : 4
Flows Aged                : 4
Inactive Timeout          : 4
```

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

Command History

Release	Modification
10.16	Command introduced on 8325P switch series
10.15	Command introduced on 9300S and 6200 Switch series
10.14	Command introduced on 8325 and 10000 Switch series.
10.11	Command introduced on 6400, 6400, 8100 and 8360 Switch series.

Command Information

Platforms	Command context	Authority
8100 8325 8325P 8325H 8360 9300 10000	`config` `config-flow-monitor`	Administrators or local user group members with execution rights for this command.

show flow congestion-monitor

`show flow congestion-monitor [<MONITOR-NAME>] [vsx-peer]`

Description

Displays flow congestion monitor configuration and status. When no congestion monitor name is provided, all flow congestion monitors are displayed.

Parameter	Description
<MONITOR-NAME>	Identifies the name of a flow congestion monitor.
`vsx-peer`	If present, the command will display the flow collector configuration for the VSX.

Usage

Possible statuses for a flow congestion monitor:

- Accepted
- Rejected
  - internal error: congestion monitor does not exist
  - feature pack is invalid or missing
  - one or more ingress interfaces must be assigned to the congestion monitor
  - an exporter must be assigned to the mointor, but no exporter is assigned
  - the state of one or more of the assigned flow exporters is rejected

Examples

Displaying information for all flow congestion monitors:

```
switch# show flow congestion-monitor
------------------------------------------------------------------------------
Flow congestion monitor 'congestion-monitor-1'
------------------------------------------------------------------------------
Description                  : Monitors tx on some interfaces
Status                       : Accepted
Flow Exporter                : exporter-1
Monitored Ingress Interfaces : 1/1/5,1/1/8-1/1/10
Monitored Queue Count        : 4 (limit 160)
Monitored Egress Interfaces and Queues:
1/1/4 queue 2,4,5
1/1/7 queue 7
```

Displaying information with no flow congestion monitors configured:

```
switch# show flow congestion-monitor
No flow congestion monitors configured.
```

Displaying information for a specific flow monitor:

```
switch# show flow congestion-monitor congestion-monitor-1
------------------------------------------------------------------------------
Flow congestion monitor 'congestion-monitor-1'
------------------------------------------------------------------------------
Description                  : Monitors tx on some interfaces
Status                       : Accepted
Flow Exporter                : exporter-1
Monitored Ingress Interfaces : 1/1/5,1/1/8-1/1/10
Monitored Queue Count        : 4 (limit 160)
Monitored Egress Interfaces and Queues:
1/1/4 queue 2,4,5
1/1/7 queue 7
```

Display information when too many egress queues are monitored:

```
switch# show flow congestion-monitor congestion-monitor-1
------------------------------------------------------------------------------
Flow congestion monitor 'congestion-monitor-1'
------------------------------------------------------------------------------
Description                  : Monitors tx on some interfaces
Status                       : Accepted
Flow Exporter                : exporter-1
```

```
Monitored Ingress Interfaces : 1/1/5,1/1/8-1/1/10
Monitored Queue Count        : 162 (limit 160)
Monitored Egress Interfaces and Queues:
1/1/4 queue 0,1,2,3,4,5,6,7
1/1/5 queue 0,1,2,3,4,5,6,7
1/1/6 queue 0,1,2,3,4,5,6,7
...
1/1/24 queue 0,1 (not applied)
! Queue 0: A maximum of 160 egress queues can be monitored.
! Queue 1: A maximum of 160 egress queues can be monitored.
```

Command History

Release	Modification
10.16	Command introduced on 8325 and 8325P

Command Information

Platforms	Command context	Authority
8325 8325P	`config` `config-flow-exporter`	Administrators or local user group members with execution rights for this command.

show flow record

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

Description

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

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

- Accepted
- Rejected (Internal error: failed to process record)
- Rejected (Mix of IPv4 and IPv6 match fields is not allowed. Specify match fields of the same IP version (IPv4 or IPv6))
- Rejected (Incomplete match fields. The mandatory match fields are: version, source address, destination address protocol, transport destination port, and transport source port)

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
    counter bytes
    counter packets

    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
    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
        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,
Status                  : Rejected (Incomplete match fields. The mandatory match
fields are: 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
counter bytes
counter packets
```

Display information with invalid Feature Pack

```
switch# show flow record record-4
--------------------------------------------------------------------------------
Flow record  'record-4'
--------------------------------------------------------------------------------
Description             : Used for dropped traffic analysis
Status                  : Accepted
Match Fields
ipv4 destination address
ipv4 protocol
ipv4 source address
transport destination port
transport source port
Collect Fields
counter bytes
counter packets
drop ingress-exceptions (configured, but not applied due to missing Feature Pack)
```

```
switch# show flow record record-5
------------------------------------------------------------------------------
Flow record  'record-5'
------------------------------------------------------------------------------
Description             : Used for dropped traffic analysis
Status                  : Accepted
Match Fields
Collect Fields
drop ingress-exceptions (configured, but not applied due to missing Feature Pack)
```

Display information with no flow records configured

```
switch# show flow record
No flow records configured
```

Command History

Release	Modification
10.17	The output of this command includes **drop ingress-exceptions** information for 5420, 6200, 6200 8325H, 9300 and 10000 Switch series.
10.16	Command introduced on 8325P switch series
10.15	Command introduced on 6200 and 9300S Switch series
10.14	Command introduced on 8325 Switch series.
10.11	Command introduced on 6400, 6400, 8100, and 8360 Switch series.

Command Information

Platforms	Command context	Authority
8100 8325 8325P 8325H 8360 9300 10000	Manager (#)	Administrators or local user group members with execution rights for this command.

show running-config

```
show running-config
     flow [collector | exporter | monitor | congestion-monitor |  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
`collector`	Display flow collector configuration.
`exporter`	Displays flow exporter configuration.
`congestion-monitor`	Displays flow congestion monitor configuration.
`monitor`	Displays flow monitor 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 4000:0:0:0:0:0:0:2 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
```

Showing **show running-config flow**:

```
switch# show running-config flow
flow collector collector-1
    listen 1.1.1.1 vrf default
    append forwarding-status
flow exporter exporter-1
    description This is an exporter
    destination 40.0.0.2
    transport udp 2055
flow record record-1
    description This is a record
```

```
    match ipv4 destination address
    match ipv4 protocol
    match ipv4 source address
    match transport destination port
    match transport source port
    collect counter bytes
    collect counter packets
    collect timestamp absolute first
    collect drop ingress-exceptions
flow monitor monitor-1
    description This is a monitor
    exporter exporter-1
    record record-1
flow congestion-monitor cng-mon-1
    description This is a monitor
    exporter exporter-1
    ingress-interface 1/1/10
    ingress-interface 1/1/11
interface lag 1
    ip flow monitor monitor-1 in
interface 1/1/1
    ip flow monitor monitor-1 in
interface 1/1/2
    ip-all flow congestion-monitor cng-mon-1 out queue 4,5,6
```

Showing **show running-config-flow** on a 9300S switch:

```
switch# show running-config flow
flow exporter exporter-1
    description This is an exporter
    destination 40.0.0.2
    transport udp 2055
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 timestamp absolute first
    collect timestamp absolute last
flow monitor monitor-1
    description This is a monitor
    exporter exporter-1
    record record-1
interface lag 1
    ip flow monitor monitor-1 in
interface 1/1/1
    ip flow monitor monitor-1 in
```

Showing **show running-config flow** on a 10000 switch:

```
switch# show running-config flow
dsm
    ipfix
    ip-all flow monitor monitor-1
```

```
flow collector collector-1
    listen 1.1.1.1 vrf default
    append egress interface
    append egress queue
flow exporter exporter-1
    description This is an exporter
    destination type traffic-insight
    destination traffic-insight ti
    transport udp 2055
flow monitor monitor-1
    description This is a monitor
    collector collector-1
    exporter exporter-1
interface 1/1/1
    ip flow monitor monitor-1 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
flow monitor monitor-1
    description This is a monitor
    collector collector-1
    exporter exporter-1
    record record-1
flow monitor monitor-1
    description This is a monitor
    exporter exporter-1
    record record-1
flow monitor monitor-1
    description This is a monitor
    collector collector-1
    exporter exporter-1
```

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
8100 8325 8325P 8325H 8360 9300 10000	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.

For 10000 Switch Series, the exported flows show the source IP address of the effective configured source in the default vrf.

For 8325 and 8325P Switch Series, the exported flows to an external collector shows the source IP address of the effective configured source in the default vrf. The exported flows to an internal collector does not utilize any source interface configuration.

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
10.16	Command introduced on 8325P switch series
10.15	Command introduced. on 9300S and 6200 Switch Series
10.14	Command introduced on 8325 and 10000 Switch series.
10.11	Command introduced on 6400, 6400, 8100, and 8360 Switch series.

Command Information

Platforms	Command context	Authority
8100 8325 8325P 8360 10000	Manager (#)	Administrators or local user group members with execution rights for this command.

Chapter 4
Queue Monitoring

AOS-CX provides multiple features that assist with monitoring queuing behavior such as Queue Monitoring and an extension of it – Congestion Event Detection.

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 drops

Queue drops is the number of outbound packets that could not be added to a queue due to congestion during a specific time window. This metric monitors changes in a queue's drop count since the last aggregation cycle and generates an event when the increase exceeds the configured threshold trigger value. The event clears when the drop count change falls below the threshold reset value.

Average queue tx rate

Represents the average transmit rate of a queue over a defined time window. This metric monitors the queue's transmit rate and triggers an event when the average rate since the last aggregation cycle exceeds the configured threshold trigger value. The event clears when the average transmit rate drops below the threshold reset value.

The percent line rate unit represents a percentage of the total line rate of the interface on which the congestion event applies.

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
10.15	Command introduced on 8100, 8325, 8325H, 8325P, 8360, 9300, 9300S and 10000 Switch series

Command Information

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 10000	`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
10.15	Command introduced on 8100 and 8360 Switch series
10.14.1000	Command introduced on 9300 Switch series
10.14	Command introduced on 8325 and 10000 Switch series

Command Information

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 10000	`config-if`	Administrators or local user group members with execution rights for this command.

queue-monitor polling-interval seconds

```
queue-monitor polling-interval seconds <SEC>
no queue-monitor polling-interval seconds <SEC>
```

Description

Configures the polling interval used to sample interface queue data. The specified interval is applied to all interfaces where queue monitoring is enabled. Changing the polling interval clears all previously collected queue statistics and restarts data collection from the time the new interval is applied.

The **no** form of the command resets the polling interval to the default value of 10 seconds.

Parameter	Description
`polling-interval seconds <SEC>`	Specifies the polling interval in seconds. Valid polling intervals are 1, 5, 10, 30, or 60 seconds.

Examples

Setting the polling interval to **1** second:

```
switch(config)# queue-monitor polling-interval seconds 1
```

Changing the polling interval while queue monitoring is enabled on one or more interfaces:

```
switch(config)# queue-monitor polling-interval seconds 20
Warning: changing the polling interval will clear collected queue monitor
statistics and 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
```

Resetting the polling interval to the default of **10** seconds when no interfaces have queue monitoring enabled:

```
switch(config)# no queue-monitor polling-interval
```

Resetting the polling interval to the default of **10** seconds when one or more interfaces have queue monitoring enabled:

```
switch(config)# no queue-monitor polling-interval
Warning: changing the polling interval will clear collected queue monitor
statistics and 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
```

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 8360v2 9300 10000	config	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}|queue-tx-rate|queue-drops}]
  list [filter [since <1-10000> days|hours|minutes|seconds][threshold {kbps <1-
      4294967295>}|{kbytes <1-4294967295>}|{packets <1-4294967295>}] [type {queue-depth
      [queue-tx-rate|queue-drops]] [slot <slot-id>]
  table [filter [threshold {kbps <1-4294967295>}|{kbytes <1-4294967295>}|{packets <1-
      4294967295>}] [type {queue-depth}|queue-tx-rate|queue-drops]]
```

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

Parameter	Description
	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. ■ 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} |queue-tx-rate|queue-drops`	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. ■ **queue-tx-rate**: Show average queue transmission (tx) rate data. ■ **queue-drops**: Show queue drops data.
`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  1/1/9      4      Average Queue Tx Rate  54321 kbps

2023-08-08 17:11:20  1/1/8      4      Average Queue Tx Rate  1234 kbps

2023-08-08 17:11:20  1/1/9      4      Average Queue Tx Rate  54321 kbps

2023-08-08 17:11:30  1/1/8      4      Average Queue Tx Rate  1234 kbps

2023-08-08 17:11:30  1/1/9      4      Average Queue Tx Rate  54321 kbps
```

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
                        Queue Depth - bytes
Time (s)       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
                        Count at Depth
Depth (KB)        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 speficied 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.
```

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
10.15	Command introduced on 8100 and 8360 Switch series
10.14.1000	Command introduced on 9300 and 9300S Switch series
10.14	Command introduced on 8325 and 10000 Switch series

Command Information

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 10000	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 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-tx-rate, queue-drops
```

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-tx-rate, queue-drops

Interface 1/1/2

Status               : Enabled
```

```
Monitored Statistics : , queue-tx-rate, queue-drops
```

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-tx-rate, queue-drops

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
8100 8325 8325H 8325P 8360 9300 10000	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,

Command	Description
	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
10.15	Command introduced on 8100 and 8360 Switch series
10.14.1000	Command introduced on 9300 and 9300S Switch series
10.14	Command introduced on 8325 and 10000 Switch series

Command Information

Platforms	Command context	Authority
8100 8325 8325H 8325P 8360 9300 10000	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.

Congestion Event Detection

The Congestion Event Detection feature is an extension of Queue Monitoring that allows users to configure alert thresholds for queue depth, queue rate, and queue drops. Alerts can be viewed on-switch and can also be sent to the switch event log on a selective basis.

Congestion Event Detection is supported on 8325, 8325H, 8325P and 10000 switch series. This feature enables users to configure queue statistics thresholds which are actively checked by the switch on interfaces where monitoring and thresholds are applied. This mechanism leverages queue monitoring or a Flow Congestion-Monitor as the source of periodic queue statistics samples for comparison against the configured thresholds. If a statistic has crossed a threshold, information about the occurrence is saved within the switch and an optional action can also be taken at the time of the occurrence. Accumulated occurrences (events) can be viewed on the switch through the CLI or requested through the REST API interface.

Congestion event configuration

The congestion event detection feature is configured with the threshold condition(s) under which an event occurrence starts, ends, and is measured. Each congestion event profile is comprised of zero or more congestion event entries.

Congestion event profile

A congestion event profile entry has a **Valid** status if all the required fields are specified with values that pass validation rules, otherwise, the congestion event profile entry status shows as **Invalid**. Each congestion event profile entry must specify the following information:

- type
- trigger threshold
- reset threshold
- units for provided values
- retention group
- action (optional)

Using congestion events

A congestion event represents a single occurrence of conditions on an interface queue where a configured threshold was violated. Configured thresholds are specified as part of a congestion event profile and are applied on a per-interface basis along with queue monitoring. When a congestion event occurs, relevant data is collected and stored as an *event occurrence* within the congestion event retention group for later analysis. The switch provides 3 retention groups for storing event occurrences, to be used however desired. One effective method is to treat them as "low", "medium", and "high" severity event occurrence storage. This method allows to quickly check the switch for specific severity events using the **show congestion-event retention-group <NUMBER>** command.

Once a congestion event retention group has reached the limit of 1000 stored events, the oldest completed event occurrence in the group is discarded in order to store a new event.

Congestion event detection commands

apply congestion-event profile

```
apply congestion-event profile [<NAME>]
no apply congestion-event profile [<NAME>]
```

Description

Applies a congestion event profile to an interface.

The **no** form removes the congestion event profile from the interface.

Parameter	Description
<NAME>	Specifies the name of a congestion event profile. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).

Usage

In order for the congestion event profile entry associated with an applied congestion event profile to be effective, the interface needs to be configured for **queue-monitor**. Once the queue statistics are being monitored, apply the congestion event profile to the interface to begin the process of actively checking the monitored queue statistics for congestion events based on the thresholds configured in the applied

congestion event profile. User-specified congestion event profiles and profile entries may be modified and deleted while applied on an interface. If a congestion event profile or entry is modified while applied on an interface, any ongoing event occurrences will end immediately.

Examples

Applying congestion event profile, **prof1**, to an interface where the queues on the interface have a monitoring source:

```
switch(config-if)# apply congestion-event profile prof1
```

Applying non-existent congestion event profile, **prof22**, to an interface:

```
switch(config-if)# apply congestion-event profile prof22
Profile prof22 does not exist.
```

Applying congestion event profile, **prof1**, to an interface where the queues on the interface do not have a queue data monitoring source. Without a source for queue statistics, no congestion event threshold checking will be performed.

```
switch(config-if)# apply congestion-event profile prof1
Warning: Queue monitoring must also be enabled using the 'queue-monitor' command
for congestion events to be detected.
```

Removing any congestion event profile from an interface.

```
switch(config-if)# no apply congestion-event profile
```

Removing congestion event profile, **prof1**, from an interface:

```
switch(config-if)# no apply congestion-event profile prof1
```

Removing congestion event profile, **prof2**, from an interface where congestion event profile, **prof1**, is currently applied.

