AOS-CX 10.12 Monitoring Guide

8320, 8100, 8325, 8360, 9300, 10000 Switch Series

a Hewlett Packard Enterprise company

Published: May 2023
Edition: 1

Copyright Information

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

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

Notices

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

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

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

Acknowledgment

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

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

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

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

UNIX® is a registered trademark of The Open Group.

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

Chapter 1
About this document

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

Applicable products

This document applies to the following products:

- Aruba 8100 Switch Series (R9W94A, R9W95A, R9W96A, R9W97A)
- Aruba 8320 Switch Series (JL479A, JL579A, JL581A)
- Aruba 8325 Switch Series (JL624A, JL625A, JL626A, JL627A)
- Aruba 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)
- Aruba 9300 Switch Series (R9A29A, R9A30A, R8Z96A)
- Aruba 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. ■ 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.

Convention	Usage
`|`	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 83xx, 9300, and 10000 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 to 4 x 10G or 4 x 25G. 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.

Diagnosing with the LEDs

This section describes LED patterns on the switch that indicate problem conditions for general switch operation troubleshooting.

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.

1. Check the table for the LED pattern you see on the switch.
2. Refer to the corresponding diagnostic tip.

Table 1: *LED error indicators for 8320*

Global status	Port LED	Diagnostic tip
Off with power cord plugged in	N/A	1
Solid amber	N/A	2
Slow flash amber	N/A	3
Slow flash amber	Slow flash amber*	4
Solid green	Off with cable connected	5
Solid green	On, but the port is not communicating	6

*The flashing behavior is an on/off cycle approximately once every 1.6 seconds.

Table 2: *LED error indicators for 8325*

PS1/PS2 LEDs	Global Status	Fan	Port LED	Diagnostic Tip
Off with power cords plugged in	-	-	-	1
On amber**	Flashing amber	-	-	2
On green	Flashing amber	On amber	-	3
On green	Flashing amber	-	Flashing amber	4
On green	On green	-	Off with cable connected	5

PS1/PS2 LEDs	Global Status	Fan	Port LED	Diagnostic Tip
On green	On green	-	On, but the port is not communicating	6

**Either the PS1 or PS2 LED is on amber, but not both.

Table 3: *Diagnostic tips*

Tip	Problem	Solution
1	Both switch power supplies are not plugged into an active AC power source.	Verify the AC power source works by plugging another device into the outlet. Or try plugging the power supplies into different outlets or try different power cords. If the problem is still not resolved, both power supplies may be faulty.
2	One of the power supplies is not plugged into an active A power source, or the power supply may have failed.	Verify that the power cord is plugged into an active power source and to the power supply. Make sure that the connections are snug. Try power cycling the switch by unplugging and plugging the power cord back into the other working power supply. If the PS1/PS2 LED is still not on, verify the AC power source works by plugging another device into the outlet or try a different power cord. If the power source and power cord are OK and this condition persists, the switch power supply may have failed. Call your Hewlett Packard Enterprise-authorized network reseller, or use the electronic support services from Hewlett Packard Enterprise to get assistance.
3	One of the switch fan assemblies may have failed.	Try disconnecting power from the switch and wait a few moments. Then reconnect the power to the switch and check the LEDs again If the error indication reoccurs, one of the fan assemblies has failed. If the ambient temperature does not exceed normal room temperature, the switch may continue to operate under this condition; but for best operation, replace the fan assembly. Call your Hewlett Packard Enterprise-authorized network reseller, or use the electronic support services from Hewlett Packard Enterprise to get assistance.
4	The network port for which the LED is flashing has experienced a self-test or initialization failure.	Try power cycling the switch. If the fault indication reoccurs: ▪ There may be a port configuration mismatch where

Tip	Problem	Solution
		a 10G transceiver is installed in a port configured for 25G, or the reverse. ■ A 10GBase-T transceiver may be installed in an incompatible port. Only ports 1, 2, 4, 5, 7, 8, 10, and 11 support 10GBase-T transceivers. ■ The transceiver may have failed. ■ The switch port may have failed. Check the switch Event Log and show interface command output for indication of the fault condition. If the port is an SFP+/SFP28 transceiver or QSFP+/QSFP28 transceiver, verify that it is one of the transceivers supported by the switch. Unsupported or unrecognized transceivers will be identified with this fault condition. For a list of supported transceivers, see the *Transceiver Guide* in the Aruba Support Portal. The transceivers are also tested when they are "hot-swapped" - installed or changed while the switch is powered on. To verify that the port has failed, remove and reinstall the transceiver without powering off the switch. If the port fault indication reoccurs, you will have to replace the transceiver. Check the event log to see why the transceiver failed. To get assistance, call your Hewlett Packard Enterprise-authorized network reseller, or use the electronic support services from Hewlett Packard Enterprise.
5	The network connection is not working properly.	Try the following procedures: ■ For the indicated port, verify that both ends of the cabling, at the switch and the connected device, are connected properly. ■ Verify that the connected device and switch are both powered on and operating correctly. ■ Verify that you have used the correct cable type for the connection: ◦ For fiber-optic connections, verify that the transmit port on the switch is connected to the receive port on the connected device and that the switch receive port is connected to the transmit port on the connected device. ◦ The cable verification process must include all patch cables from any end devices, including the switch, to any patch panels in the cabling path. ■ Verify that the port has not been disabled through

Tip	Problem	Solution
		a switch configuration change. Use the console interface or, if you have configured an IP address on the switch, use the web browser interface to determine the state of the port and re-enable the port if necessary. ■ Verify that the switch port configuration matches the configuration of the attached device. For example, if the switch port is configured as "Full-duplex", the port on the attached device also MUST be configured as "Full-duplex". If the configurations do not match, the results could be an unreliable connection, or no link at all. ■ If the other procedures do not resolve the problem, try using a different port or a different cable.
6	The port may be improperly configured, or the port may be in a "blocking" state by the normal operation of the Spanning Tree, LACP, or IGMP features.	Use the switch console to see if the port is part of a dynamic trunk (through the LACP feature), if Spanning Tree is enabled on the switch, and if the port may have been put into a "blocking" state by those features. The `show lacp interfaces` command displays the port status for the LACP feature; the show spanning tree command displays the port status for Spanning Tree. Also check the Port Status screen using the `show interfaces` command to see if the port has been configured as "disabled". Other switch features that may affect the port operation include VLANs, IGMP, and port group settings. Use the switch console to see how the port is configured for these features. Ensure that the device at the other end of the connection is indicating a good link to the switch. If it is not, the problem may be with the cabling between the devices or the connectors on the cable.

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 external flow collectors. IPFIX is configurable via CLI or REST. 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.

Compatibility with Application Recognition and Traffic Insight

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

Supported Platform

The following table list the scales and supported platforms for IPFIX.

Table 1: *Scale and supported platforms for IPFIX*

Platform	IPFIX	Maximum Flows	Maximum pps	Maximum TCAM
8100	Yes	15,536	2,500	16,384
8360	Yes	64,688	2,500	65,536

TCAM is shared between multiple features and is allocated based on first-in, first-out principle. This could result in a scale impact if the environment already has other features using TCAM and the usage is going beyond the available or allocated TCAM limit.

- In case of TCAM overflow, an error or warning message is displayed under the event logs.

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. A maximum of sixteen flow monitors can be created. There is a limit of two flow exporters that can be applied to a single flow monitor.

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. A single flow exporter can be assigned to one or more flow monitors, and multiple flow exporters can be assigned to a single flow monitor.