```
switch(config-if)# no apply congestion-event profile prof2
The profile to remove does not match the currently configured profile.
```

Command History

Release	Modification
10.16	Command introduced

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 10000	`config-if`	Administrators or local user group members with execution rights for this command.

congestion-event profile

```
congestion-event profile <NAME>
no congestion-event profile <NAME>
```

Description

Creates a new congestion event profile and switches to the **config-cng-prof** context for the profile.

The no form deletes the profile and removes it from all applied interfaces.

If the specified profile already exists, this command switches to the **config-cng-prof** context for the named profile.

Parameter	Description
`<name>`	Name of the congestion event profile, up to 64 characters.

Usage

A congestion event profile contains zero or more congestion event entries that specify the thresholds to use for identifying conditions of interest that should be saved as an event occurrence. Use **show congestion-event profile [NAME]** to view the status of all congestion event profiles or their settings.

Examples

Creating a congestion-event profile named **prof1**:

```
switch(config)# congestion-event profile prof1
```

Deleting a congestion event profile named **prof1**:

```
switch(config)# no congestion-event profile prof1
```

Creating a congestion event profile with an invalid name:

```
switch(config)# congestion-event profile prof##
Invalid name. Please enter a string of up to 64 alphanumeric,
underscore, hyphen, and period characters.
```

Command History

Release	Modification
10.16	Command introduced

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 10000	`config`	Administrators or local user group members with execution rights for this command.

event-config-id

```
event-config-id {<event-type>} | {trigger <units> <value>} | [reset <units> <value>] |
[retention-group <retention-group>] | [action <action> <action-parameter>]
no event-config-id [<event-type>] [trigger <units> <value>] | [reset <units> <value>] |
[retention-group <retention-group>] | [action <action> <action-parameter>]
```

Description

Creates a new congestion event profile entry associated with the current congestion event profile.

The **no** form deletes the congestion event entry from the congestion event profile.

Parameter	Description
`event-type`	Specifies the type of congestion event. Congestion event types include the following: ■ queue-depth ■ queue-drop ■ queue-tx-rate
`[trigger <units> <value>]`	Specifies the units and value for the queue statistic trigger threshold.
`[reset <units> <value>]`	Specifies the units and value for the queue statistic reset threshold. If not provided, the reset value and units will be assigned the same values as the trigger units and value. **Note**: Units for the trigger and reset values must be equal to or below the trigger value to have a valid congestion event entry.
`[retention-group <retention-group-id>]`	(Optional) Specifies the retention group used to store congestion event occurrences for the congestion event entry. If no retention group is provided, congestion event retention group 1 will be used.
`[action <action> <action-parameter>]`	Specifies an action to take when a congestion event occurs. Available actions and their corresponding parameters: ■ eventlog: Logs an event to the system eventlog which may be seen with show events.

Parameter	Description
	The action parameter specifies at which level to log a message: ■ info (informational level) ■ warn (warning level) ■ crit (critical level) **Note**: If no action is specified, no action is taken with event occurrences.

Examples

Configuring congestion event entries for queue-depth:

```
switch(config-cng-prof)# queue-depth trigger kbytes 3000 reset kbytes 2000
retention-group 2 action eventlog crit
switch(config-cng-prof)# 7 queue-depth trigger kbytes 5000 reset kbytes 1000
```

Removing the congestion event entry with ID 7:

```
switch(config-cng-prof)# no 7
```

Configuring congestion event entries for queue-drops:

```
switch(config-cng-prof)# 2 queue-drops trigger packets 20 reset packets 10
retention-group 3
```

Removing the congestion event entry with ID 2:

```
switch(config-cng-prof)# no 2
```

Configuring congestion event entries for queue-tx-rate:

```
switch(config-cng-prof)# 5 queue-tx-rate trigger kbps 200000 retention-group 1
action eventlog info
switch(config-cng-prof)# queue-tx-rate trigger pps 7800 reset pps 5000 retention-
group 3
```

Removing the congestion event entry with ID 5:

```
switch(config-cng-prof)# no 5
```

Command History

Release	Modification
10.16	Command introduced

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 10000	`config`	Administrators or local user group members with execution rights for this command.

show congestion-event profile

`show congestion-event profile [<NAME>]`

Description

Shows the status of all congestion event profiles, or the settings of a specific congestion event profile.

Parameter	Description
<NAME>	Specifies the name of a congestion event profile. Range: 1 to 64 alphanumeric characters, including period (.), underscore (_), and hyphen (-).

Usage

The status of each congestion event entry can be one of the following:

- **Valid** - the congestion event entry is valid.
- **Invalid** - the congestion event entry is invalid.

Use the **show congestion-event profile *<NAME>*** command to see details about each congestion event entry and reasons for rejection.

Examples

Showing the status of all congestion event profiles. The numbers provided in **Config Entry Status** field show the count of how many entries are currently in each status.

```
switch# show congestion-event profile
Name         Config Entry Status
----------   --------------------
prof1        2 Valid
prof2        1 Valid, 1 Invalid
prof3        -
```

Showing the status of congestion event profile **prof1**:

```
switch# show congestion-event profile prof1
Profile Name       : prof1
Applied Interfaces : 1/1/1-1/1/3
Retention
ID State   Event Type            Trigger       Reset         Group Action
-- ------  --------------        ---------     ----------    ----  ------
*1 Invalid Queue Depth           -             -              -    eventlog/crit
2  Valid   Queue Depth           12345 kbytes  12345 kbytes  2    -
```

```
3  Valid   Queue Drops           1234 pkts     123 pkts      1    -
*4 Invalid Queue Drops           600 pkts      50 pkts       -    eventlog/info
7  Valid   Average Queue Tx Rate 75% line rate 50% line rate 1    eventlog/error
8  Valid   Average Queue Tx Rate 12345 pps     1234 pps      2    -
9  Valid   Average Queue Tx Rate 54321 kbps    54321 kbps    2    -

*1: Missing value: trigger, reset, retention group.
*4: Missing value: retention group.
```

Command History

Release	Modification
10.16	Command introduced

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show interface congestion-event

```
show interface [<IFNAME>] congestion-event
```

Description

Displays the congestion event status on an interface.

Parameter	Description
<IFNAME>	Specifies the name of the interface.

Examples

Displaying the congestion event status on an interface that does not have a congestion event profile applied:

```
switch# show interface 1/1/1 congestion-event
Interface 1/1/1 does not have a congestion event profile applied.
```

Displaying the congestion event status on an interface that is not currently monitoring queue statistics and has a congestion event profile applied:

```
switch# show interface 1/1/1 congestion-event
Congestion Event Profile : cng_event_prof1
Queue 0-7
Aggregation Interval :
```

```
Data Source          :
Event Types          :
Status               : Not Applied (No queue monitoring source for event data)
```

Displaying the congestion event status on an interface that is down:

```
switch# show interface 1/1/1 congestion-event
Congestion Event Profile : cng_event_prof1
Queue 0-7
Aggregation Interval :
Data Source          :
Event Types          :
Status               : Not Applied (Interface is down)
```

Displaying the congestion event status on an interface where the applied congestion event profile contains only invalid entries:

```
switch# show interface 1/1/1 congestion-event
Congestion Event Profile : cng_event_prof1
Queue 0-7
Aggregation Interval :
Data Source          :
Event Types          :
Status               : Not Applied (2 invalid entries in congestion-event profile)
```

Displaying the congestion event status on an interface that is actively collecting queue statistics with queue monitoring, but some entries in the applied profile are invalid:

```
switch# show interface 1/1/1 congestion-event
Congestion Event Profile : cng_event_prof1
Queue 0-7
Aggregation Interval : 10s
Data Source          : queue-monitor
Event Types          : queue-depth, queue-drops, queue-tx-rate
Status               : Partially Applied (1 invalid entry in congestion-event
profile)
```

Displaying the congestion event status on an interface that is actively collecting queue statistics with queue monitoring:

```
switch# show interface 1/1/1 congestion-event
Congestion Event Profile : cng_event_prof1
Queue 0-7
Aggregation Interval : 10s
Data Source          : queue-monitor
Event Types          : queue-depth, queue-drops, queue-tx-rate
Status               : Applied
```

Displaying congestion event status on all interfaces:

```
switch# show interface congestion-event
Interface 1/1/1
Congestion Event Profile : cng_event_prof1
```

```
Queue 0-7
Aggregation Interval : 10s
Data Source          : queue-monitor
Event Types          : queue-depth, queue-drops, queue-tx-rate
Status               : Applied
Interface 1/1/2
Congestion Event Profile : cng_event_prof1
Queue 0-7
Aggregation Interval : 10s
Data Source          : queue-monitor
Event Types          : queue-depth, queue-drops, queue-tx-rate
Status               : Applied
```

Command History

Release	Modification
10.16	Command introduced

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show congestion-event

```
show congestion-event
   filter [since <time-value> | {seconds|minutes|hours|days}] | [state
   {ongoing|complete}] | [interface <ifname|ifrange>] | [queue <queues>][retention-group
   <retention-group>] | [type {queue-tx-rate|queue-drops|queue-depth}]
```

Description

Shows congestion events that have occurred.

Parameter	Description
`[since <time-value> {seconds|minutes|hours|days}]`	Filters the output by only displaying event occurrences that were updated equal to or later than the time given, where **<time-value>** and **{seconds|minutes|hours|days}** specifies how far backward to look for events, relative to the current time.
`[state {ongoing|complete}]`	Filters event occurrences based on whether the event occurrence is ongoing or complete (ended).

Parameter	Description
`[interface <ifname|ifrange>]`	Specifies the name of the interface or range of interfaces.
`[queue <queues>]`	Filters event occurrences based on the queue(s) where the event occurred. Range: 0 to 7.
`[retention-group <retention-group>]`	Filters event occurrences based on the retention group where they are stored.
`[type {queue-tx-rate|queue-drops|queue-depth}]`	Filters event occurrences based on the type of event that occurred.

Usage

For more detailed information on an event occurrence, use the **show congestion-event <id>** command.

Examples

Displaying all congestion event occurrences, but none matching the provided filters have occurred yet:

```
8325(config-cng-prof)# show congestion-event
Matching Event Count: 0
```
```
Display all of the congestion event occurrences.
```
```
switch# show congestion-event
Matching Event Count: 4

O - Ongoing
C - Complete

ID   Start time       Duration     Type                    Interface  Queue   Value
-------------------------------------------------------------------------------
---------
124  10 seconds ago   10s (O)      Queue Depth             1/1/1      3
123456 kbytes
122  20 minutes ago   20m (O)      Average Queue Tx Rate   1/1/2      1       9283
kbps
123  3 minutes ago    50m (C)      Queue Depth             1/1/4      6
728394 kbytes
121  1 hour ago       2m (C)       Queue Drops             1/1/7      7       6233
pkts
```

Displaying all congestion event occurrences that were updated in the last 5 minutes on **interfaces 1/1/1 and 1/1/4**:

```
switch# show congestion-event filter since 5 minutes interface 1/1/1,1/1/4
Matching Event Count: 2

O - Ongoing
C - Complete
```

```
ID   Start time       Duration   Type           Interface  Queue   Value
-------------------------------------------------------------------------------
--
124  10 seconds ago   10s (O)    Queue Depth     1/1/1      3       123456 kbytes
123  2 days ago       2m (C)     Queue Depth     1/1/4      6       728394 kbytes
```

Displaying all congestion event occurrences that are still ongoing:

```
switch# show congestion-event filter state ongoing
Matching Event Count: 2

O - Ongoing
C - Complete

ID   Start time      Duration  Type                 Interface Queue   Value
-------------------------------------------------------------------------------
--
124  10 seconds ago  10s (O)   Queue Depth           1/1/1     3       123456
kbytes
122  20 minutes ago  20m (O)   Average Queue Tx Rate 1/1/2     1       9283 kbps
```

Displaying all congestion event occurrences on **interface 1/1/4**:

```
switch# show congestion-event filter interface 1/1/4
Matching Event Count: 1

O - Ongoing
C - Complete

ID   Start time    Duration     Type           Interface  Queue   Value
-------------------------------------------------------------------------------
--
123  3 hours ago   10s (C)      Queue Depth    1/1/4      6       728394 kbytes
```

Displaying all congestion event occurrences on interface 1/1/1 and queue 3:

```
switch# show congestion-event filter interface 1/1/1 queue 3
Matching Event Count:1

O - Ongoing
C - Complete

ID   Start time       Duration     Type           Interface  Queue   Value
-------------------------------------------------------------------------------
--
124  10 minutes ago   10m (O)      Queue Depth    1/1/1      3       123456
kbytes
```

Displaying all congestion event occurrences stored in **retention group 3** and type **queue-depth**:

```
switch# show congestion-event filter retention-group 3 type queue-depth
Matching Event Count: 1

O - Ongoing
C - Complete
```

```
ID   Start time       Duration    Type           Interface  Queue   Value
-------------------------------------------------------------------------------
--
124  10 seconds ago   10s (O)     Queue Depth    1/1/1      3       123456 kbytes
```

If the congestion event entry includes an action to log to the system event log, display the event log to view start and end information for logged occurrences. Displaying event log:

```
switch# show events
2024-11-01T23:43:11.778434+00:00 8325 switchd_agent[3436]: Event|1502|LOG_
INFO|AMM|1/1|Congestion event 123: start - interface 1/1/1 queue 3 - queue-depth:
5100 kbytes
2024-11-01T23:43:21.999842+00:00 8325 switchd_agent[3436]: Event|1503|LOG_
INFO|AMM|1/1|Congestion event 123: end - interface 1/1/1 queue 3 - lasted 10
seconds with peak value 65656 kbytes
```

Command History

Release	Modification
10.16	Command introduced

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show congestion-event <id>

```
show congestion-event <id>
```

Description

Shows detailed output of a single congestion event occurrence.

Parameter	Description
<id>	Specifies the congestion event occurrence to display.

Examples

Displaying details of an ongoing Average Queue Tx Rate congestion event occurrence:

```
switch# show congestion-event 123
Congestion Event Occurrence 123
```

```
Event Configuration
    Data Source          : queue-monitor
    Aggregation Interval : 10s
    Type                 : Average Queue Tx Rate
    Trigger              : 12345 pps
    Reset                : 1234 pps
    Retention Group      : 2

Event Data
    Interface            : 1/1/1
    Queue                : 3
    Start Time           : 10 minutes ago (2024-03-09 10:02:24)
    End Time             : N/A (Ongoing Event)
    Duration             : 10m
    Peak Time            : 2024-03-09 10:11:34
    Peak                 : 123456 pps
```

Displaying details of a completed Queue Depth congestion event occurrence:

```
switch# show congestion-event 124
Congestion Event Occurrence 124

Event Configuration
    Data Source          : queue-monitor
    Aggregation Interval : 10s
    Type                 : Queue Depth
    Trigger              : 12345 kbytes
    Reset                : 1234 kbytes
    Retention Group      : 1

Event Data
    Interface            : 1/1/2
    Queue                : 5
    Start Time           : 20 seconds ago (2024-03-09 10:02:24)
    End Time             : 10 seconds ago (2024-03-09 10:02:34)
    Duration             : 10s
    Peak Time            : 2024-03-09 10:02:29
    Peak                 : 123456 kbytes
```

Displaying details of a completed Queue Drop congestion event occurrence:

```
switch# show congestion-event 125
Congestion Event Occurrence 125

Event Configuration
    Data Source          : queue-monitor
    Aggregation Interval : 10s
    Type                 : Queue Drops
    Trigger              : 12345 packets
    Reset                : 1234 packets
    Retention Group      : 1

Event Data
    Interface            : 1/1/2
    Queue                : 5
    Start Time           : 7 hours ago (2024-03-09 10:05:24)
    End Time             : 7 hours ago (2024-03-09 10:12:44)
    Duration             : 7m
```

```
    Peak Time             : 2024-03-09 10:05:29
    Peak                  : 123456 packets
    Sum                   : 234567 packets
```

Displaying the details of congestion event occurrence 125 when this congestion event occurrence ID is invalid:

```
switch# show congestion-event 125
Congestion event occurrence 125 is not valid.
```

An ID is invalid if the occurrence ID is not being used yet. If the ID *is* valid and displaying "invalid", the congestion event retention group deleted the occurrence to make room for new occurrences.

Command History

Release	Modification
10.16	Command introduced

Command Information

Platforms	Command context	Authority
8325 8325H 8325P 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

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). *Required*
L4-PORT	Specify the L4 Port (1-65535). *Optional*
PROTOCOL	Specify the transport protocol number (1-255) *Optional*
VRF-NAME	Specify the VRF. *Optional*
DOMAIN-NAME	Specify the Domain Name. *Optional*
APP-NAME	Specify the Application Name. *Optional*

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                  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        : FlowTrace_Capable           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   : lag20->1/3/2
Ingress VLAN       : 20                          Egress VLAN        : 30
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
```

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                   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        : FlowTrace_Capable           Device ID          : SGXXXXXXCL
System MAC         : f8:cc:cc:cc:cc: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        : 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          : SGXXXXXXBL
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          : SGXXXXXXBL
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 the 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
8100 8360	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 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 mistmatch 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",
```

```
                                "out-broadcast-pkts": "8918",
                                "out-multicast-pkts": "255712",
                                "out-octets": "22279174",
                                "out-pkts": "264630",
                                "resets": "1025"
                            },
                            "cpu": false,
                            "description": "AP1",
                            "enabled": true,
                            "ifindex": 1,
                            "logical": false,
                            "management": false,
                            "mtu": 1500,
                            "name": "1/1/1",
                            "oper-status": "UP",
                            "type": "iana-if-type:ethernetCsmacd"
                        }
                    }
                ]
            }
        }
    ]
}
```


Monitor all interface counters

Monitor all interface counters 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/state/counters --encoding json_ietf
```

Monitor all interface counters 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/state/counters
--mode once --encoding json_ietf --tls-version 1.2 -d
```

Example output:


```
{
        "source": "<your-switch-ip>:9339",
        "subscription-name": "default-1759857987",
        "timestamp": 1759859458737119840,
        "time": "2025-10-07T11:50:58.73711984-06:00",
        "prefix": "openconfig:interfaces",
        "updates": [
                {
                        "Path": "interface[name=1/1/11]/state/counters",
                        "values": {
                          "interface/state/counters": {
                             "openconfig-interfaces:carrier-transitions": "0",
                             "openconfig-interfaces:resets": "0"
                          }
                }
        },
```

```
        {
                "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
8100 8320 8325 8325H 8360 9300 10000	`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
8100 8320 8325 8325H 8360 9300 10000	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
8100 8320 8325 8325H 8360 9300 10000	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.

Chapter 5
Inband Flow Analyzer (IFA)

This feature is applicable only on AOS-CX8325 and 9300 Switch Series.

Inband Flow Analyzer (IFA) is a feature that allows the user to monitor a network for faults and performance issues. IFA samples a user-defined flow or set of flows of interest and generates special probe packets to collect telemetry data from an end-to-end path and per-hop data path information. For IFA to work, L3 traffic must be flowing through the network, therefore the configuration and policies for L3 traffic flow must be configured beforehand.

Only probe packets are supported.

Inband Flow Analyzer Process

Inband Flow Analyzer defines the following processing nodes:

- Initiator node
- Transit node
- Terminator node

An IFA zone has no restrictions on the types of nodes utilized between two hosts. Generally, there will be an initiator (ingress), a terminator (egress), and zero or more transit nodes.

The following example shows a simple IFA flow created between two hosts. This flow contains the three different nodes: initiator, transit and terminator nodes.

Figure 1 *Inband Flow Analyzer flow*

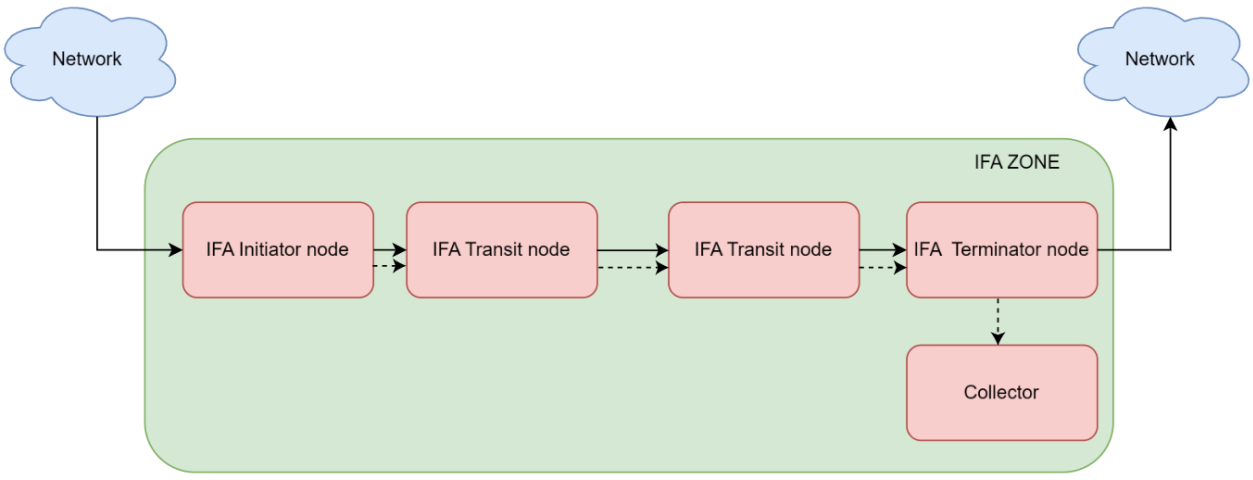

The following table describes the different functions for each node on an IFA flow.

Node	Functions
Initiator	■ Samples flow traffic. ■ Creates an IFA probe packet for each sample by adding the IFA header and metadata to it. ■ Changes probe packet protocol number to 253. ■ Supports: Unicast IPv4 and IPv6, UDP and TCP protocols, tagged and untagged packets, and Aggregation links (LAG) and multichassis LAG (MC-LAG). Both original packet and IFA probe sampled packed use same LAG port member to exit. **NOTE:** An IFA initiator configuration can only be applied in the context of a route-only port or on the global config context. IFA cannot initiate flows on non-routing interfaces. ■ On the 9300 and 9300S switch series, IP traffic coming in through an L2 bridge and matching a configured IFA class is initiated.
Transit	■ Identifies IFA probe packets. ■ Checks for probe packet validity. ■ Appends node metadata after the probe packet's IFA metadata header. ■ Forwards probe packet. An IFA packet that is L2 bridged will attach its local metadata.
Terminator	■ Moves the packet to the CPU, where its data is used to generate IFA metrics information. ■ An IFA packet that is L2 bridged attaches its local metadata. ■ Copies the IFA probe packet to the CPU, where its data is used to generate IFA metrics information. ■ Drops the probe packet.

For more information on the Inband Flow Analyzer feature, including use cases, best practices and troubleshooting information, refer to the video on the HPE Aruba Networking Airheads Broadcasting Channel.

Inband Flow Analyzer Packet Headers

An IFA probe packet contains the following headers:

- IFA Header
- IFA Metadata Header
- IFA Metadata Stack

At the initiator node, an IFA L3 probe packet format is as follows:

Figure 1 *IFA Header*

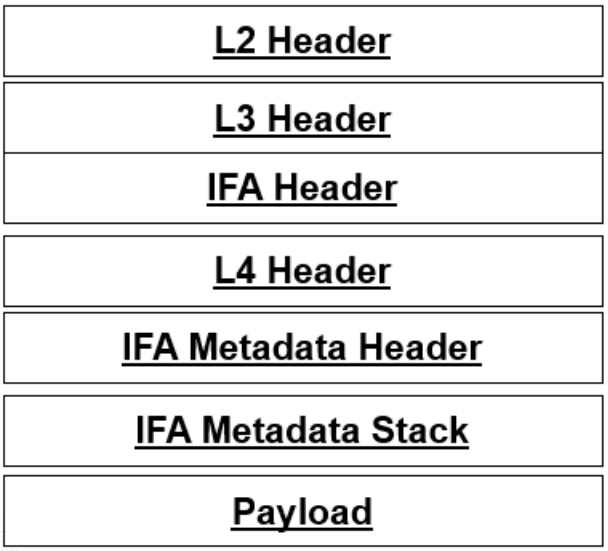

IFA Header

The IFA Header is added at the initiator node and contains the following fields:

Version (4bits)	GNS (4 bits)	Protocol Type (8bits)	Flags (8 bits)	Max Length (8bits)

IFA Header Field	Description
Version	Version of the IFA header. Set to 2.0.
GNS	Global Name Space. Set to 0xF.
Protocol Type	IP header protocol copied from the packet IP header.
Flags	Unused
Max Length	Maximum allowed length of metadata stack in multiples of 4 octets.

IFA Metadata Header

The IFA Metadata Header is also added at the initiator node and contains the following fields:

Figure 2 *IFA Metadata Header*

Request Vector (8bits)	Action Vector (8bits)	Hop Limit (8bits)	Current Length (8bits)

IFA Metadata Header Field	Description
Request Vector	Unused.
Action Vector	Unused.
Hop Limit	Maximum number of allowed hops in an IFA flow. This value is set by the initiator node. The hop limit is decremented at each transit node, the terminator node reduces the hop limit as well. If the hop limit of the probe packet is 0, the current node does not insert metadata.
Current Length	Current length of the metadata stack in multiples of 4 octets. If the length equals or exceeds the maximum length, the transit node stops inserting metadata.

IFA Metadata Stack

The IFA Metadata Stack is added at every node (initiator, transit and terminator) as long as the hop limit and length allows it. It contains the following fields:

Figure 3 *IFA Metadata Stack*

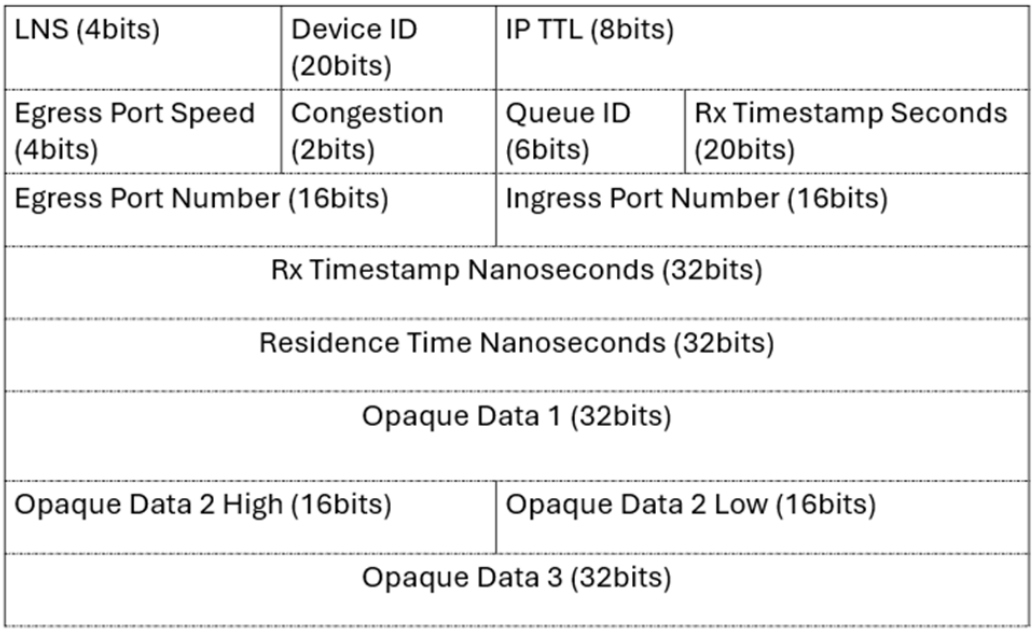

IFA Metadata Stack Field	Description
LNS	Local Name Space. Sets to 1.
Device ID	User-configured device ID that identifies an IFA node in a metadata stack.
IP TTL	IP time-to-live value at each hop.
Egress Port Speed	Egress port speed is mapped with the IFA metadata. Encodings are 0–10Gbps, 1–25Gbps, 2–40Gbps, 3–50Gbps, 4–100Gbps, 5–200Gbps, and 6–400Gbps.

IFA Metadata Stack Field	Description
	For example, if an egress port speed is 40Gbps, the speed field of the IFA packet is set to 2.
Congestion	Indicates whether the packet has experienced congestion. User must enable Explicit Congestion Notification (ECN) on the egress port.
Queue ID	Egress port queue ID.
Rx Timestamp Seconds	Received packet timestamp (in seconds).
Egress Port Number	Egress hardware port number.
Ingress Port Number	Ingress hardware port number.
Rx Timestamp Nanoseconds	Received timestamp in nanoseconds.
Residence Time Nanoseconds	Per-hop latency in nanoseconds.
Opaque Data 1	Egress queue transmission bytes.
Opaque Data 2 High	Reserved.
Opaque Data 2 Low	Depth of the packet queue in cells.
Opaque Data 3	Queue pool available in cells.

Configuration tasks list

To configure Inband Flow Analyzer (IFA), use the following commands:

To access the Flow Telemetry Profile context in the switch, use the following command:

```
switch(config)# flow-telemetry-profile
```

To configure the IFA device ID, use the following command:

```
switch(config-flow-telemetry-profile)# ifa-device-id {auto | <VALUE>}
```

To configure the IFA hop-limit, use the following command (Default: 10):

```
switch(config-flow-telemetry-profile)# ifa-hop-limit <VALUE>
```

To configure the IFA Max Metadata Stack length, use the following command (Default: 80):

```
switch(config-flow-telemetry-profile)# ifa-max-metadata-stack-length <VALUE>
```

IFA hop-limit and IFA Max Metadata Stack length are applicable only for initiator.

To configure the IFA sampling rate, use the following command:

```
switch(config-flow-telemetry-profile)# ifa-sampling-rate <VALUE>
```

To verify running configurations, use the following commands:

```
switch# show running-config

switch# show running-config all

switch# show running-config flow-telemetry-profile

switch# show running-config current-context

switch# show running-config flow ifa

switch# show running-config interface <INTERFACE_NAME>
```

To accesses the IFA Initiator monitor context in the switch and create a new IFA Initiator monitor, use the following command:

```
switch(config)# flow ifa-initiator-monitor <NAME>
```

To configure the flow filter class, use the following command:

```
switch(config-flow-ifa-initiator-monitor)# flow-filter-class <FILTER-CLASS-NAME>
```

To apply a flow IFA Initiator Monitor, use the following command:

```
switch(config)# {ip|ipv6} flow ifa-initiator-monitor <NAME>
```

To apply the IFA Transit monitor for all IP traffic, use the following command:

```
switch(config)# ip-all flow ifa-transit-monitor
```

To apply the IFA Terminator monitor for all IP traffic, use the following command:

```
switch(config)# ip-all flow ifa-terminator-monitor
```

To configure IFA IPv4 and IPv6 classes, use the following command:

```
switch(config)# class {ifa-ip|ifa-ipv6} <NAME>
```

To configure Class IPv4 or IPv6 filter match, use the following commands in the **config-class-ifa-ip** context:

```
<sequence_number>
    {match|ignore}
    {udp|tcp}
    {any|<src_ip_address>[/{<prefix_length>|<subnet_mask>}]}
```

```
    [any|{eq|gt|lt} <source_port_number>|range <port_number> <port_number>]
    {any|<dst_ip_address>[/{<prefix_length>|<subnet_mask>}]}
    [any|{eq|gt|lt} <destination_port_number>|range <port_number> <port_number>]
```

To verify flow IFA metrics, use the following commands:

```
switch# show flow ifa metrics
switch# show flow ifa metrics brief
```

To verify flow IFA flow monitor, use the following commands:

```
switch# show flow ifa monitor-status all
```

IFA support on VSX

- All VSX modes are supported: MCLAG, Active Gateway, and Active Forwarding.
- VSX Sync configuration is supported.
- All IFA roles are supported: Initiator, Transit, and Terminator.
- If both VSX members are configured as terminator:
  - IFA traffic is terminated on the first VSX member where it is received.
  - It is recommended to review the metrics on both VSX members.
- If an IFA flow needs to be initiated on the VSX, both VSX members must be configured as Initiator. If traffic goes through the ISL, the node must also be configured as Transit to capture its metrics. Otherwise, some packages will not be captured and flagged as IFA.
- Since each VSX member samples flow packets independently.
  - Initialization of IFA Flows depends on which VSX member receives the flow.
  - Sampling counts will start based on local IFA class entries settings, and the actual traffic traversing the VSX member.
- Both VSX members should be configured as Transit in order to include the device IFA metrics. This is required for traffic that could pass through the ISL. Otherwise, at the termination point, IFA metric information of the VSX member through which the sampled IFA packets traversed will be missed.

Supported scale

CoPP class

IFA copies packets to the CPU using the flow-telemetry class. This class has the following characteristics:

- Rate: 2048 kbps
- Priority: 0
- Burst size: 16 bytes

Initiators sampled packets rate

Packet samplers have a maximum rate defined for each platform. This will limit the amount of flows that can be monitored simultaneously. In other words, the total rate of all packets sampled by all initiators cannot exceed the rate of the sampler.

Following are the platforms supporting IFA and their maximum rates:

- 9300: 10 Gbps
- 9300S: 10 Gbps

Max number of flows tracked on a Terminator

Terminator nodes can display a maximum of 1000 flows as defined by the **ifa_max_flow_metrics** capability. The value of metrics is set to 1000. The sampler on the initiator samples packets from all the flows at the specified sampling rate. Due to this random nature, it becomes statistically difficult to monitor a large amount of flows simultaneously. For example, with a sampling rate of 1/5000 and 1000 flows, the probability of sampling a packet for each flow is 1/5000000.

Flows table memory usage

The IPFIX process measures data memory usage while tracking different number of flows. Measurements are taken for up to 250 flows, and values beyond this are estimated. For this measurement, a topology with three nodes is used.

Flows tracked	Memory used (MB)
0	6.59082
100	6.59472
200	6.60449
250	6.61767

Projected data:

Flows tracked	Memory used (MB)
350	6.62575
450	6.63494
550	6.64412
650	6.65330
750	6.66249
850	6.67167
950	6.68085
1000	6.68544

Max number of filter entries for all Initiators

The maximum amount of filter entries supported is 50. These can be distributed between several initiators. For example, a user can define one initiator with 50 class entries or 50 initiators with one class entry each.

This limit assumes an isolated environment where no other features are configured on the switch. This means all hardware traffic filtering resources are available. If other features are configured, configuring new IFA monitors may fail due to unavailable hardware resources. If the system has enough resources available, more monitors can be configured, but behavior is undefined.

IFA monitors memory usage

Memory used by 50 IFA Initiator monitors is measured.

Flows tracked	Memory used (MB)
0	6.60205
50	7.01953

Terminator flow table ageing

The flow table uses a basic aging mechanism that deletes the flow if they are not updated within 5 to 6 minutes. This means that if an IFA packet for a flow is not received during that time, the flow is removed from the table.

TCAM entry resources

When configuring the switch for Inband Flow Analyzer (IFA), it is important to consider Classifier Policy and Access Control Lists configurations, as they may interfere with the desired outcome.

Hardware support for IFA is enabled through the custom configuration of the policy TCAM ASIC module.

The TCAM module is used by multiple features for different purposes, typically involving matching traffic on some combination of header fields and taking the appropriate action if there us a match. This includes security features such as access lists (ACLs) and policies.

Because of its multi-functionality and energy requirements on the system, the TCAM is a limited, shared resource. Since IFA is enabled by programming rules in the TCAM ASIC module, creating new IFA monitors, new Initiation classes or any other TCAM user rules are dependent on the availability of TCAM resources, or **entries** at any given moment.

If the TCAM is near capacity or full, the ability to configure IFA monitors as desired will be limited. Conversely, if there are large quantities of IFA entries, are already configured, it will limit the ability of other TCAM users to acquire the resources that they need to create their own entries. The switch operator must find the correct balance of usage of the TCAM for all intended purpose.

For more information, see ACLs and Classifiers Policy Guide.

TCAM resource usage

Use the **show resources** command to monitor the allocated IFA TCAM entries.