Destinations

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

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

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

Flow Records

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

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

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

The supported sets of match fields are:

1. All IPv4:

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

2. All IPv6:

- IPv6 version
- IPv6 destination address

- IPv6 protocol
- Transport destination port
- Transport source port

Configuring IP Flow Information Export

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

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

Step one: Create Flow Records

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

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

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

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

```
switch(config)# show flow record
--------------------------------------------------------------------------------
Flow record 'flowRecordv4'
--------------------------------------------------------------------------------
Match Fields
ipv4 destination address
ipv4 protocol
ipv4 source address
ipv4 version
transport destination port
```

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

-------------------------------------------------------------------------------
Flow record 'flowRecordv6'
-------------------------------------------------------------------------------
Match Fields
ipv6 destination address
ipv6 protocol
ipv6 source address
ipv6 version
transport destination port
transport source port
Collect Fields
application name
counter bytes
counter packets
timestamp absolute first
timestamp absolute last
```

Step two: Configure flow exporter(s)

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

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

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

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

Next, configure the flow exporter for Traffic Insight

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

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

```
switch(config)# show flow exporter flowExpTI

------------------------------------------------------------------------------
Flow exporter 'flowExpTI'
------------------------------------------------------------------------------
Status                   : Accepted
Export Protocol          : ipfix
Destination Type         : Traffic Insight
Destination              : TI
Transport Configuration
Protocol             : udp
Port                 : 4739
```

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

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

Step three: Configure a monitor(s)

First, configure an IPv4 flow monitor.

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

Next, configure an IPv6 flow monitor.

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

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

```
switch(config-flow-monitor)# show flow monitor
-------------------------------------------------------------------------------
Flow monitor 'flowMonv4'
```

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

-------------------------------------------------------------------------------
Flow monitor 'flowMonv6'
-------------------------------------------------------------------------------
Status                    : Accepted
Flow Record               : flowRecordv6
Flow Exporter(s)          : flowExternal
Cache Configuration
Inactive Timeout     : 30
Active Timeout       : 1800
```

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

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

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

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

Next, apply flow monitor to IPv4 and IPv6 interfaces

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

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

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

FAQs and Troubleshooting

- When IPFIX is used with Application Recognition, these features do not support LAGs or MCLAGs (VSX LAGs).
- The following messages are displayed to indicate an illegal argument:
  - % The flow exporter <EXPORTER-NAME> does not exist.
  - % The flow record <RECORD-NAME> does not exist.
  - % The flow monitor <MONITOR-NAME> does not exist.
  - Invalid destination IP address or hostname entered.
  - Unable to create the flow exporter. The maximum allowed number of flow exporters (16) has been reached.
  - Unable to create the flow record. The maximum allowed number of flow records (16) has been reached.
  - Unable to create the flow monitor. The maximum allowed number of flow monitors (16) 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

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}
   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. Each flow monitor supports a maximum of two different flow exporter configurations, sending flow records to up to two destinations.

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.

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

Examples

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

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

Related Commands

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

Command History

Release	Modification
10.11	Command introduced.

Command Information

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

flow monitor

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

Description

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.

Parameter	Description
`<name>`	Name of the flow monitor , up to 64 characters.
`cache timeout active|inactive <timeout>`	Use the cache timeout parameter to define an active or inactive timeout for the flow monitor. A flow monitor closes a flow session that is active for longer than the active timeout or inactive for longer than the inactive timeout. The supported timeout ranges for both the active timeout and inactive timeout are 30-604800 seconds, and the default is 30 seconds.
`description`	A description up to 256 characters long, including spaces.
`exporter <name>`	Assign a flow exporter to a flow monitor. Each flow monitor supports a maximum of two different flow exporters, sending flow records to up to two destinations.
`record <name>`	Assigns a flow record to a flow monitor.

Examples

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

```
switch(config)# flow monitor monitor-1
switch(config-flow-monitor)# description Monitor for analyzing basic ipv4 traffic
switch(config-flow-monitor)# exporter flow-exporter-1
switch(config-flow-monitor)# exporter flow-exporter-2
switch(config-flow-monitor)# record flow-record-1
switch(config-flow-monitor)# cache timeout inactive 300
switch(config-flow-monitor)# cache timeout active 1500
```

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

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

Related Commands

Command	Description
flow exporter	Define how a flow monitor exports the flow reports.
flow record	Define data to be included in a flow record by configuring flow record match and collect fields
flow monitor	Enable flow monitoring on inbound traffic coming into an interface by assigning a flow monitor to that interface.

Command History

Release	Modification
10.11	Command introduced.

Command Information

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

flow record

```
flow record <name>
   match
      ipv4|ipv6 {protocol|version}|{source|destination address}
      transport {source|destination} port
   collect
      application name
      counter {packets|bytes}
      timestamp absolute {first|last}
   description <description>
```

Description

Define data to be included in a flow record by configuring flow record match and collect fields. The **match** attributes define what makes the traffic flow unique. Traffic with matching attributes (for example, traffic coming from the same interface, sent to the same destination with the same protocol) are classified as a single flow. Information for some or all of the matched settings can be collected and exported to a destination defined by the flow exporter assigned to the flow monitor.

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

Parameter	Description
`<name>`	Name of the flow monitor , up to 64 characters.
`match`	match traffic according to one or more of the following key attributes: ■ **ipv4**: match traffic on an IPv4 network ■ **ipv6**: match traffic on an IPv6 network ■ **protocol**: Match traffic using the same IP protocol ■ **version**: Match traffic using the same IP version ■ **source**: Match traffic from the same source ■ **destination**: Match traffic to the same destination ■ **address**: Match traffic by source or destination IP address ■ **transport**: Match traffic by source or destination transport type ■ **port**: Match traffic by source or destination transport port
`description`	A description for the flow record up to 256 characters long, including spaces
`collect`	Configures data fields to be included a flow record. ■ **application name:** Include the application name as a non-key field in a flow record ■ **counter packets**: Collect counter data for packets in the flow ■ **counter bytes**: Collect counter data for bytes in the flow ■ **timestamp absolute first**: Collect absolute timestamp of the first packet observed.

Examples

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

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

Removing the IPv4 destination match field from the flow record defined in the previous example.

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

Adding counter and timestamp collect fields to flow record flow-record-1.

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

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

Related Commands

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

Command History

Release	Modification
10.11	Command introduced.

Command Information

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

ipv4|ipv6 flow monitor

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

Description

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

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

Examples

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

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

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

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

Related Commands

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

Command History

Release	Modification
10.11	Command introduced.

Command Information

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

show flow exporter

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

Description

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

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

- Accepted
- Rejected (Internal error: exporter does not exist)
- Rejected (Internal error: destination type does not exist)
- Rejected (Destination type is 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 (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)

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

Examples

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

```
switch# show flow exporter exporter-1
-------------------------------------------------------------------------------
Flow exporter 'exporter-1'
-------------------------------------------------------------------------------
Description             : Exports to the first collector
Status                  : Accepted
Export Protocol         : ipfix
Destination Type        : Hostname or IP address
Destination             : 192.168.0.1
Transport Configuration
    Protocol            : UDP
    Port                : 9995
```

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

-------------------------------------------------------------------------------
Flow exporter 'exporter-1'
-------------------------------------------------------------------------------
Reports sent            : 14961
```

Related Commands

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

Command History

Release	Modification
10.11	Command introduced.

Command Information

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

show flow monitor

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

Description

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

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

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

Parameter	Description
`<name>`	Name of the flow monitor.
`statistics`	Display additional flow and cache statistics.

Examples

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

```
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
```

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

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

Related Commands

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

Command History

Release	Modification
10.11	Command introduced.

Command Information

Platforms	Command context	Authority
8100 8360	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

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

```
switch# show flow record record-1
-------------------------------------------------------------------------------
Flow record  'record-1'
-------------------------------------------------------------------------------
Description             : Used for IPv4 traffic analysis
Status                  : Accepted
Match Fields
    ipv4 destination address
```

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

Related Commands

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

Command History

Release	Modification
10.11	Command introduced.

Command Information

Platforms	Command context	Authority
8100 8360	config config-flow-exporter	Administrators or local user group members with execution rights for this command.

show tech ipfix

```
show tech ipfix
```

Description

Shows the IPFIX configuration settings.

Examples

The example shows the IPFIX configuration settings.