IFA initiator monitor

```
switch# show resources

Resource Usage:

Mod   Description
          Resource                                  Width    Used Reserved    Free
-------------------------------------------------------------------------------
1/1   Global
      Total
          Destination Field Processor Entries           0        0    1024
1/1-0 Ports 9-12,21-24
          Ingress Control Plane Policing
              Ingress TCAM Entries                  3     435     6144
          Ingress IFA Initiator
              Ingress TCAM Entries                  3       9     6144
      Total
          Ingress TCAM Entries                          444    12288   12288
          Egress TCAM Entries                             0        0    2048
          VLAN Field Processor Entries                    0        0    1024
1/1-1 Ports 1-8
          Ingress Control Plane Policing
              Ingress TCAM Entries                  3     435     6144
          Ingress IFA Initiator
              Ingress TCAM Entries                  3       9     6144
      Total
          Ingress TCAM Entries                          444    12288   12288
          Egress TCAM Entries                             0        0    2048
          VLAN Field Processor Entries                    0        0    1024
1/1-2 Ports 13-20
          Ingress Control Plane Policing
              Ingress TCAM Entries                  3     435     6144
          Ingress IFA Initiator
              Ingress TCAM Entries                  3       9     6144
      Total
          Ingress TCAM Entries                          444    12288   12288
          Egress TCAM Entries                             0        0    2048
          VLAN Field Processor Entries                    0        0    1024
1/1-3 Ports 25-32
          Ingress Control Plane Policing
              Ingress TCAM Entries                  3     435     6144
          Ingress IFA Initiator
              Ingress TCAM Entries                  3       9     6144
      Total
          Ingress TCAM Entries                          444    12288   12288
          Egress TCAM Entries                             0        0    2048
          VLAN Field Processor Entries                    0        0    1024
```

IFA initiators	Qty TCAM entries	Total entries	Bank reserved
1 (first)	9		6144
	- 6 (General IFA Initiator)		
	- 3 (Sample flow)	9	
2	3 (Sample flow)	12	6144
3	3 (Sample flow)	15	6144
4	3 (Sample flow)	18	6144

IFA initiators	Qty TCAM entries	Total entries	Bank reserved
...			
498	3 (Sample flow)	1500	6144
499	3 (Sample flow)	1503	6144
500	3 (Sample flow)	1506	6144
...			
510	3 (Sample flow)	1536	6144

IFA transit monitor

```
switch# show resources
Resource Usage:

Mod   Description
          Resource                                  Width    Used Reserved    Free
-------------------------------------------------------------------------------
1/1   Global
      Total
          Destination Field Processor Entries           0        0    1024
1/1-0 Ports 1-20
          Ingress Control Plane Policing
              Ingress TCAM Entries              3     432     6144
          Ingress IFA Transit Terminator
              Ingress TCAM Entries              1       2     2048
      Total
          Ingress TCAM Entries                        434     8192    8192
          Egress TCAM Entries                           0        0    2048
          VLAN Field Processor Entries                  0        0    1024
1/1-1 Ports 21-40
          Ingress Control Plane Policing
              Ingress TCAM Entries              3     432     6144
          Ingress IFA Transit Terminator
              Ingress TCAM Entries              1       2     2048
      Total
          Ingress TCAM Entries                        434     8192    8192
          Egress TCAM Entries                           0        0    2048
          VLAN Field Processor Entries                  0        0    1024
```

IFA transit	Qty TCAM entries	Total entries	Bank reserved
1 (first)	2	2	2048

IFA terminator monitor

9300 Switch Series

```
switch# show resources

Resource Usage:

Mod   Description
```

```
          Resource                                Width    Used Reserved    Free
-------------------------------------------------------------------------------
1/1   Global
      Total
          Destination Field Processor Entries              0        0    1024
1/1-0 Ports 9-12,21-24
          Ingress Control Plane Policing
              Ingress TCAM Entries                3     435     6144
          Ingress IFA Transit Terminator
              Ingress TCAM Entries                1       2     2048
      Total
          Ingress TCAM Entries                          437     8192   16384
          Egress TCAM Entries                             0        0    2048
          VLAN Field Processor Entries                    0        0    1024
1/1-1 Ports 1-8
          Ingress Control Plane Policing
              Ingress TCAM Entries                3     435     6144
          Ingress IFA Transit Terminator
              Ingress TCAM Entries                1       2     2048
      Total
          Ingress TCAM Entries                          437     8192   16384
          Egress TCAM Entries                             0        0    2048
          VLAN Field Processor Entries                    0        0    1024
1/1-2 Ports 13-20
          Ingress Control Plane Policing
              Ingress TCAM Entries                3     435     6144
          Ingress IFA Transit Terminator
              Ingress TCAM Entries                1       2     2048
      Total
          Ingress TCAM Entries                          437     8192   16384
          Egress TCAM Entries                             0        0    2048
          VLAN Field Processor Entries                    0        0    1024
1/1-3 Ports 25-32
          Ingress Control Plane Policing
              Ingress TCAM Entries                3     435     6144
          Ingress IFA Transit Terminator
              Ingress TCAM Entries                1       2     2048
      Total
          Ingress TCAM Entries                          437     8192   16384
          Egress TCAM Entries                             0        0    2048
          VLAN Field Processor Entries                    0        0    1024
```

IFA terminator	Qty TCAM entries	Total entries	Bank reserved
1 (first)	2	2	2048

9300S Switch Series

```
switch# show resources

Resource Usage:

Mod   Description
          Resource                                Width    Used Reserved    Free
-------------------------------------------------------------------------------
1/1   Global
      Total
          Destination Field Processor Entries              0        0    1024
1/1-0 Ports 1-20
```

```
          Ingress Control Plane Policing
              Ingress TCAM Entries                3     435     6144
          Egress IFA Terminator Recirculate
              Egress TCAM Entries                 1       2      512
          Ingress IFA Transit Terminator
              Ingress TCAM Entries                1       3     2048
      Total
          Ingress TCAM Entries                          438     8192     8192
          Egress TCAM Entries                             2      512     1536
          VLAN Field Processor Entries                    0        0     1024
1/1-1 Ports 21-40
          Ingress Control Plane Policing
              Ingress TCAM Entries                3     435     6144
          Egress IFA Terminator Recirculate
              Egress TCAM Entries                 1       2      512
          Ingress IFA Transit Terminator
              Ingress TCAM Entries                1       3     2048
      Total
          Ingress TCAM Entries                          438     8192     8192
          Egress TCAM Entries                             2      512     1536
          VLAN Field Processor Entries                    0        0     1024
```

Ingress

IFA transit	Qty TCAM entries	Total entries	Bank reserved
1 (first)	3	3	2048

Egress

IFA transit	Qty TCAM entries	Total entries	Bank reserved
1 (first)	2	2	2048

Mirroring and sFlow

IFA uses a hardware mirror session that does not interfere with Mirroring or sFlow.

Updating class entries for an active IFA initiator monitor

When an IFA Initiator monitor is configured, the system will program the hardware using the class entries as configured at that moment. Any updates made to the class entries for the applied IFA initiators will be ignored, until the monitor is reconfigured. To make changes to class entries, it is required to:

1. Update the class entries using the **class {ifa-ip|ifa-ipv6} *<NAME>*** command.
2. Remove the IFA Initiator Monitor using the **no {ip|ipv6} flow ifa-initiator-monitor *[<NAME>]*** command.
3. Apply the IFA Initiator Monitor using the **{ip|ipv6} flow ifa-initiator-monitor *[<NAME>]*** command.

If an Initiator is applied without any Class, the same steps are required with the addition of configuring the flow filter Class, on the corresponding Initiator before re-applying the IFA Initiator Monitor.

This sections show an example of the steps to follow when configuring Class entry updates applied for a given IFA initiator.

Step 1: Check the current system configuration, and IFA nodes status **Show running-configuration** after configuring an IPv4 IFA Initiator Monitors to the system.

```
switch# show running-config flow ifa
Current configuration:
!
!
class ifa-ip class-ifa
  1 match udp any 10.10.10.10
flow ifa-initiator-monitor ifa-init
    flow-filter-class class-ifa
ip flow ifa-initiator-monitor ifa-init
!
!
switch# show flow ifa monitor-status all
*** System ***
initiator_ipv4
     Target State   : Active
```

Step 2: Update the class entries with the missing match/ignore entries.

```
switch(config)# class ifa-ip class-ifa
switch(config-class-ifa-ip)# 2 match tcp any 10.10.10.10
```

Step 3: Remove the IFA Initiator Monitor, and check the IFA node state until it is un-programmed:

```
switch(config)# no ip flow ifa-initiator-monitor ifa-init
switch(config)# exit
!
switch# show flow ifa monitor-status all
*** System ***
initiator_ipv4
     Target State   : Un-programming
!
!
switch# show flow ifa monitor-status all
switch#
```

Step 4: Apply the IFA Initiator Monitor and check its state.

```
switch(config)# ip flow ifa-initiator-monitor ifa-init
switch(config)# exit
!
switch# show running-config flow ifa
Current configuration:
!
!
class ifa-ip class-ifa
  1 match udp any 10.10.10.10
  2 match tcp any 10.10.10.10   ! New Class entry
flow ifa-initiator-monitor ifa-init
    flow-filter-class class-ifa
ip flow ifa-initiator-monitor ifa-init
!
```

```
!
switch# show flow ifa monitor-status all
*** System ***
initiator_ipv4
     Target State   : Programming
!
!
switch# show flow ifa monitor-status all
*** System ***
initiator_ipv4
     Target State   : Active
```

Important considerations

- Interoperability with other vendor switches is not supported.
- IFA is not supported for broadcast, unknown unicast and multicast traffic.
- IFA multicast flow sampling configuration is not supported.
- Only UDP or TCP protocols are supported.
- Initiator supports IPv4 and IPV6 unicast traffic with the following exceptions:
  - IPv6 with extension header is not supported.
  - TCP packets with options are not supported.
  - 1588 UDP frames are not supported.
- Initiating IFA traffic on a LAG interface is not supported.
- Tunneled traffic such as VXLAN and GRE, is not supported. Therefore, ECMP over VXLAN is not supported.
- The IFA initiator node ingress port will always report an incorrect HW ASIC ID.
- On an initiator node, when the probe packet egresses, it will contain an additional 40 bytes consisting of: 4 byte IFA header, 4 byte IFA Metadata Header and 32 byte IFA Metadata Stack.
- On transit nodes, a 32 byte IFA Max Metadata Stack is inserted at every hop, causing the probe packet size to increase after every hop. The hop-limit can be configured at the initiator node to ensure that the size of the IFA packets will not exceeds the specified value.
- IFA terminator nodes local metadata is not added to the probe packet. This is not applicable only for 9300S switch series.
- Hardware resources (TCAM and mirror sessions) are required for IFA functionality. If resources are exhausted, IFA will not work properly and failure status will be reflected in a show command on the CLI or via REST.
- User cannot configure different sampling rates for different flows on the IFA initiator interface.
- The IFA initiator node uses a sampling profile and a mirror instance. These resources will be removed from the system when an IFA initiator is started.
- To obtain end-to-end metrics, all devices traversed by the flow need to have the IFA feature enabled. Otherwise results can be skewed or undefined.
- When an IFA Initiator monitor is configured, the system programs the hardware using the class entries as configured at that time. Any updates made to the class entries for the applied IFA initiators will be ignored, until the monitor is reconfigured. For more details on updating class entries, see Updating class entries for an active IFA initiator monitor.

Debugging

To enable the debug logging, use the following command:

```
switch# debug flowtelemetry ifa
```

Inband Flow Analyzer (IFA) commands

class

```
class {ifa-ip | ifa-ipv6} <NAME>
no class {ifa-ip | ifa-ipv6} <NAME>
```

Description

Creates a new IFA class for IPv4 or IPv6 traffic.

The **no** form of this command removes an IFA class for IPv4 or IPv6 traffic.

Parameter	Description
`ifa-ip`	Specifies the IFA based Internet Protocol v4 (IFA-IPv4) .
`ifa-ipv6`	Specifies the IFA based Internet Protocol v6 (IFA-IPv6) .
<NAME>	Specifies the IFA class name.

Examples

Configuring an IFA IPv4 class:

```
switch(config)# class ifa-ip IFA_Class
```

Configuring an IFA IPv6 class:

```
switch(config)# class ifa-ipv6 IFA_IPv6_Class
```

Removing an IFA class for IPv4 or IPv6 traffic:

```
switch(config)# no class ifa-ip IFA_Class
```

```
switch(config)#  no class ifa-ipv6 IFA_IPv6_Class
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config`	Administrators or local user group members with execution rights for this command.

class ipv4/ipv6 filter

```
[<sequence_number>]
       {match|ignore}
       {udp|tcp}
       {any|<src_ip_address>[/{<prefix_length>|<subnet_mask>}]}
       [any|{eq|gt|lt} <source_port_number>|range <port_number> <port_number>]
       {any|<dst_ip_address>[/{<prefix_length>|<subnet_mask>}]}
       [any|{eq|gt|lt} <destination_port_number>|range <port_number> <port_number>]
no <sequence_number>
```

Description

Configures an IFA IPv4 or IPv6 Flow Filter Class based on 5-tuple. The IFA Flow Filter can filter by protocols UDP and TCP, source IP, source port, destination IP, and destination port.

The **no** form of this command removes the IFA Flow Filter Class based on the entry sequence number.

Parameter	Description
<sequence_number>	Specifies a sequence number for the match number. Range: 1 to 4294967295.
`udp`	Specifies the user datagram protocol.
`tcp`	Specifies the transmission datagram protocol.
<src_ip_address>	Specifies the source IP host address.
<prefix_length>	Specifies the source or destination prefix length.
<subnet_mask>	Specifies the source or destination network mask.
<dst_ip_address>	Specifies the destination IP host address.
`any`	Specifies any source or destination IP address.
`eq`	Specifies the Layer 4 port.
`gt`	Specifies the Layer 4 port greater than the indicated port.
lt	Specifies the Layer 4 port lesser than the indicated port.

Examples

Configuring an IFA IPv4 Flow Filter **match** with all the available parameters (5-tuple):

```
switch(config-class-ifa-ip)# 1 match tcp 10.10.10.10 eq 5 1.1.1.1 eq 10
```

Configure an IFA Flow Filter **ignore** with all the available parameters (5-tuple):

```
switch(config-class-ifa-ip)# 1 ignore tcp 10.10.10.10 eq 5 1.1.1.1 eq 10
```

Configuring an IFA IPv6 Flow Filter **match** with all the available parameters (5-tuple):

```
switch(config-class-ifa-ip)# 1 match tcp 10:10::10:10 eq 5 1:1::1:1 eq 10
```

Configuring an IFA IPv6 Flow Filter Match with source and destination IP:

```
switch(config-class-ifa-ip)# 1 match 10:10::10:10 1:1::1:1
```

Removing an IFA Flow Filter Class match:

```
switch(config-class-ifa-ip)# no 1
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config-class-ifa-ip`	Administrators or local user group members with execution rights for this command.

flow ifa-initiator-monitor

```
flow ifa-initiator-monitor <NAME>
no flow ifa-initiator-monitor
```

Description

Accesses the IFA initiator monitor context in the switch. This command creates a new IFA initiator monitor with the specified name.

The **no** form of this command exits the IFA initiator monitor context in the switch. It deletes the IFA initiator monitor with the specified name.

Parameter	Description
<NAME>	Specifies the IFA Monitor Initiator name.

Examples

Creating the Flow IFA initiator monitor **initiator-**1:

```
switch(config)# flow ifa-initiator-monitor initiator-1
```

Deleting the Flow Initiator monitor **initiator-2**

```
switch(config)# no flow ifa-initiator-monitor initiator-2
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config`	Administrators or local user group members with execution rights for this command.

flow-filter-class

```
flow-filter-class <NAME>
no flow-filter-class
```

Description

Assigns a Flow Filter Class to a specific IFA Initiator Flow Monitor. Each class can classify traffic based on the IPv4 or IPv6 header information, using special IFA IPv4 and IPv6 class types.

The **no** form of this command removes a Flow Filter Class from a specific IFA initiator monitor.

Parameter	Description
<NAME>	Specifies the flow filter class name.

Examples

Assigning a Flow Filter Class to an IFA Initiator Monitor on the system:

```
switch(config-flow-ifa-initiator-monitor)# flow-filter-class filter-class-1
```

Removing a Flow Filter Class:

```
switch(config-flow-ifa-initiator-monitor) # no flow-filter-class class-4
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config-flow-ifa-initiator-monitor`	Administrators or local user group members with execution rights for this command.

flow-telemetry-profile

```
flow-telemetry-profile
no flow-telemetry-profile
```

Description

Accesses the Flow Telemetry Profile context in the switch.

The **no** form of this command resets all the Flow Telemetry Profile configuration to defaults.

Examples

Accessing the Flow telemetry configuration context:

```
switch(config)# flow-telemetry-profile
```

Resetting all Flow Telemetry Profile settings to its defaults:

```
switch(config)# no flow-telemetry-profile
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config`	Administrators or local user group members with execution rights for this command.

ifa-device-id

```
ifa-device-id {auto| <VALUE>}
no ifa-device-id {auto| <VALUE>}
```

Description

Configures the IFA Device ID for the flow telemetry profile. The default setting is **auto** which configures the device identifier to match the low 20 bits of the switch's MAC address.

The **no** form of this command resets the IFA Device ID for the flow Telemetry Profile to the default value **auto**.

Parameter	Description
`auto`	Resets the IFA Device ID of Flow Telemetry Profile to default.
<VALUE>	Specifies the interval range for IFA Device ID. Range: 0 to 1048575.

Examples

Configuring the IFA Device ID on the system:

```
switch(config-flow-telemetry-profile)# ifa-device-id 10
```

Configuring IFA Device ID **auto** on the system:

```
switch(config-flow-telemetry-profile)# ifa-device-id auto
```

Resetting the IFA Device ID for flow telemetry profile to default value **auto**:

```
switch(config-flow-telemetry-profile)# no ifa-device-id 10
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config-flow-telemetry-profile`	Administrators or local user group members with execution rights for this command.

ifa-hop-limit

```
ifa-hop-limit <VALUE>
no ifa-hop-limit <VALUE>
```

Description

Configures the Inband Flow Telemetry (IFA) Hop Limit. The value used on IFA Metadata Header for created IFA probe packet running as the Initiator. Validate the number of IFA hops supported at the IFA Terminator node, before adjusting this value at the Initiator. In the case of AOS-CX, up to 10 IFA Hops are supported on a Terminator node.

The **no** form of this command resets the IFA hop limit to its default value.

Parameter	Description
<VALUE>	Specifies the IFA hop limit for the Flow Telemetry Profile. Default: 10

Examples

Configuring the IFA hop limit on the system:

```
switch(config-flow-telemetry-profile)# ifa-hop-limit 10
```

Resetting the IFA Hop Limit value to the default value:

```
switch(config-flow-telemetry-profile)# no ifa-hop-limit
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config-flow-telemetry-profile`	Administrators or local user group members with execution rights for this command.

ifa-max-metadata-stack-length

```
ifa-max-metadata-stack-length <VALUE>
no ifa-max-metadata-stack-length <VALUE>
```

Description

Configures the Inband Flow Telemetry Maximum Metadata Stack Length value in the IFA Metadata Header for created IFA probe packet running as the initiator. In the case of AOS-CX, only LNS 1 is supported. LNS 1 adds an IFA Metadata Stack with a size of 8 octets per hop.

Ensure that the IFA Max Metadata Stack Length aligns with the IFA hop limit value.

The **no** form of this command resets the IFA Max Metadata Stack Length for flow telemetry profile to its default value.

Parameter	Description
<VALUE>	Specifies the maximum metadata stack length numeric value in multiples of four octets. The default value is 80. Range: 8 to 248.

Examples

Configuring the IFA Max Metadata Stack Length value on the system:

```
switch(config-flow-telemetry-profile)# ifa-max-metadata-stack-length 240
```

Resetting the IFA Max Metadata Stack Length for flow telemetry profile to the default value **80:**

```
switch(config-flow-telemetry-profile)# no ifa-max-metadata-stack-length
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config-flow-telemetry-profile`	Administrators or local user group members with execution rights for this command.

ifa-sampling-rate

```
ifa-sampling-rate <VALUE>
no ifa-sampling-rate <VALUE>
```

Description

Configures the IFA sampling rate in multiples of a thousand packets to initiate cloned packets on filtered flows that match any configured initiator class, among all active IFA initiator monitors. For example, if the rate is set to 4, 1 out of every 4000 packets will be sampled.

The **no** form of this command resets the sampling rate to its default value.

Parameter	Description
<VALUE>	Specifies the IFA sampling rate. The default value is 4. Range: 4 to 4200000.

Examples

Configure the IFA sampling rate for the flow telemetry profile:

```
switch(config-flow-telemetry-profile)# ifa-sampling-rate 240
```

Resetting the IFA sampling rate for the flow telemetry profile to its default value **4**:

```
switch(config-flow-telemetry-profile)# no ifa-sampling-rate 240
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config-flow-telemetry-profile`	Administrators or local user group members with execution rights for this command.

ip-all flow ifa-terminator-monitor

```
ip-all flow ifa-terminator-monitor
no ip-all flow ifa-terminator-monitor
```

Description

Applies the IFA terminator Monitor for all IP traffic. The IFA terminator monitor configuration is applied globally to the system, monitoring both IPv4 and IPv6 traffic. Only one role (transit or terminator) can be actived at a time. It overwrites the other command **ip-all flow ifa-transit-monitor** in case both are executed.

The **no** form of this command unconfigures the flow IFA terminator Monitor for all IP traffic.

Examples

Applying the flow IFA terminator Monitor for all IP traffic:

```
switch(config)# ip-all flow ifa-terminator-monitor
```

Unconfiguring a Flow IFA terminator Monitor

```
switch(config)# no ip-all flow ifa-terminator-monitor
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config`	Administrators or local user group members with execution rights for this command.

ip-all flow ifa-transit-monitor

```
ip-all flow ifa-transit-monitor
no ip-all flow ifa-transit-monitor
```

Description

Applies the IFA Transit Monitor for all IP traffic probed by IFA. The IFA transit monitor configuration is applied globally to the system, monitoring both IPv4 and IPv6 traffic. Only one role (transit or terminator) can be actived at a time. It overwrites the other command **ip-all flow ifa-terminator-monitor** in case both are executed.

The **no** form of this command unconfigures the flow IFA Transit Monitor for all IP traffic.

Examples

Applying the IFA Transit Monitor for all IP traffic:

```
switch(config)# ip-all flow ifa-transit-monitor
```

Unconfiguring a Flow IFA Transit Monitor

```
switch(config)# no ip-all flow ifa-transit-monitor
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config`	Administrators or local user group members with execution rights for this command.

{ip|ipv6} flow ifa-initiator-monitor

```
{ip|ipv6} flow ifa-initiator-monitor <NAME>
no {ip|ipv6} flow ifa-initiator-monitor <NAME>
```

Description

Applies the IFA initiator monitor to the **config-if** or **config** context. IPv4 and IPv6 are supported. The monitor configuration should match the class IP version.

When a user has both system-level and port-level IFA initiator monitors, the port settings do not overwrite the system settings. Instead, both monitors work together. The system-initiator monitor samples packets to capture IFA probe packets, while any port-initiator monitors also operate simultaneously.

Supported contexts are system and interface

The **no** form of this command removes the IFA initiator monitor to the **config-if** or **config** context

Parameter	Description
<NAME>	Specifies the IFA Monitor Initiator name.

Examples

Applying the IFA Initiator Monitor to the system.

```
switch(config)# ip flow ifa-initiator-monitor initiator-1
```

```
switch(config)# ipv6 flow ifa-initiator-monitor initiator-1
```

Applying the IFA Initiator Monitor to a port range:

```
switch(config-if-<1/1/1-1/1/10>)# ip flow ifa-initiator-monitor initiator-1
```

```
switch(config-if-<1/1/1-1/1/10>)# ipv6 flow ifa-initiator-monitor initiator-1
```

Applying the IFA Initiator Monitor to a specific port:

```
switch(config-if-<1/1/1>)# ip flow ifa-initiator-monitor initiator-2
```

```
switch(config-if-<1/1/1>)# ipv6 flow ifa-initiator-monitor initiator-2
```

Removing the IFA Initiator Monitor to the system:

```
switch(config)# no ip flow ifa-initiator-monitor initiator-1
```

```
switch(config)# no ipv6 flow ifa-initiator-monitor initiator-2
```

```
switch(config)# no ip flow ifa-initiator-monitor
```

```
switch(config)# no ipv6 flow ifa-initiator-monitor
```

Removing the IFA Monitor Initiator to a port range:

```
switch (config-if-<1/1/1-1/1/10>)# no ip flow ifa-initiator-monitor initiator-1
```

```
switch (config-if-<1/1/1-1/1/10>)# no ip flow ifa-initiator-monitor
```

```
switch (config-if-<1/1/1-1/1/10>)# no ipv6 flow ifa-initiator-monitor initiator-1
```

```
switch (config-if-<1/1/1-1/1/10>)# no ipv6 flow ifa-initiator-monitor
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	`config` `config-if`	Administrators or local user group members with execution rights for this command.

show flow ifa metrics

show flow ifa metrics [brief]

Description

Displays full information about current IFA flows. Timestamps are displayed in UTC. If **brief** is specified, then the flow IFA metrics summary is displayed.

Examples

Showing flow IFA metrics summary:

```
switch# show flow ifa metrics brief

IFA metrics data will be available only if there is a local active IFA terminator.
It can be validated with "show flow ifa monitor-status all"

|IP destination  |IP source      |Dst port   |Src port |IP protocol  |VRF   |
-------------------------------------------------------------------------------
-------------------------------------------------
 10.10.100.2      10.10.10.5      2000       1000       tcp           <not
available>
 2001:db8:3333    DDDD:EEEE:FFFF  8192       4096       udp           <not
available>
 200:3::03         200:3::15      2000       1000       tcp           <not
available>
 10.10.100.2      10.10.10.12     8192       4096       udp           <not
available>
```

The VRF is shown as **<not available>** because IFA is not performing VRF detection.

Showing flow IFA metrics on initiator or transit only node:

```
switch# show flow ifa metrics brief
```

```
IFA metrics data will be available only if there is a local active IFA terminator.
It can be validated with "show flow ifa monitor-status all"
```

Showing flow IFA metrics on 9300 Switch Series::

```
switch# show flow ifa metrics

Total IFA packets Received: The number of IFA packets received for this flow that
                            haven not been interrupted within a 5-minute time
range.
Network IFA Path Latency: The cumulative latency across all IFA hops.


Egress Queue Depth Cells: Indicates the queue memory usage at the moment the IFA
packet
                          passed through the switch. To determine the queue depth
in bytes,
                          multiply this value by the block size 254 bytes.
                          Monitoring this alongside latency and link speed in the
IFA reports
                          can aid in predicting latency based on queue depth.
Congestion Packet Counter: A monotonically increasing counter indicating the
number of packets,
                           out of the total received, that were marked as ECN-
congested.

IFA metrics data will be available only if there is a local active IFA terminator.
It can be validated with "show flow ifa monitor-status all"

Local IFA terminator metadata is not available. IFA metrics won't show local IFA
hop.

Total IFA Packets Received: 20
Destination IP             : 10.10.100.2
Source IP                  : 10.10.10.13
Destination Port           : 8192
Source Port                : 4096
IP Protocol                : udp
VRF                        : <not available>
Network Path Latency       : 1629 nanoseconds
ECN Congestion Hop         : <None>
Max Hop Latency
   Device ID: 878080
   Latency  : 875 nanoseconds
Terminator Timestamp       : 2024-07-02 19:21:30 UTC
Number of Hops             : 2
IFA Hop Path (Device ID)   : 908765 -> 878080

Device ID  Ingress ASIC Port Egress ASIC Port Egress Port Speed(Gbps) Egress Queue
Egress Queue TX Bytes Egress Queue Depth Cell Residence Time(ns) RX Timestamp(UTC)
  Congestion Packet Counter
---------- ----------------- ---------------- ----------------------- ------------
--------------------- ---------------------- ------------------ -----------------
-- -------------------------
908765     255               26               40                      1
7133048               1                      754                2024-07-02
19:21:28 0
878080     62                34               40                      1
8581664               1                      875                2024-07-02
19:21:30 0
```

```
================================================================================
================================================================================
==============================

Total IFA Packets Received: 916
Destination IP            : 2001:db8:3333:4444:5555:6666:7777:8888
Source IP                 : 200:3::03
Destination Port          : 2000
Source Port               : 1000
IP Protocol               : tcp
VRF                       : <not available>
Network Path Latency      : 358433092 nanoseconds
ECN Congestion Hop        : 908765
Max Hop Latency
   Device ID: 908765
   Latency  : 358432202 nanoseconds
Terminator Timestamp      : 2024-07-02 19:44:41 UTC
Number of Hops            : 3
IFA Hop Path (Device ID)  : 908765 -> 978670 -> 878080

Device ID  Ingress ASIC Port Egress ASIC Port Egress Port Speed(Gbps) Egress Queue
Egress Queue TX Bytes Egress Queue Depth Cell Residence Time(ns) RX Timestamp(UTC)
  Congestion Packet Counter
---------- ----------------- ---------------- ----------------------- ------------
-------------------- ---------------------- ------------------ -----------------
-- -------------------------
908765     255               26               40                      1
84878144             200734                   358432202          2024-07-02
19:44:38 915
978670     62                34               40                      1
94820192             1                        890                2024-07-02
19:44:40 915
878080     68                22               40                      1
80317792             1                        888                2024-07-02
19:44:41 915
```

Showing flow IFA metrics on 9300S Switch Series:

```
switch# show flow ifa metrics

Total IFA packets Received: The number of IFA packets received for this flow that
                            haven not been interrupted within a 5-minute time
range.
Network IFA Path Latency: The cumulative latency across all IFA hops


Egress Queue Depth Cells: Indicates the queue memory usage at the moment the IFA
packet
                          passed through the switch. To determine the queue depth
in bytes,
                          multiply this value by the block size 318 bytes.
                          Monitoring this alongside latency and link speed in the
IFA reports
                          can aid in predicting latency based on queue depth.
Congestion Packet Counter: A monotonically increasing counter indicating the
number of packets,
                           out of the total received, that were marked as ECN-
congested.

IFA metrics data will be available only if there is a local active IFA terminator.
```

```
It can be validated with "show flow ifa monitor-status all"
Total IFA Packets Received: 20
Destination IP            : 10.10.100.2
Source IP                 : 10.10.10.13
Destination Port          : 8192
Source Port               : 4096
IP Protocol               : udp
VRF                       : <not available>
Network Path Latency      : 1629 nanoseconds
ECN Congestion Hop        : <None>
Max Hop Latency
   Device ID: 878080
   Latency  : 875 nanoseconds
Terminator Timestamp      : 2024-07-02 19:21:30 UTC
Number of Hops            : 2
IFA Hop Path (Device ID)  : 908765 -> 878080

Device ID  Ingress ASIC Port Egress ASIC Port Egress Port Speed(Gbps) Egress Queue
Egress Queue TX Bytes Egress Queue Depth Cell Residence Time(ns) RX Timestamp(UTC)
  Congestion Packet Counter
---------- ----------------- ---------------- ----------------------- ------------
-------------------- ---------------------- ------------------ -----------------
-- -------------------------
908765     255               26               40                      1
7133048              1                       754                2024-07-02
19:21:28 0
878080     62                34               40                      1
8581664              1                       875                2024-07-02
19:21:30 0


================================================================================
================================================================================
==============================

Total IFA Packets Received: 916
Destination IP            : 2001:db8:3333:4444:5555:6666:7777:8888
Source IP                 : 200:3::03
Destination Port          : 2000
Source Port               : 1000
IP Protocol               : tcp
VRF                       : <not available>
Network Path Latency      : 358433092 nanoseconds
ECN Congestion Hop        : 908765
Max Hop Latency
   Device ID: 908765
   Latency  : 358432202 nanoseconds
Terminator Timestamp      : 2024-07-02 19:44:41 UTC
Number of Hops            : 3
IFA Hop Path (Device ID)  : 908765 -> 978670 -> 878080

Device ID  Ingress ASIC Port Egress ASIC Port Egress Port Speed(Gbps) Egress Queue
Egress Queue TX Bytes Egress Queue Depth Cell Residence Time(ns) RX Timestamp(UTC)
  Congestion Packet Counter
---------- ----------------- ---------------- ----------------------- ------------
-------------------- ---------------------- ------------------ -----------------
-- -------------------------
908765     255               26               40                      1
84878144             200734                  358432202          2024-07-02
19:44:38 915
978670     62                34               40                      1
94820192             1                       890                2024-07-02
```

```
19:44:40 915
878080     68                  22                40                     1
80317792             1                           888                    2024-07-02
19:44:41 915
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	Manager (#)	Administrators or local user group members with execution rights for this command.

show flow ifa monitor-status all

```
show flow ifa monitor-status all
```

Description

Displays information about currently configured IFA monitors.

Examples

Showing flow IFA monitor-status on an initiator node with IPv4 and IPv6 configured and no ports configured:

The VRF is shown as **<not available>** because IFA is not performing VRF detection.

Showing flow IFA metrics on initiator or transit only node:

```
switch# show flow ifa monitor-status all

*** System ***

initiator_ipv4
     Target State   : Active

initiator_ipv6
     Target State   : Active
```

Showing flow IFA monitor-status on an initiator node with IPv4 and IPv6 configured and some ports are configured:

```
switch# show flow ifa monitor-status all

*** System ***
```

```
initiator_ipv4
     Target State   : Active

initiator_ipv6
     Target State   : Active

Port 1/1/1
initiator_ipv4
     Target State   : Active

initiator_ipv6
     Target State   : Programming

Port 1/1/4
initiator_ipv6
     Target State   : Un-programming

Port 1/1/7
initiator_ipv4
     Target State   : Active
```

Command History

Release	Modification
10.15	Command introduced.

Command Information

Platforms	Command context	Authority
9300	Manager (#)	Administrators or local user group members with execution rights for this command.

show running-config

```
show running-config [all | flow-telemetry-profile | current-context | flow ifa |
interface <INTERFACE-NAME ]
```

Description

Shows the current running configuration.

Examples

Showing all the running configuration:

```
switch# show running-config all
Current configuration:
!
!
flow-telemetry-profile
    ifa-device-id auto
    ifa-hop-limit 10
    ifa-max-metadata-stack-length 80
    ifa-sampling-rate 4
```

Showing the flow-telemetry-profile configuration:

```
switch# show running-config flow-telemetry-profile

flow-telemetry-profile
    ifa-device-id 30
    ifa-hop-limit 55
    ifa-max-metadata-stack-length 90
    ifa-sampling-rate 13
```

Show running configuration after configuring some IFA initiator monitors to the system:

```
switch# show running-config flow ifa

Current configuration:
!
!
class ifa-ip class-ifa
class ifa-ipv6 class-ifa-ipv6
flow ifa-initiator-monitor ifa-init
    flow-filter-class class-ifa
flow ifa-initiator-monitor ifa-init2
    flow-filter-class class-ifa-ipv6
flow ifa-initiator-monitor ifa-init3
```

Show running configuration after configuring IFA initiator monitor to a specific port:

```
switch# show running-config flow ifa

Current configuration:
!
!
flow ifa-initiator-monitor ifa-init-port
interface 1/1/1
    ip flow ifa-initiator-monitor ifa-init-port
```

Show running configuration after configuring IFA initiator monitor to a port range:

```
switch# show running-config flow ifa

Current configuration:
!
!
flow ifa-initiator-monitor ifa-init-port
interface 1/1/1
    ip flow ifa-initiator-monitor ifa-init-port
interface 1/1/2
    ip flow ifa-initiator-monitor ifa-init-port
interface 1/1/3
    ip flow ifa-initiator-monitor ifa-init-port
interface 1/1/4
    ip flow ifa-initiator-monitor ifa-init-port
interface 1/1/5
    ip flow ifa-initiator-monitor ifa-init-port
```

Show running with IFA Monitor configuration combined:

```
switch# show running-config

Current configuration:
!
!
class ifa-ip class-ifa
class ifa-ipv6 class-ifa-ipv6
flow ifa-initiator-monitor ifa-init-ipv6
    flow-filter-class class-ifa-ipv6
flow ifa-initiator-monitor ifa-init
    flow-filter-class class-ifa
ip-all ifa-transit-monitor
ip flow ifa-initiator-monitor ifa-init
ipv6 flow ifa-initiator-monitor ifa-init-ipv6
interface 1/1/1
    ip flow ifa-initiator-monitor ifa-init
interface 1/1/2
    ipv6 flow ifa-initiator-monitor ifa-init-ipv6
interface 1/1/3
    ip flow ifa-initiator-monitor ifa-init
    ipv6 flow ifa-initiator-monitor ifa-init-ipv6
```

Command History

Release	Modification
10.15	Command introduced

Command Information

Platforms	Command context	Authority
9300	Manager (#)	Administrators or local user group members with execution rights for this command.

Configuring IFA using REST APIs

Before using the REST API, it is necessary to log in to the device. For more information on login and logout, configuring REST and REST sessions, refer to the REST API Guide.

Configure the IFA flow telemetry profile

To configure **Flow_Telemetry_Profile** IFA related setting, it is required to modify **Flow_Telemetry_Profile** attributes for existing with name **default**. This is the only row allowed in Flow_Telemetry_Profile. Avoid using **DELETE** and **POST** operations, as it may result on adding invalid **Flow_Telemetry_Profile settings** for the system.

- Attribute **ifa-device-id** accepts: **auto** or a 20 bits integer represented as a string.
- Attribute **ifa_hop_limit** accepts: integer from 1 to 255.
- Attribute **ifa_max_metadata_stack_length accepts: integer** from 8 to 248.
- Attribute **ifa_sampling_rate** accepts: integer from 4 to 4200000. Configured value will be multiplied by 1000, before configuring on hardware.

REST API information

Get the Flow_Telemetry_Profile writable attributes:

- **URI**: /rest/{version}/system/flow_telemetry_profiles/default
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**: selector=writable
- **Request body**: empty
- **Response body**:

```
{
  "ifa_device_id": "auto",
  "ifa_device_id": "4040",
  "ifa_hop_limit": 10,
  "ifa_max_metadata_stack_length": 80,
  "ifa_sampling_rate": 4
}
```

Get all Flow_Telemetry_Profile writable attributes to form a PUT request.

- **Description**: Get the Flow_Telemetry_Profile writable attributes to use them for a PUT operation.
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/get_system_flow_telemetry_profiles
- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable
- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable" \
-H  "accept: application/json"
```

CLI equivalent command: show running-config

Then, use the returned JSON, to form a PUT request:

- URI: /rest/{version}/system/flow_telemetry_profiles
- Valid versions: v10.15 and higher
- Operation: PUT
- Query parameters: empty
- Request body:

```
      {
  "ifa_device_id": "4040",
  "ifa_hop_limit": 10,
  "ifa_max_metadata_stack_length": 80,
  "ifa_sampling_rate": 4
}
```

- **Description**: Configure Flow_Telemetry_Profile ifa_hop_limit together with all writable attributes current value returned on the GET operation, with a PUT operation
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/ut_system_flow_telemetry_profiles__Flow_Telemetry_Profile_name_

- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default
- **Curl example**:

```
curl {IP} -X PUT -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_hop_limit\": 5,
\"ifa_device_id\":\"4040\",
\"ifa_max_metadata_stack_length\": 80,
\"ifa_sampling_rate\": 4
}"
```

CLI equivalent commands:

- ifa-device-id {auto|<*VALUE*>}
- ifa-hop-limit <*VALUE*>
- ifa-max-metadata-stack-length <*VALUE*>
- ifa-sampling-rate <*VALUE*>

Configure IFA device ID

To configure IFA **ifa-device-id** it is required to write the attribute **Flow_Telemetry_Profile::ifa_device_id** for the row with name default, that can be modified in the **Flow_Telemetry_Profile** resource. Attribute **ifa-device-id** accepts: **auto** or a 20 bits integer represented as a string.