```
switch#show tech ipfix
=====================================================
Show Tech executed on Tue Apr 11 02:43:06 2023
=====================================================
=====================================================
[Begin] Feature ipfix
=====================================================
*********************************
Command : show flow exporter
*********************************
```

```
--------------------------------------------------------------------------------
Flow exporter 'ipfix'
--------------------------------------------------------------------------------
Status                  : Accepted
Export Protocol         : ipfix
Destination Type        : Traffic Insight
Destination             : t1
Transport Configuration
Protocol            : udp
Port                : 4739
--------------------------------------------------------------------------------
Flow exporter 'V6E1'
--------------------------------------------------------------------------------

....

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

Command History

Release	Modification
10.11	Command introduced.

Command Information

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

diag-dump ipfix basic

```
diag-dump ipfix basic
```

Description

Displays diagnostic information for IPFIX.

Examples

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

....

:- IPFIX Monitor v6ti completed -
- End of IPFIX Monitor Cache dump -
```

```
--------------------------------------------------------------------------------
[End] Daemon ipfixd
--------------------------------------------------------------------------------
--------------------------------------------------------------------------------
[Start] Daemon ops-switchd
--------------------------------------------------------------------------------
Key format: <traffic_type>_<coalescence_id>_<agent_id>_<asic_port>
Key                              TCAM Entry ID    Count
-------------------------------- ---------------- -----
1_1532781829_3_20                0xffff7c7e7a00   1
1_3217499901_1_12                0xffff91187580   1
1_3217499901_1_13                0xffff91183d80   1
1_3217499901_1_14                0xffff91186e80   1

....

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

Command History

Release	Modification
10.11	Command introduced.

Command Information

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

boot set-default

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

Description

Sets the default operating system image to use when the system is booted.

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

Example

Selecting the primary image as the default boot image:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

boot system

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

Description

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

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

Usage

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

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

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

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

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

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

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

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

Examples

Rebooting the system from the configured default operating system image:

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

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

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

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

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

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

Canceling a system reboot:

```
switch# boot system

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

This will reboot the entire switch and render it unavailable
until the process is complete.
Continue (y/n)? n
Reboot aborted.
switch#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

show boot-history

`show boot-history [all]`

Description

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

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

Usage

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

`Index`

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

`Boot ID`

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

`Current Boot, up for <SECONDS> seconds`

For the current boot, the `show boot-history` command shows the number of seconds the module has been running on the current software.

`Timestamp boot reason`

For previous boot operations, the `show boot-history` command shows the time at which the operation occurred and the reason for the boot. The reason for the boot is one of the following values:

`<DAEMON-NAME>` `crash`

The daemon identified by <*DAEMON-NAME*> caused the module to boot.

`Kernel crash`

The operating system software associated with the module caused the module to boot.

`Reboot requested through database`

The reboot occurred because of a request made through the CLI or other API.

`Uncontrolled reboot`

The reason for the reboot is not known.

Examples

Showing the boot history of the active management module:

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

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

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

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

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

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

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

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

Index : 2
```

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

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

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

Chapter 5
Switch system and hardware commands

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

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

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

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

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

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

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

The system supports server connection through data and management ports.

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

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

External storage commands

address

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

Description

Specifies the NAS IP address or hostname.

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

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

Examples

Creating the logfiles storage volume with IP address 10.1.1.1:

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

Deleting an external storage volume named logfiles:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 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 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 8325 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.

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

Platforms	Command context	Authority
8360 9300 10000		

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 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 8360 9300 10000	config-external-storage-*<VOLUME-NAME>*	Administrators or local user group members with execution rights for this command.

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

IP-SLA guidelines

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

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

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

Limitations with VoIP SLAs

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

IP-SLA commands

http

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

Description

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

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

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

Examples

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 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>|<HOSTNAME>} [source {<SOURCE-IPV4-ADDR> | <IFNAME>}]
     [name-server <IPV4-ADDR-DNS-SERVER>] [payload-size <PAYLOAD-SIZE>]
     [tos <TYPE-OF-SERVICE>] [probe-interval <PROBE-INTERVAL>]
```

Description

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

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

Parameter	Description
`name-server <IPV4-ADDR-DNS-SERVER>`	Specifies the DNS server for destination hostname resolution.
`payload-size <PAYLOAD-SIZE>`	Specifies the payload size of an SLA probe. Range: 0 to 1440.
`tos <TYPE-OF-SERVICE>`	Specifies the type of serve to be used in the probe packets. Range: 0 to 255.
`probe-interval <PROBE-INTERVAL>`	Specifies the probe interval in seconds. Range: 5 to 604800.

Examples

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8360 9300 10000	`config-ip-sla-<IP-SLA-NAME>`	Administrators or local user group members with execution rights for this command.

ip-sla

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

Description

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

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

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

Examples

Creating an IP-SLA:

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

Deleting an IP-SLA:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 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> | <IFNAME>}][vrf <VRF-NAME>]
no ip-sla responder <SLA-NAME> {udp-echo | tcp-connect | udp-jitter-voip} <PORT-NUM>
     [source {<SOURCE-IPV4-ADDR> | <IFNAME>}][vrf <VRF-NAME>]
```

Description

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

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

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

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

Examples

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8360 9300 10000	config	Administrators or local user group members with execution rights for this command.

show ip-sla responder

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

Description

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

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

Examples

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8360 9300 10000	config	Administrators or local user group members with execution rights for this command.

show ip-sla responder results

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

Description

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

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

Examples

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

    IP-SLA Type         : Udp-echo
    VRF Name            : Default
    Source IP           : 2.2.2.1
    Source Port         : 8000
    Responder Port      : 8888
    Responder IP        : 2.2.2.3
```

```
    Responder Interface :
    Responder Status    : Running
    Packets Received    : 2
    Packets Sent        : 2
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 8360 9300 10000	`config`	Administrators or local user group members with execution rights for this command.

show ip-sla *<SLA-NAME>*

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

Description

Shows the given IP-SLA source configuration and status.

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

Examples

```
switch# show ip-sla xyz results

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

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

```
       Reception Error                 : 0
       Transmission Error              : 0

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

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

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


switch(config)# show ip-sla jitter-sla results
    IP-SLA session status
       IP-SLA Name                     : jitter-sla
       IP-SLA Type                     : udp-jitter-voip
       Destination Host Name/IP Address: 2.2.2.1
       Destination Port                : 8888
       Source IP Address/IFName        :
       Source Port                     : 5555
       Status                          : Running

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

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

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

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

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

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

    Voice Scores:
       MOS  Score                      : 4.38   ICPIF                  : 0


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


switch(config)# show ip-sla http-sla
    IP-SLA Name             : http-sla
    Status                  : Running
    IP-SLA Type             : http
    VRF                     : ipslasrc
    Source IP               : 2.2.2.2
    Source Interface        :
    Domain Name Server      : 10.10.10.2
    Probe Interval(seconds) : 90
    HTTP Request Type       : GET
    HTTP/HTTPS URL          : abcd.com/ws/home
    Cache                   : Enabled
    HTTP Proxy URL          :
    HTTP Version Number     : 1.1
    ```

IP-SLA status description
    ```
    | Status                  | Description                                          |
    |-------------------------|------------------------------------------------------|
    | Running                 | SLA is fully operational                             |
    | Bind Error              | Another service is using the same source port  |
    | Interface Down          | Interface status is not up
    | Dns Resolution Error    | Failed to resolve destination hostname         |
    | No Route                | No available route to the responder            |
    | Internal Error          | Unexpected error prevents SLA session          |
    | Disabled                | SLA is disabled                                |
    |Configuration Incomplete | Configuration is not complete to enable the SLA|
    ```
IP SLA session cumulative counters description
    ```
    | Status                          | Description
                              |
    |---------------------------------|-----------------------------------------
```