REST API information

Get the **Flow_Telemetry_Profile ifa_max_metadata_stack_length** attribute:

- **URI**: /rest/{version}/system/flow_telemetry_profiles/default
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**: selector=writable
- **Request body**: empty
- **Response body**:

```
{
  "ifa_device_id": "auto",
  ...
}
```

Get all Flow_Telemetry_Profile **writable** attributes to form a **PUT** request.

- **Description**: Get the Flow_Telemetry_Profile writable attributes to use them for a PUT operation.
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/get_system_flow_telemetry_profiles
- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable
- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable" \
-H  "accept: application/json"
```

CLI equivalent command: show running-config

Then, use the returned JSON, to form a PUT request:

- URI: /rest/{version}/system/flow_telemetry_profiles
- Valid versions: v10.15 and higher
- Operation: PUT
- Query parameters: empty
- Request body:

```
      {
  "ifa_device_id": "4040",
  "ifa_hop_limit": 10,
  "ifa_max_metadata_stack_length": 80,
  "ifa_sampling_rate": 4
}
```

- **Description**: Configure Flow_Telemetry_Profile **ifa_device_id** together with all writable attributes current value returned on the GET operation, with a PUT operation
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/ut_system_flow_telemetry_profiles__Flow_Telemetry_Profile_name_
- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default
- **Curl example**:

```
curl {IP} -X PUT -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_device_id\":\"5050\",
\"ifa_hop_limit\": 10,
\"ifa_max_metadata_stack_length\": 80,
\"ifa_sampling_rate\": 4
}"
```

CLI equivalent command: ifa-device-id {auto|<*VALUE*>}

Alternatively, **ifa_device_id** can be modified with a PATCH request (instead of a GET and a PUT) as follows:

- **URI**: /rest/{version}/system/flow_telemetry_profiles
- **Valid versions**: v10.15 and higher
- **Operation**: PATCH
- **Query parameters**: empty
- **Request body**:

```
      {
  "ifa_device_id": "5050"
}
```

- **Description**: Update **ifa_device_id** with a PATCH operation
- **Swagger**:

```
https://{IP}/api/{version}/#/Flow_Telemetry_Profile/patch_system_flow_
telemetry_profiles__Flow_Telemetry_Profile_name_
```

- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default
- **Curl example**:

```
curl -X PATCH -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_device_id\":\"2020\"
}"
```

CLI equivalent command: ifa-device-id {auto|<*VALUE*>}

Configure IFA hop-limit

To configure IFA **ifa-hop-limit** it is required to write the attribute **Flow_Telemetry_Profile::ifa_hop_limit** for the row with name **default**, that can be modified in the **Flow_Telemetry_Profile** resource. Attribute ifa_hop_limit accepts: accepts: integer from 1 to 255.

REST API information

Get the Flow_Telemetry_Profile **ifa-hop-limit** attribute:

- **URI**: /rest/{version}/system/flow_telemetry_profiles/default
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**: selector=writable
- **Request body**: empty
- **Response body**:

```
{
  "ifa_hop_limit": 10,
  ...
}
```

- Get all Flow_Telemetry_Profile **writable** attributes to form a **PUT** request.
- **Description**: Get the Flow_Telemetry_Profile writable attributes to use them for a PUT operation.
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/get_system_flow_telemetry_profiles

- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable
- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable" \
-H  "accept: application/json"
```

CLI equivalent command: **show running-config**

Then, use the returned JSON, to form a PUT request:

- URI: /rest/{version}/system/flow_telemetry_profiles
- Valid versions: v10.15 and higher
- Operation: PUT
- Query parameters: empty
- Request body:

```
{
  "ifa_device_id": "4040",
  "ifa_hop_limit": 10,
  "ifa_max_metadata_stack_length": 80,
  "ifa_sampling_rate": 4
}
```

- **Description**: Configure Flow_Telemetry_Profile **ifa-hop-limit** together with all writable attributes current value returned on the GET operation, with a PUT operation
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/ut_system_flow_telemetry_profiles__Flow_Telemetry_Profile_name_
- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default
- **Curl example**:

```
curl {IP} -X PUT -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_device_id\":\"5050\",
\"ifa_hop_limit\": 10,
\"ifa_max_metadata_stack_length\": 80,
\"ifa_sampling_rate\": 4
}"
```

CLI equivalent command: ifa-hop-limit <*VALUE*>

Alternatively, **ifa_hop_limit** can be modified with a PATCH request (instead of a GET and a PUT) as follows:

- **URI**: /rest/{version}/system/flow_telemetry_profiles
- **Valid versions**: v10.15 and higher
- **Operation**: PATCH

- **Query parameters**: empty
- **Request body**:

```
{
  "ifa_hop_limit": 5
}
```

- **Description**: Update **ifa_device_id** with a PATCH operation
- **Swagger**:

```
https://{IP}/api/{version}/#/Flow_Telemetry_Profile/patch_system_flow_
telemetry_profiles__Flow_Telemetry_Profile_name_
```

Full URI: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default

Curl example:

```
curl -X PATCH -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_hop_limit\":8
}"
```

CLI equivalent commands: ifa-hop-limit <*VALUE*>

Configure IFA max-metadata-stack-length

To configure IFA **ifa_max_metadata_stack_length** it is required to write the attribute **Flow_Telemetry_Profile::ifa_max_metadata_stack_length** for row with name **default**, that can be modified in the **Flow_Telemetry_Profile** resource. Attribute **ifa_max_metadata_stack_length** accepts: integer from 8 to 248.

REST API information

Get the Flow_Telemetry_Profile **ifa_max_metadata_stack_length** attribute:

- **URI**: /rest/{version}/system/flow_telemetry_profiles/default
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**: selector=writable
- **Request body**: empty
- **Response body**:

```
{
 "ifa_max_metadata_stack_length": 80,
 ...
}
```

Get all Flow_Telemetry_Profile **writable** attributes to form a **PUT** request.

- **Description**: Get the Flow_Telemetry_Profile writable attributes to use them for a PUT operation.
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/get_system_flow_telemetry_profiles
- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable
- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable" \
-H  "accept: application/json"
```

CLI equivalent command: **show running-config**

Then, use the returned JSON, to form a PUT request:

- URI: /rest/{version}/system/flow_telemetry_profiles
- Valid versions: v10.15 and higher
- Operation: PUT
- Query parameters: empty
- Request body:

```
      {
  "ifa_device_id": "4040",
  "ifa_hop_limit": 10,
  "ifa_max_metadata_stack_length": 80,
  "ifa_sampling_rate": 4
}
```

- **Description**: Configure Flow_Telemetry_Profile **ifa_max_metadata_stack_length** together with all writable attributes current value returned on the GET operation, with a PUT operation
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/ut_system_flow_telemetry_profiles__Flow_Telemetry_Profile_name_
- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default
- **Curl example**:

```
curl {IP} -X PUT -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_hop_limit\": 5,
\"ifa_device_id\":\"4040\",
\"ifa_max_metadata_stack_length\": 80,
\"ifa_sampling_rate\": 4
}"
```

CLI equivalent command: ifa-max-metadata-stack-length *<VALUE>*

Alternatively,Flow_Telemetry_Profile ifa_max_metadata_stack_length can be modified with a PATCH request (instead of a GET and a PUT) as follows:

- **URI**: /rest/{version}/system/flow_telemetry_profiles
- **Valid versions**: v10.15 and higher
- **Operation**: PATCH
- **Query parameters**: empty
- **Request body**:

```
       {
  "ifa_max_metadata_stack_length": 40
}
```

- **Description**: Update **ifa_max_metadata_stack_length** with a PATCH operation
- **Swagger**:

```
https://{IP}/api/{version}/#/Flow_Telemetry_Profile/patch_system_flow_
telemetry_profiles__Flow_Telemetry_Profile_name_
```

- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default
- **Curl example**:

```
curl -X PATCH -b /tmp/cookies \

"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_max_metadata_stack_length\":48
}"
```

CLI equivalent commands: ifa-max-metadata-stack-length <*VALUE*>

Configure IFA sampling-rate

To configure IFA **ifa_sampling_rate** it is required to write the attribute **Flow_Telemetry_Profile::ifa_sampling_rate** for row with name **default**, that can be modified in the **Flow_Telemetry_Profile** resource. Attribute **ifa_sampling_rate** accepts: integer from 4 to 4200000. Configured value will be multiplied by 1000, before configuring on hardware.

REST API information

Get the Flow_Telemetry_Profile **ifa_sampling_rate** attribute:

- **URI**: /rest/{version}/system/flow_telemetry_profiles/default
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**: selector=writable
- **Request body**: empty
- **Response body**:

```
{
  "ifa_sampling_rate": 4,
 ...
}
```

Get all Flow_Telemetry_Profile **writable** attributes to form a **PUT** request.

- **Description**: Get the Flow_Telemetry_Profile writable attributes to use them for a PUT operation.
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/get_system_flow_telemetry_profiles
- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable
- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles?selector=writable" \
-H  "accept: application/json"
```

CLI equivalent command: **show running-config**

Then, use the returned JSON, to form a **PUT** request:

- URI: /rest/{version}/system/flow_telemetry_profiles
- Valid versions: v10.15 and higher
- Operation: PUT
- Query parameters: empty
- Request body:

```
      {
 "ifa_sampling_rate": 4,
  "ifa_device_id": "4040",
  "ifa_hop_limit": 10,
  "ifa_max_metadata_stack_length": 80
}
}
```

- **Description**: Configure Flow_Telemetry_Profile **ifa_sampling_rate** together with all writable attributes current value returned on the GET operation, with a PUT operation
- **Swagger**: https://{IP}/api/{version}/#/Flow_Telemetry_Profile/ut_system_flow_telemetry_profiles__Flow_Telemetry_Profile_name_
- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default
- **Curl example**:

```
curl {IP} -X PUT -b /tmp/cookies \

"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_sampling_rate\": 600,
\"ifa_device_id\":\"5050\",
```

```
\"ifa_hop_limit\": 10,
\"ifa_max_metadata_stack_length\": 80
}"
```

CLI equivalent command: ifa-sampling-rate <*VALUE*>

Alternatively,Flow_Telemetry_Profile **ifa_sampling_rate** can be modified with a PATCH request (instead of a GET and a PUT) as follows:

- **URI**: /rest/{version}/system/flow_telemetry_profiles
- **Valid versions**: v10.15 and higher
- **Operation**: PATCH
- **Query parameters**: empty
- **Request body**:

```
      {
 "ifa_sampling_rate": 600
}
```

- **Description**: Update **ifa_sampling_rate** with a PATCH operation
- **Swagger**:

```
https://{IP}/api/{version}/#/Flow_Telemetry_Profile/patch_system_flow_
telemetry_profiles__Flow_Telemetry_Profile_name_
```

- **Full URI**: https://{IP}/rest/{version}/system/flow_telemetry_profiles/default
- **Curl example**:

```
curl -X PATCH -b /tmp/cookies \
"https://{IP}/rest/{version}/system/flow_telemetry_profiles/default" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_sampling_rate\":800
}"
```

CLI equivalent commands: ifa-sampling-rate <*VALUE*>

Configure Flow IFA initiator monitor

To configure a Flow IFA initiator monitor it is required to create a new **IFA_Initiator_Flow_Monitor** resource.

REST API information

- **URI**: /rest/{version}/system/ifa_initiator_flow_monitor
- **Valid versions**: v10.15 and higher
- **Operation**: POST

- **Query parameters**: empty
- **Request body**:

```
{
  "name": "<IFA_INITIATOR_MONITOR_NAME>"
}
```

- **Response body**: empty
- **Description**: Create a custom queue profile
- **Swagger**: https://{IP}/api/{version}/#/IFA_Initiator_Flow_Monitor/post_system_ifa_initiator_flow_monitors
- **Full URI**: https://{IP}/rest/{version}/system/ifa_initiator_flow_monitor
- **Curl example**:

```
curl -X POST -b /tmp/cookies \
"https://{IP}/rest/{version}/system/ifa_initiator_flow_monitor" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
'{
   "name": "{IFA_INITIATOR_MONITOR_NAME}"
}'
```

CLI equivalent command: flow ifa-initiator-monitor <*NAME*>

Configure flow filter for an IFA initiator monitor

To configure a Filter Class for a Flow IFA initiator monitor it is required to execute a PATCH request for **IFA_Initiator_Flow_Monitor**.

REST API information

Filter Class for a Flow IFA initiator monitor accepts class with types **ifa-ipv4** and **ifa-ipv6**.

- **URI**: /rest/{version}/system/ifa_initiator_flow_monitors
- **Valid versions**: v10.15 and higher
- **Operation**: PATCH
- **Query parameters**: selector=writable
- **Request body**:

```
{
  "class": "/rest/v10.15/system/classes/{CLASS_NAME},{CLASS_TYPE}"
}
```

Description: Update IFA_Initiator_Flow_Monitor **class** with a **PATCH** operation

Swagger:

```
https://{IP}/api/{version}/#/IFA_Initiator_Flow_Monitor/patch_system_ifa_
initiator_flow_monitors
```

```
__IFA_Initiator_Flow_Monitor_name_
```

Full URI:

```
https://{IP}/rest/{version}/system/ifa_initiator_flow_monitors/{IFA_INITIATOR_
MONITOR_NAME}
```

Curl example:

```
curl -X PATCH -b /tmp/cookies \
"https://{IP}/rest/{version}/system/ifa_initiator_flow_monitors/{IFA_INITIATOR_
MONITOR_NAME}" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"class\":\"/rest/v10.15/system/classes/{CLASS_NAME},{CLASS_TYPE}\"
}"
```

CLI equivalent command: flow-filter-class

Apply IFA initiator flow monitor

To apply an IFA initiator flow monitor, it is required to execute a **PATCH** request for **System** or **Port**, depending on the context that is going to be applied.

REST API information

Flow IFA initiator monitor, can be applied for ipv4 or ipv6 IP address versions.

Request for System

- **URI**: /rest/{version}/system
- **Valid versions**: v10.15 and higher
- **Operation**:PATCH
- **Query parameters**: empty
- **Request body**:

```
{
  "ifa_initiator_flow_monitor": {"{IP_VERSION}":"/rest/{version}/system/ifa_
initiator_flow_monitors
/{IFA_INITIATOR_MONITOR_NAME}"}
}
```

- **Description**: Update IFA_Initiator_Flow_Monitor class with a **PATCH** operation
- **Swagger**: https://{IP}/api/{version}/#/System/patch_system
- **Full URI**: https://{IP}/rest/{version}/system
- **Curl example**:

```
curl -X PATCH -b /tmp/cookies \
"https://{IP}/rest/{version}/system/" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_initiator_flow_monitor\":{\"{IP_VERSION}\":\"/rest/{version}/system/ifa_
initiator_flow_monitors/{IFA_INITIATOR_MONITOR_NAME}\"}
}"
```

Request for Port

Port name **1/1/1** should be used as **/1%2F1%2F1**. Port name **1/1/1:1** should be used as **/1%2F1%2F1A1**.

- **URI**: /rest/{version}/system/interfaces/{INTERFACE_NAME}
- **Valid versions**: v10.15 and higher
- **Operation**:PATCH
- **Query parameters**: empty
- **Request body**:

```
{
  "ifa_initiator_flow_monitor": {"{IP_VERSION}":"/rest/{version}/system/ifa_
initiator_flow_monitors/{IFA_INITIATOR_MONITOR_NAME}"}
}
```

- **Description**: Update IFA_Initiator_Flow_Monitor class with a **PATCH** operation
- **Swagger**: https://{IP}/api/{version}/#/Interface/patch_system_interfaces__Interface_name_
- **Full URI**:https://{IP}/rest/{version}/system/interfaces/{INTERFACE_NAME}
- **Curl example**:

```
curl -X PATCH -b /tmp/cookies \
"https://{IP}/rest/{version}/system/interfaces/{INTERFACE_NAME}" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_initiator_flow_monitor\":{\"{IP_VERSION}\":\"/rest/{version}/system/ifa_
initiator_flow_monitors/{IFA_INITIATOR_MONITOR_NAME}\"}
}"
```

Once several IFA flow monitors are configured for different IP versions, **PUT** operations are required to do partial modifications on **ifa_initiator_flow_monitor**, for example removing the configurations for a single IP version.

Apply IFA ip-all monitor behavior

To configure IFA ip-all monitor behavior, it is required to execute a PATCH request for **System**.

Accepted values are **transit**, **terminator** or **none**. The **none** value is used to unconfigure IFA

REST API information

Configuring system as an IFA transit.

Request for System

- **URI**: /rest/{version}/system
- **Valid versions**: v10.15 and higher
- **Operation**:PATCH
- **Query parameters**: empty
- **Request body**:

```
{
  "ifa_flow_monitor": "transit"
}
```

- **Description**: Update IFA_Initiator_Flow_Monitor class with a **PATCH** operation
- **Swagger**: https://{IP}/api/{version}/#/System/patch_system
- **Full URI**: https://{IP}/rest/{version}/system
- **Curl example**:

```
curl -X PATCH -b /tmp/cookies \
"https://{IP}/rest/{version}/system/" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{
\"ifa_flow_monitor\":\"transit\"
}"
```

Configuring system as an IFA terminator.

Request for System

- **URI**: /rest/{version}/system
- **Valid versions**: v10.15 and higher
- **Operation**:PATCH
- **Query parameters**: empty
- **Request body**:

```
{
  "ifa_flow_monitor": "terminator"
}
```

- **Description**: Update IFA_Initiator_Flow_Monitor class with a **PATCH** operation
- **Swagger**: https://{IP}/api/{version}/#/System/patch_system
- **Full URI**: https://{IP}/rest/{version}/system
- **Curl example**:

```
curl -X PATCH -b /tmp/cookies \
"https://{IP}/rest/{version}/system/" \
-H  "accept: */*" \
-H  "Content-Type: application/json" \
-d \
"{\"ifa_flow_monitor\":\"terminator\"
```

```
}"
```

Show flow IFA metrics

To retrieve IFA metrics on the terminator node, it is required to read **Inband_Flow_Analyzer_Metrics**.