```
------------------------------|
    |Probes Timed-out               | Total numbers of probes failed to receive
response.                      |
    |Bind Error                     | Total numbers of probes transmission failed
as source port not available.|
    |Destination Address Unreachable | Total numbers of probes transmission failed
due to route unavailable.    |
    |DNS Resolution Failures        | Total numbers of probes failed due to DNS
resolution failure.            |
    |Reception Error                | Total numbers of probes failed due to
internal error in reception.       |
    |Transmission Error             | Total numbers of probes failed due to
internal errr in transmission.     |
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 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 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 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> | <HOSTNAME>} <PORT-NUM> [source {<SOURCE-IPV4-ADDR> |
     <IFNAME>} [source-port <PORT-NUM>]] [name-server <IPV4-ADDR-DNS-SERVER>]
     [probe-interval <PROBE-INTERVAL>]
```

Description

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

Parameter	Description
{*<DEST-IPV4-ADDR>* | *<HOSTNAME>*}	Selects the destination IPv4 address for the IP-SLA or the hostname of the destination.

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

Examples

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 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>|<HOSTNAME>} <PORT-NUM> [source {<SOURCE-IPV4-ADDR> |
     <IFNAME>} [source-port <PORT-NUM>]] [name-server <IPV4-ADDR-DNS-SERVER>] [payload-
size
     <PAYLOAD-SIZE>] [tos <TYPE-OF-SERVICE>] [probe-interval <PROBE-INTERVAL>]
```

Description

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

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

Examples

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 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> | <HOSTNAME>} <PORT-NUM> [codec-type <CODEC-TYPE>]
     [advantage-factor <VALUE>] [source {<SOURCE-IPV4-ADDR> | <IFNAME>} [source-port
<PORT-NUM>]]
     [name-server <IPV4-ADDR-DNS-SERVER>][probe-interval <PROBE-INTERVAL>] [tos <TYPE-OF-
SERVICE>]
```

Description

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

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

Examples

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

```
codec-type g711a source 1/1/1
    switch(config-ipsla-1)# udp-jitter-voip  https://device.arubanetworks.com 8080
advantage-factor 10 codec-type g711a source 2.2.2.1
    switch(config-ipsla-1)# udp-jitter-voip  https://device.arubanetworks.com 8080
advantage-factor 10 codec-type g711a source 1/1/1
    switch(config-ipsla-1)# udp-jitter-voip  https://device.arubanetworks.com 8080
advantage-factor 10 codec-type g711a name-server 10.10.10.2 probe-interval 120
source 10.1.1.1 source-port 8888 tos 10
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 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.

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 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.
`human-readable`	Shows statistics rounded to the nearest power of 1000, for example, 1K, 345M, 2G. This is available only in the CLI interface output.
`non-zero`	Shows only non zero statistics.
`LAG`	Shows LAG interface information.
`monitor`	Continuously monitor interface statistics.
`LOOPBACK`	Shows loopback interface information.
`TUNNEL`	Shows tunnel interface information.
`VLAN`	Shows VLAN interface information.
<LAG-ID>	Specifies the LAG number. Range: 1-256
<LOOPBACK-ID>	Specifies the LOOPBACK number. Range: 0-255

Parameter	Description
<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 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 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.

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

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
...
```

Command History

Release	Modification
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.

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

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

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

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

Mirroring 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 Aruba 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 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>*]

Parameter	Description
	[;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 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
All platforms	`config-mirror-<SESSION-ID>`	Administrators or local user group members with execution rights for this command.

destination interface

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

Description

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

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

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

Usage

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

Configuring a different destination interface in an enabled mirroring session causes all mirrored traffic to use the new destination interface. This action might cause a temporary suspension of mirrored source traffic during the reconfiguration.

Examples

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

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

Replacing the existing destination with different interface:

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

Removing a destination:

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

Switch	Destination interface limit per mirror session (4 possible sessions)
8320	1
8325	1
8360	64
9300	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
```

Replacing the existing destination with a different VRF:

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

Removing the destination:

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

Platforms	Command context	Authority
8100 8320 8325 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

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

Usage

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

Examples

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

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

Listing saved capture files:

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

Reading my_capture_file.pcap:

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

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

Removing my_capture_file.pcap:

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

Command History

Release	Modification
10.08	Command introduced

Command Information

Platforms	Command context	Authority
8100 8320 8325 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)#
```

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

Applies only to the Aruba 8100 and 8360 Switch Series.