IFA metrics data will be available only if there is a local active IFA terminator. It can be validated with Show flow IFA monitor status via REST .

REST API information

Get the **Inband_Flow_Analyzer_Metrics key attributes**, equivalent to **Show flow IFA metrics brief** command:

- **URI**: /rest/{version}/system/inband_flow_analyzer_metrics?attributes=&depth=1
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**: attributes=destination_ip,ip_protocol,source_port,source_ip,destination_port,vrf&depth=1
- **Request body**: empty
- **Response body**:

```
{
 "{PROTOCOL},{SRC-PORT},{DST-PORT},{SRC-IP},{DST-IP},{VRF}":
"/rest/{version}/system/inband_flow_analyzer_metrics/{PROTOCOL},{SRC-PORT},{DST-
PORT},{SRC-IP},{DST-IP},{VRF-NAME}"

  }
```

- **Description**: Get the IFA flow metrics attributes.
- **Swagger**: https://{IP}/api/{version}/#/Inband_Flow_Analyzer_Metrics/get_system_inband_flow_analyzer_metrics
- **Full URI**: https://{IP}/rest/{version}/system/inband_flow_analyzer_metrics?attributes=destination_ip,ip_protocol,source_port,source_ip,destination_port,vrf&depth=1
- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system/inband_flow_analyzer_
metrics?attributes=destination_ip,
ip_protocol,source_port,source_ip,destination_port,vrf&depth=1" \
-H  "accept: application/json"
```

Number of entries on the response depends on the IFA Flows being tracked by IFA Terminator node

CLI equivalent command:show flow ifa metrics brief

Get all Inband_Flow_Analyzer_Metrics attributes, equivalent to the full show command:

- **URI**: /rest/{version}/system/inband_flow_analyzer_metrics?depth=2
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**: depth=2

- **Request body**: empty
- **Response body**:

```
{
  "{PROTOCOL},{SRC-PORT},{DST-PORT},{SRC-IP},{DST-IP},{VRF-NAME}": {
    "destination_ip": "{DST-IP}",
    "destination_port": {DST-PORT},
    "ecn_congestion_hop": {CONGESTED-DEVICE-ID},
    "hop_metric": {
      "1": {
        "congestion": {CONGESED-PACKETS-COUNT},
        "device_id": {DEVICE-ID},
        "egress_port": {PORT-ID},
        "egress_port_speed": "{PORT-SPEED}",
        "egress_queue": {QUEUE-ID},
        "egress_queue_depth": {QUEUE-DEPTH},
        "egress_queue_tx_bytes": {QUEUE-TX-BYTES},
        "ingress_port": {PORT-ID},
        "residence_time": {NS-RESIDENCE-TIME},
        "rx_timestamp": {RX-TIMESTAMP}
      },
      "2": {
        "congestion": {CONGESED-PACKETS-COUNT},
        "device_id": {DEVICE-ID},
        "egress_port": {PORT-ID},
        "egress_port_speed": "{PORT-SPEED}",
        "egress_queue": {QUEUE-ID},
        "egress_queue_depth": {QUEUE-DEPTH},
        "egress_queue_tx_bytes": {QUEUE-TX-BYTES},
        "ingress_port": {PORT-ID},
        "residence_time": {NS-RESIDENCE-TIME},
        "rx_timestamp": {RX-TIMESTAMP}
      },
      "3": {
        "congestion": {CONGESED-PACKETS-COUNT},
        "device_id": {DEVICE-ID},
        "egress_port": {PORT-ID},
        "egress_port_speed": "{PORT-SPEED}",
        "egress_queue": {QUEUE-ID},
        "egress_queue_depth": {QUEUE-DEPTH},
        "egress_queue_tx_bytes": {QUEUE-TX-BYTES},
        "ingress_port": {PORT-ID},
        "residence_time": {NS-RESIDENCE-TIME},
        "rx_timestamp": {RX-TIMESTAMP}
      }
    },
    "ip_protocol": "{PROTOCOL}",
    "network_path_latency": {NS-PATH-LATENCY},
    "single_hop_max_latency": {
      "device_id": {DEVICE-ID},
      "max_latency": {NS-RESIDENCE-TIME}
    },
    "source_ip": "{SRC-IP}",
    "source_port": {SRC-PORT},
    "terminator_timestamp": {RX-TIMESTAMP},
    "total_ifa_packets": {TOTAL_PACKETS},
    "vrf": "{VRF_NAME}"
  }
}
```

Example:

```
{
  "udp,15,14,2.2.2.2,1.1.1.1,": {
    "destination_ip": "1.1.1.1",
    "destination_port": 14,
    "ecn_congestion_hop": null,
    "hop_metric": {
      "1": {
        "congestion": 0,
        "device_id": 768955,
        "egress_port": 32,
        "egress_port_speed": "100_GBPS",
        "egress_queue": 1,
        "egress_queue_depth": 260931,
        "egress_queue_tx_bytes": 65537,
        "ingress_port": 2,
        "residence_time": 69632,
        "rx_timestamp": 1715274267491805200
      },
      "2": {
        "congestion": 0,
        "device_id": 978670,
        "egress_port": 80,
        "egress_port_speed": "400_GBPS",
        "egress_queue": 2,
        "egress_queue_depth": 129939,
        "egress_queue_tx_bytes": 262401,
        "ingress_port": 5,
        "residence_time": 81920,
        "rx_timestamp": 1715274267613912300
      },
      "3": {
        "congestion": 0,
        "device_id": 1048575,
        "egress_port": 96,
        "egress_port_speed": "10_GBPS",
        "egress_queue": 3,
        "egress_queue_depth": 129955,
        "egress_queue_tx_bytes": 327937,
        "ingress_port": 6,
        "residence_time": 86016,
        "rx_timestamp": 1715274263157783600
      }
    },
    "ip_protocol": "udp",
    "network_path_latency": 237568,
    "single_hop_max_latency": {
      "device_id": 1048575,
      "max_latency": 86016
    },
    "source_ip": "2.2.2.2",
    "source_port": 15,
    "terminator_timestamp": 1715274263157783600,
    "total_ifa_packets": 30,
    "vrf": "<not available>"
  }
}
```


- **Description**: Get the IFA flow metrics attributes.
- **Swagger**: https://{IP}/api/{version}/#/Inband_Flow_Analyzer_Metrics/get_system_inband_flow_analyzer_metrics

- **Full URI**: https://{IP}/rest/{version}/system/inband_flow_analyzer_metrics?depth=2
- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system/inband_flow_analyzer_metrics?depth=2"
```

Number of entries on the response depends on the IFA Flows being tracked by IFA Terminator node. Number of Hops found on each entry depend on the IFA Flows being tracked by IFA Terminator node.

CLI equivalent command:show flow ifa metrics

It is also possible to subscribe for notifications using the following subscription body:

```
{
    "type": "subscribe",
    "topics": [
        {
            "name": "/rest/latest/system/inband_flow_analyzer_metrics?depth=2"
        }
    ]
}
```

Show flow IFA monitor status

To retrieve IFA monitor status it is required to read the attributes **System::ifa_flow_monitor_current_state**, **Flow_Telemetry_Profile::ifa_flow_monitor_target_state** and **Port::ifa_flow_monitor_current_state, Port::ifa_flow_monitor_target_state**.

REST API information

Get the System attributes:

- **URI**: /rest/{version}/system?attributes=ifa_flow_monitor_current_state,ifa_flow_monitor_target_state
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**:attributes=ifa_flow_monitor_current_state,ifa_flow_monitor_target_state
- **Request body**: empty
- **Response body**:

```
{
  "ifa_flow_monitor_current_state": {
    "{IFA-MONITOR-IP-VERSION}": "{IFA-MONITOR-STATUS}"
  },
  "ifa_flow_monitor_target_state": {
    "{IFA-MONITOR-IP-VERSION}": "{IFA-MONITOR-STATUS}"
  }
}
```

- **Description**: Get the System attributes.
- **Swagger**: https://{IP}/api/{version}/#/System/get_system
- **Full URI**: https://{IP}/rest/{version}/system?attributes=ifa_flow_monitor_current_state,ifa_flow_

monitor_target_state

- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system?attributes=ifa_flow_monitor_current_state,ifa_
flow_monitor_target_state" \
-H  "accept: application/json"
```

IFA possible system monitors are initiator_ipv4, initiator_ipv6, transit_ip_all, and terminator_ip_all.

CLI equivalent command: Show flow IFA monitor-status

Get the Port attributes:

- **URI**: /rest/{version}/system/interfaces?attributes=ifa_flow_monitor_current_state,ifa_flow_monitor_target_state&depth=2
- **Valid versions**: v10.15 and higher
- **Operation**: GET
- **Query parameters**: attributes=ifa_flow_monitor_current_state,ifa_flow_monitor_target_state&depth=2
- **Request body**: empty
- **Response body**:

```
{
  "{PORT-NAME}": {
    "ifa_flow_monitor_current_state": {
      "{IFA-MONITOR-IP-VERSION}": "{IFA-MONITOR-STATUS}"
    },
    "ifa_flow_monitor_target_state": {
      "{IFA-MONITOR-IP-VERSION}": "{IFA-MONITOR-STATUS}"
    }
  },
  "{PORT-NAME}": {
    "ifa_flow_monitor_current_state": {},
    "ifa_flow_monitor_target_state": {}
  }
}
```

- **Description**: Get the Port attributes.
- **Swagger**: https://{IP}/api/{version}/#/Interface/get_system_interfaces
- **Full URI**: https://{IP}/rest/{version}/system/interfaces?attributes=ifa_flow_monitor_current_state,ifa_flow_monitor_target_state&depth=2
- **Curl example**:

```
curl -X GET -b /tmp/cookies \
"https://{IP}/rest/{version}/system/interfaces?attributes=ifa_flow_monitor_
current_state,ifa_flow_monitor_target_state&depth=2" \
-H  "accept: application/json"
```

IFA possible port monitors are initiator_ipv4 and initiator_ipv6.

CLI equivalent command: Show flow IFA monitor-status

Chapter 6
Boot commands

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.

.

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

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

Usage

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

The output of this command includes the following information:

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

Examples

Showing the boot history of the active management module:

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

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

Index : 1
```

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

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

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

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
...
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

Chapter 7
External storage

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
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320	config-external-storage-<*VOLUME-NAME*>	Operators or Administrators or local user group members with execution

Platforms	Command context	Authority
8325 8325H 8325P 8360 9300 10000		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
8100 8320 8325 8325H 8325P 8360 9300 10000	`config-external-storage-<VOLUME-NAME>`	Operators or Administrators or local user group members with execution rights for this command. Operators can execute this command from the operator context (>) only.

external-storage

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

Description

Creates or updates an external storage volume.

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

Examples

Creating the logfiles storage volume:

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

Deleting the logfiles storage volume:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8325H 8325P 8360 9300 10000	`config`	Administrators or local user group members with execution rights for this command.

password (external-storage)

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

Description

Sets the password for network attached storage server login.

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

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

Examples

Creating a volume named logfiles with password Xj#9:

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

Creating a volume named bak1 with a prompted plaintext password:

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

Clearing the password for volume logfiles:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8325H 8325P 8360 9300 10000	config-external-storage-*<VOLUME-NAME>*	Administrators or local user group members with execution rights for this command.

show external-storage

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

Description

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

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

Examples

```
switch# show external-storage

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8325H 8325P 8360 9300 10000	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

show running-config external-storage

`show running-config external-storage`

Description

Shows the running configuration of the external storage.

Examples

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

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

```
    directoty /home
    enable
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320 8325 8325H 8325P 8360 9300 10000	`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
8100 8320 8325 8325H 8325P 8360 9300 10000	`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
8100 8320 8325 8325H 8325P 8360	`config-external-storage-<VOLUME-NAME>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
9300 10000		

Chapter 8
IP-SLA

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
8100 8320 8325 8325H 8325P 8360 9300 10000	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.

Parameter	Description
<URL>	Specifies HTTPS URL address of syntax. https://<HOST NAME/IP-ADDRESS>:<PORT>/<PATH>.
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.
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
8100 8320 8325 8325H 8325P 8360 9300 10000	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.

Parameter	Description
`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 source interface to use for sending IP-SLA probes.
`timeout <TIMEOUT>`	Specifies the interval before a probe is timed out. Range: 5 to 604800.
`tos <TYPE-OF-SERVICE>`	Specifies the type of serve 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
8100 8320 8325	`config-ip-sla-<IP-SLA-NAME>`	Administrators or local user group members with execution rights for this command.

Platforms	Command context	Authority
8325H 8325P 8360 9300 10000		

ip-sla

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

Description

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

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

Parameter	Description
<IP-SLA-NAME>	Specifies an IP-SLA profile name. Length: 1 to 64 characters.

Examples

Creating an IP-SLA:

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

Deleting an IP-SLA:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8325H 8325P 8360 9300 10000	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
<SLA-NAME>	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 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
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320 8325 8325H 8325P 8360 9300 10000	`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
8100 8320 8325 8325H 8325P 8360 9300 10000	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

tcp-connect

```
tcp-connect {<DEST-IPV4-ADDR>|<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.

Parameter	Description
	Range: 60 to 7200.
`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
8100 8320 8325 8325H 8325P 8360 9300 10000	`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
8100 8320 8325 8325H 8325P 8360 9300 10000	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.
advantage-factor <*ADVANTAGE-FACTOR*>	Selects the value for the advantage factor. Default value is 0. Range: 0 to 20.
codec-type <*CODEC-TYPE*>	Selects the codec-type for the Voip IP-SLA test.
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: 240 to 7200.
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: 120 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 source port for the IP-SLA probes.
tos <*TYPE-OF-SERVICE*>	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
8100 8320 8325 8325H 8325P 8360 9300 10000	config-ip-sla-*<IP-SLA-NAME>*	Administrators or local user group members with execution rights for this command.

vrf

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

Description

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

The **no** form of the command removes VRF from SLA.

Parameter	Description
<VRF-NAME>	Specifies a VRF name. Length: Default: default. 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
8100 8320 8325 8325H 8325P 8360 9300 10000	`config-ip-sla-<IP-SLA-NAME>`	Administrators or local user group members with execution rights for this command.

show interface

```
show interface [<IFNNAME>|<IFRANGE>] [brief | physical]
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]
show interface vxlan <VXLAN-ID> [brief | physical]
show interface vxlan <VXLAN-ID> [brief | physical]
```

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.
`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.

Parameter	Description
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 (the **persona** item is only available on the HPE Aruba Networking 10000 Switch Series):

```
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
 Persona: access
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
MDI 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 shut down during a VSX split (the **persona** item is only available on the HPE Aruba Networking 10000 Switch Series):

```
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:
 Persona: access
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 the monitor information:

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

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

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

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

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

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

Showing information about extended counters:

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

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

Showing interface link-status:

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


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

Showing interface loopback 1 link-status:

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

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

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

Showing interface link-status:

```
switch# show interface link-status
------------------------------------------------------------------------
Port           Type            Physical    Link        Link Flaps  Last
                               Link State  Transitions Ignored     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 `monitor` parameter.
10.10	Added `human-readable` parameter.
10.09	Added persona information for the 10000 Switch Series.
10.07 or earlier	--

Command Information

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

show interface statistics

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

Description

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

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

Examples

Showing statistics of all interfaces:

```
show interface statistics
-------------------------------------------------------------------------------------------------------------------------------------------------------
Interface                RX Bytes  RX Packets  RX Drops  TX Bytes  TX Packets  TX Drops  RX Broadcast RX Multicast  TX Broadcast TX Multicast  RX Pause  TX 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.

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

                                                  help: ?, quit: q
```

```
Help for Interface Monitor

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

c    Clear interface statistics
        Does not apply to rates

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

Delay:2
                                                                    help: ?, qui
```

Showing monitor error statistics in human-readable format:

```
show interface 1/1/1-1/1/3,1/1/30-1/1/32 error-statistics monitor human-readable
-------------------------------------------------------------------------------------------------
Interface                RX Errors    TX Errors    RX Giants     RX Runts      CRC/FCS  Collisions
-------------------------------------------------------------------------------------------------
1/1/1                            0            0            0            0            0           0
1/1/2                            0            0            0            0            0           0
1/1/3                            0            0            0            0            0           0
1/1/30                           0            0            0            0            0           0
1/1/31                           0            0            0            0            0           0
1/1/32                           0            0            0            0            0           0


Human-readable                                                          help: ?, quit: q
```

```
Help for Interface Monitor

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

c    Clear interface statistics
        Does not apply to rates

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

Delay:2
                                                                help: ?, quit: q
```

Command History

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

Command Information

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

Chapter 9
Mirroring

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
- 8320, 8325/8325H/8325P, 9300,9300S, and 10000 Switch series have a limit of 2 simultaneously enabled Mirroring Sessions with LAG sources, regardless of direction.
- Ability to mirror only received packets. (8320, 8325, 8325P, 8325H, 10000 Switch series only)
- Support for LAG member interfaces as source for RX (ingress) mirroring (8325, 8325P, 8325H, 9300, 9300S, 10000, Switch series only)
- Supports for LAG member interfaces as source for TX (egress) mirroring | 8325, 8325P, 8325H, 9300, 10000, 10040 Switch series ony)
- Ability for a given Mirror source in one session to act as source in another Mirror Session (8100, 82xx, 83xx, 9xxx 100xx Switch series only)
- Support for a Layer 2/bridged Link Aggregation Group (LAG) as Session destination (8xxx, 100xx Switch series only)
- Support for a Layer 3/Route-only Link Aggregation Group (LAG) as Session destination (8xxx, 100xx Switch series only)

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.

Mirroring statistics and sFlow

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.

Mirroring and sFlow configuration on the same port is supported.

Limitations

The following limitations apply when configuring multiple mirroring sessions on a switch:

- CPU generated packets egressing on a routed L3 interface will not be mirrored to the destination port.
- Untagged egress packets that get mirrored will have the native VLAN tag in the mirrored packet. These extra bytes can cause traffic loss at the mirror destination when running line rate traffic.
- True egress mirroring is not supported on 832x platforms. Egress mirroring takes place at the ingress. The packets that may get dropped at the egress might also have been mirrored at ingress. Traffic will be mirrored even before the policy actions are processed at the egress.
- Packets mirrored to CPU from a Layer-3 Route Only Port (ROP) will have a VLAN tag with the VLAN ID set to the internal VLAN ID assigned to that port.
- 832x platforms have 4 mirror ASIC resources that can be used among the different mirror sessions. Each direction in a mirror session will consume 1 mirror ASIC resource. Hence, a user can have up to 4 unidirectional mirror sessions or 2 bi-directional mirror sessions active at any given time. If there are no mirror ASIC resources available when attempting to enable a mirror session, the 'Operation Status' field of `show mirror` command for session ID will have the status set to 'platform_session_limit_reached.'
- The mirror destination port among the active mirror sessions must be unique i.e. 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.

- The interface/LAG used to transmit ERSPAN packets cannot be a source in *any* mirror session.
- The interface/LAG used to transmit ERSPAN packets must be unique per ERSPAN mirror session. If a change in the L3 topology causes multiple ERSPAN mirror sessions to use the same egress interface/LAG to transmit the ERSPAN packets, then only one session will work. The other session(s) will go into an error state.

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

Applies only to the HPE Aruba Networking 8100 and 8360 Switch Series.

`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
8100 8360	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 passowrd:
Connected to 10.0.0.2.
sftp > put my_capture_file.pcap file.pcap
Uploading my_capture_file.pcap to /root/file.pcap
my_capture_file.pcap                          100%   156   219.8KB/s   00:00
Copied successfuly.
```

Command History

Release	Modification
10.08	Command introduced

Command Information

Platforms	Command context	Authority
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320 8325 8325H 8325P 8360 9300 10000	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
8100 8320 8325 8325H 8325P 8360 9300 10000	`config-mirror-<SESSION-ID>`	Administrators or local user group members with execution rights for this command.

destination interface

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

Description

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

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

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

Usage

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

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

Examples

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

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

Replacing the existing destination with different interface:

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

Removing a destination:

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

Switch	Destination interface limit per mirror session (4 possible sessions)
8320	1
8325/8325H/8325P	1
8360	64
9300/9300S	1
10000	1

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

destination tunnel

```
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
8100 8320 8325 8325H 8325P 8360 9300 10000	`config-mirror-<SESSION-ID>`	Administrators or local user group members with execution rights for this command.

diagnostic

```
diagnostic

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

Description

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

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

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

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

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

Example

Performing diagnostic:

```
switch# diagnostic

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

diag utilities tcpdump

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

Description

Captures traffic received or transmitted over a network.

Parameter	Description
`command <TEXT>`	Captures packets based on a specified tcpdump command string.
`delete file <FILE-NAME>`	Deletes specified tcpdump list files.
`list-files`	Lists all the tcpdump capture files saved on the device.
`vrf <VRF-NAME>`	Captures packets on the specified VRF. If no VRF is named, the default is used.
`count <COUNT-NUM>`	Runs the tcpdump command until the specified number of packets are captured. Range: 1-2147483647.
`proto <PROTO-NUM>`	Captures packets of a particular type based on IP protocol number. Range: 0-255.
`host-ip <IP-ADDR>`	Captures packets matching with the source or destination IP address.
`source-ip <IP-ADDR>`	Captures packets from the specified IP address.
`destination-ip <IP-ADDR>`	Captures packets sent to the specified IP address.
`host-port <PORT>`	Captures packets matching with the source or destination port.
`source-port <PORT>`	Captures packets from the specified IP port.
`destination-port <PORT>`	Captures packets sent to the specified IP port.
`verbosity <LEVEL>`	Captures packets of the specified verbosity. Range: level1-level4. If no verbosity is specified, the default is level1.
`print <DATA>`	Captures the data of each packet. The maximum is 262144 bytes
`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
8100 8320 8325 8325H 8325P 8360 9300 10000	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

Configuring and enabling a mirroring session:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

mirror session

`mirror session <SESSION-ID>`

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

Description

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

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

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

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

Examples

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

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

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

mirror endpoint

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

Description

Creates the specified mirror endpoint or enters its context if it already exists. The specifics of a mirror endpoint are created or altered while in the mirror endpoint context and the mirror endpoint is enabled or disabled from this context. It may be possible to support different encapsulations by different ASICs. For example, UDP for PVOS compatibility. Termination of GRE encapsulation is also supported.

The **no** form of this command removes an existing mirror endpoint. An enabled mirror endpoint is automatically disabled first before removal.

Parameter	Description
<NAME>	Specifies mirror endpoint name.

Examples

Creating a mirror endpoint named test :

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

Deleting mirror endpoint named test:

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

Configuring a mirror endpoint named test :

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

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

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

Command History

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

Command Information

Platforms	Command context	Authority
8100 8360	`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 Statu**s colum indicates the status of the mirroring session. Operation status is one of the following values:

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

Examples

Showing summary information about all configured mirroring sessions:

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

Showing detailed information about a single mirroring session:

```
switch# show mirror 3
 Mirror Session: 3
 Admin Status: disable
 Operation Status: disabled
 Comment: Monitor router port ingress-only traffic
 Source: interface 1/1/2 rx
 Destination: interface 1/1/3
switch#
```

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

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

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

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

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

```
switch# show mirror 1
Mirror Session: 1
Admin Status: enable
Operation Status: disabled
```

```
Source: interface 1/1/1 (unknown lag) both
Destination: cpu
Output Packets: 0
Output Bytes: 0
```

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

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

show mirror endpoint

`show mirror endpoint [<NAME>]`

Description

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

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

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

Examples

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

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

Showing the details of enabled mirror endpoint test:

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8360	Operator (>) or Manager (#)	Administrators or local user group members with execution rights for this command.

shutdown

Applies only to the HPE Aruba Networking 8100 and 8360 Switch Series.

```
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
8100 8360	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
8100 8360	`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)
8320	128
8325/8325H/8325P	128
8360	64
9300/9300S	128
10000	72

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

You can configure the same source interface in multiple mirroring sessions, if required.

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

Examples

Configuring a mirrored traffic source interface:

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

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

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

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

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

Removing the first source interface:

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

Configuring a source interface to mirror received packets only:

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

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

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

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

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

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

source vlan

Applies only to the HPE Aruba Networking 8100 and 8360 Switch Series.

```
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.

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 theingress 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
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8360	`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 *tcpdump* capture file to be copied.
<REMOTE_URL>	Specifies a URL to copy the command output. Valid URL syntaxes: ■ `sftp://USER@{IP|HOST}[:PORT]/FILE` ■ `tftp://{IP|HOST}[:PORT][;blocksize=VAL]/FILE`

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 6200 6300 6400 8100 8325 8360 9300 10000	`config`	Administrators or local user group members with execution rights for this command.

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 *tcpdump* capture file to be deleted.

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
5420 6200	`config`	Administrators or local user group members with execution rights for this command.

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 *EtherType* 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 *tcpdump* 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 6200 6300 6400 8100 8325 8360 9300 10000	`config`	Administrators or local user group members with execution rights for this command.

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 *tcpdump* capture files.

Command History

Release	Modification
10.17	Command introduced.

Command Information

Platforms	Command context	Authority
5420 6200 6300 6400 8100 8325 8360	`config`	Administrators or local user group members with execution rights for this command.

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 6200 6300 6400 8100 8325 8360 9300 10000	`config`	Administrators or local user group members with execution rights for this command.

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

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

Limitations with breakout cable support

- The JL720A Aruba 8360-48XT4C models (ordered SKU #s JL706A/JL707A) do not support split ports.

Breakout cable support commands

split

```
split [<COUNT>][<SPEED>][confirm]
no split [confirm]
```

Description

Splits a port into multiple interfaces. Only ports capable of supporting breakout cables or SR4/eSR4/eDR4 optics can be split.

Parameter	Description
<COUNT>	Specifies the number of child interfaces to activate upon splitting the port. Default: 4.
<SPEED>	Specifies the speed for the child interfaces.
`confirm`	Specifies the confirmation of port splitting.

Usage

- Some switch interfaces support different split counts depending on the installed transceiver. For these interfaces, the number of child interfaces to activate can be specified. If omitted, the default is 4. For transceivers capable of supporting multiple split modes, the closest mode with enough lanes is used.

- Some transceivers also support multiple split modes with different speeds. For example, 2x200G or 2x100G. When a speed is not specified, the highest available speed for the split count is used. To select a different split mode with a lower speed, the desired speed must be specified.

When the current transceiver does not support the configured split speed, the interface will remain down with an `Invalid speed` error.

Ports that are not splittable in the currently applied interface profile can be configured as split, but remain in a **down** state until a compatible profile is selected and the switch rebooted to apply it.

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

Model	Description	Port info
8320 Series		
■ JL479A ■ JL579A ■ JL581A	8320 48 10/6 40 X472 5 2 Bundle 8320 32 40G X472 5 2 Bundle 8320 48 T/6 40 X472 5 2 Bundle	49-54 (40G) 5-28 (40G - center 24 ports) 49-54 (40G)
8325 Series		
■ JL635A	8325-48Y8C 48p 25G 8p 100G switch 8325-48Y8C FB 6 F 2 PS Bundle	49-56 (40G or 100G)
■ JL624A ■ JL625A ■ JL626A ■ JL627A ■ JL636A	8325-48Y8C BF 6 F 2 PS Bundle JL626A 8325-32C FB 6 F 2 PS Bundle 8325-32C BF 6 F 2 PS Bundle 8325-32C 32p 100G switch	49-56 (40G or 100G) 49-56 (40G or 100G) 1-32 (40G or 100G) 1-32 (40G or 100G) 1-32 (40G or 100G)
8360 32Y4C models		
JL717A (base system) ■ JL700A Port-to-Power model ■ JL701A Power-to-Port model	8360-32Y4C switch 8360-32Y4C switch	33-36 (40G or 100G) 33-36 (40G or 100G)
8360 16Y2C models		
JL718A (base system) ■ JL702A Port-to-Power model ■ JL703A Power-to-Port model	8360-16Y2C switch 8360-16Y2C switch	17-18 (40G or 100G) 17-18 (40G or 100G)
8360 48XT4C models		
JL720A (base system) ■ JL706A Port-to-Power model ■ JL707A Power-to-Port model	8360-48XT4C switch 8360-48XT4C switch	NO SUPPORT for Split ports
8360-12C models		
JL721A (base system) ■ JL708A Port-to-Power model ■ JL709A Power-to-Port model	8360-12C switch 8360-12C switch	1-12 (40G or 100G) 1-12 (40G or 100G)
8360 24XF2C models		
JL722A (base system)	8360-24XF2C switch	25-26 (40G or 100G)

Model	Description	Port info
▪ JL710A Port-to-Power model ▪ JL711A Power-to-Port model	8360-24XF2C switch	25-26 (40G or 100G)
9300 Switch Series ▪ R8Z96A	9300-32D 32p 100/200/400G QSFP-DD p 10G SFP+ Switch	Porte 1-32 are splittable. 400G SR8 can split two, four, or eight ways (8-way splitting available on the 9300-32D only). Some ports are splittable depending on the interface profile use. See the Installation and Getting Started Guide for more information.
9300S Switch Series ▪ S0F95A ▪ S0F96A	9300S 32P QSFP28 100G 8p QSFP-DD 400G TAA Switch 9300S 32P QSFP28 100G 8p QSFP-DD 400G Switch	9-12. 17-24, and 29-32 ports are splittable. 400G SR8 can split two or four ways QSFP-DD ports are capable of operating at: ▪ 100G ports ▪ 40G ports ▪ Split into 4 individual 25G or 10G ports Some ports are splittable depending on the interface profile use. See the Installation and Getting Started Guide for more information.
10000 Switch Series ▪ R8P13A	10000-48Y6C FB6F2PS Bundle	Ports 49-54 are splittable. QSFP28 ports are capable of operating as: ▪ 100G ports ▪ 40G ports ▪ Split into 4 individual 25G or 10G ports
▪ R8P14A	10000-48Y6C BF6F2PS Bundle	Ports 49-54 are splittable. QSFP28 ports are capable of operating as: ▪ 100G ports ▪ 40G ports ▪ Split into 4 individual 25G or 10G ports

Examples

Before splitting an interface (example on a 8325 Series Switch):

```
switch(config)# show interface 1/1/56 brief
--------------------------------------------------------------------------------
Port    Native  Mode   Type      Enabled Status Reason                Speed Desc
        VLAN
--------------------------------------------------------------------------------
1/1/56  --      routed QSFP+DA1  no      down   Administratively down --    --
```

After splitting:

```
switch(config)# interface 1/1/56
switch(config-if)# split
This command will disable the specified port, clear its configuration,
and split it into multiple interfaces.

Continue (y/n)? y

8325(config-if)# show interface 1/1/56,1/1/56:1-1/1/56:4 brief
--------------------------------------------------------------------------------
Port      Native Mode   Type       Enabled Status  Reason           Speed   Desc
          VLAN                                                     (Mb/s)
--------------------------------------------------------------------------------
1/1/56    --     routed QSFP+DA1   no      down    Interface split  --      --
1/1/56:1  --     routed QSFP+DA1   yes     up                       10000   --
1/1/56:2  --     routed QSFP+DA1   yes     up                       10000   --
1/1/56:3  --     routed QSFP+DA1   yes     up                       10000   --
1/1/56:4  --     routed QSFP+DA1   yes     up                       10000   --
```

Unsplitting a port on a switch that does not require a reboot:

```
switch(config)# interface 1/1/1
switch(config-if)# no split
This command will disable the split interfaces for this port and clear
their configuration.

Continue (y/n)? y
```

Splitting an interface two ways on a 9300 Series Switch using the default for speed, taking the port capability (400G), and assuming 200G:

```
switch(config)# interface 1/1/1
switch(config-if)# split 2
This command will disable the specified port, clear its configuration,
and split it into multiple interfaces.

Continue (y/n)? y

switch(config-if)# show interface brief
--------------------------------------------------------------------------------
Port     Native  Mode   Type       Enabled Status  Reason    Speed   Description
         VLAN                                               (Mb/s)
--------------------------------------------------------------------------------
1/1/1:1   --     routed 400G-SR8    yes     up              200000
1/1/1:2   --     routed 400G-SR8    yes     up              200000
```

Changing the interface to 2x100G mode:

```
switch(config)# interface 1/1/1
switch(config-if)# split 2 100g
This command will clear the configuration for all split interfaces of
this port.

Continue (y/n)? y

switch(config-if)# show interface brief
--------------------------------------------------------------------------------
Port     Native  Mode   Type       Enabled Status  Reason    Speed   Description
         VLAN                                                (Mb/s)
--------------------------------------------------------------------------------
1/1/1:1   --     routed 400G-SR8    yes     up              100000
1/1/1:2   --     routed 400G-SR8    yes     up              100000
```

Command History

Release	Modification
10.10.1000	Added parameters: ***<COUNT>*** , **<SPEED>**
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8325H 8325P 8360 9300 10000	config-if	Administrators or local user group members with execution rights for this command.

Chapter 12
Aruba AirWave

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

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

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

SNMP support and AirWave

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

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

SNMP on the switch

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

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

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

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

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

Supported features with AirWave and the AOS-CX switch

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

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

Configuring the AOS-CX switch to be monitored by AirWave

Prerequisites

Aruba AirWave is active on the network.

Procedure

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

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

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

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

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

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

switch(config)# snmpv3 context Admin
```

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

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

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

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

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

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

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

AirWave commands

logging

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

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

Description

Enables syslog forwarding to a remote syslog server.

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

Parameter	Description
{*<IPV4-ADDR>* | *<IPV6-ADDR>* | *<HOSTNAME>*}	Selects the IPv4 address, IPv6 address, or host name of the remote syslog server. Required.
[udp [*<PORT-NUM>*] | tcp [*<PORT-NUM>*]]	Specifies the UDP port or TCP port of the remote syslog server to receive the forwarded syslog messages.
udp [*<PORT-NUM>*]	Range: 1 to 65535. Default: 514
tcp [*<PORT-NUM>*]	Range: 1 to 65535. Default: 1470
tls [*<PORT-NUM>*]	Range: 1 to 65535. Default: 6514
include-auditable-events	Specifies that auditable messages are also logged to the remote syslog server.
severity *<LEVEL>*	Specifies the severity of the syslog messages: ■ alert: Forwards syslog messages with the severity of alert (6) and emergency (7). ■ crit: Forwards syslog messages with the severity of critical (5) and above. ■ debug: Forwards syslog messages with the severity of debug (0) and above. ■ emerg: Forwards syslog messages with the severity of emergency (7) only. ■ err: Forwards syslog messages with the severity of err (4) and above ■ info: Forwards syslog messages with the severity of info (1) and above. Default. ■ notice: Forwards syslog messages with the severity of notice (2) and above. ■ warning: Forwards syslog messages with the severity of warning (3) and above.
auth-mode	Specifies the TLS authentication mode used to validate the certificate. ■ certificate: Validates the peer using trust anchor certificate based authentication. Default. ■ subject-name: Validates the peer using trust anchor certificates as well as subject-name based authentication.
legacy-tls-renegotiation	Enables the TLS connection with a remote syslog server supporting legacy renegotiation.
vrf *<VRF-NAME>*	Specifies the VRF used to connect to the syslog server. Optional. Default: default

Examples

Enabling the syslog forwarding to remote syslog server 10.0.10.2:

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

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

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

Disabling syslog forwarding to a remote syslog server:

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

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
All platforms	`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
<VRF-NAME>	Specifies the VRF on which the SNMP agent listens for incoming requests. The SNMP agent can listen on either the `mgmt` or `default` VRF. If configured for both, the SNMP agent listens on `default`, which has a higher priority.

Example

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

snmpv3 context

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

Description

Creates an SNMPv3 context on the specified VRF.

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

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

Parameter	Description
	characters, excluding space and question mark (?).
`vrf <VRF-NAME>`	Specifies the VRF associated with the context. Default: default.
`community <STRING>`	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
<NAME>	Specifies the SNMPv3 username. Range 1 - 32 printable ASCII characters, excluding space and question mark.
`auth` *<AUTH-PROTOCOL>*	Specifies the authentication protocol used to validate user logins. Available options are: `md5` or `sha`.
`auth-pass {plaintext | ciphertext}` *<AUTH-PWORD>*	Specifies the SNMPv3 user password. Range for `plaintext` is 8 - 32 printable ASCII characters, excluding space and question mark. Range for `ciphertext` 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.
`priv` *<PRIV-PROTOCOL>*	Specifies the SNMPv3 security protocol (encryption method). Available options are: `aes` or `des`.
`priv-pass {plaintext | ciphertext}` *<PRIV-PWORD>*	Specifies the SNMPv3 user privacy passphrase. Range for `plaintext` is 8 - 32 printable ASCII characters, excluding space and question mark. Range for `ciphertext` 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.

Chapter 13

Support and Other Resources

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./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.