```
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
```

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.

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

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

Usage

Admin Status indicates the configured status. Admin Status is one of the following values:

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

Operation Status indicates the status of the mirroring session. Operation Status is one of the following values:

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

Examples

Showing summary information about all configured mirroring sessions:

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

Showing detailed information about a single mirroring session:

```
switch# show mirror 3
 Mirror Session: 3
 Admin Status: disable
 Operation Status: disabled
 Comment: Monitor router port ingress-only traffic
 Source: interface 1/1/2 rx
 Destination: interface 1/1/3
 Output Packets: 0
 Output Bytes: 0
switch#
```

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

show mirror endpoint

Applies only to the Aruba 8100 and 8360 Switch Series.

```
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 audit:

```
switch# show mirror endpoint audit
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: 1/1/1-1/1/10,lag1
Output Packets: 123456789
Output Bytes: 8912345678
```

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 Aruba 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

Applies only to the Aruba 8100 and 8360 Switch Series.

```
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	128
8360	64
9300	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 Aruba 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.

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 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.

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

Model	Description	Port info
Aruba 8320 Series ■ JL479A ■ JL579A ■ JL581A	Aruba 8320 48 10/6 40 X472 5 2 Bdl Aruba 8320 32 40G X472 5 2 Bdl Aruba 8320 48 T/6 40 X472 5 2 Bdl	49-54 (40G) 5-28 (40G - center 24 ports) 49-54 (40G)
Aruba 8360 32Y4C models JL717A (base system) ■ JL700A Port-to-Power model ■ JL701A Power-to-Port model	Displayed by `show system` Aruba 8360-32Y4C switch Aruba 8360-32Y4C switch	33-36 (40G or 100G) 33-36 (40G or 100G)
Aruba 8360 16Y2C models JL718A (base system) ■ JL702A Port-to-Power model ■ JL703A Power-to-Port model	Displayed by `show system` Aruba 8360-16Y2C switch Aruba 8360-16Y2C switch	17-18 (40G or 100G) 17-18 (40G or 100G)
Aruba 8360 48XT4C models JL720A (base system) ■ JL706A Port-to-Power model ■ JL707A Power-to-Port model	Displayed by `show system` Aruba 8360-48XT4C switch Aruba 8360-48XT4C switch	NO SUPPORT for Split ports
Aruba 8360-12C models JL721A (base system) ■ JL708A Port-to-Power model ■ JL709A Power-to-Port model	Displayed by `show system` Aruba 8360-12C switch Aruba 8360-12C switch	1-12 (40G or 100G) 1-12 (40G or 100G)
Aruba 8360 24XF2C models JL722A (base system) ■ JL710A Port-to-Power model ■ JL711A Power-to-Port model	Displayed by `show system` Aruba 8360-24XF2C switch Aruba 8360-24XF2C switch	25-26 (40G or 100G) 25-26 (40G or 100G)
Aruba 8325 (JL635A)	Aruba 8325-48Y8C 48p 25G 8p 100G Swch	49-56 (40G or 100G)
Aruba 8325 (JL624A)	Aruba 8325-48Y8C FB 6 F 2 PS Bdl	49-56 (40G or 100G)
Aruba 8325 (JL625A)	Aruba 8325-48Y8C BF 6 F 2 PS Bdl	49-56 (40G or 100G)
Aruba 8325 (JL626A)	JL626A Aruba 8325-32C FB 6 F 2 PS Bdl	1-32 (40G or 100G)
Aruba 8325 (JL627A)	Aruba 8325-32C BF 6 F 2 PS Bdl	1-32 (40G or 100G)
Aruba 8325 (JL636A)	Aruba 8325-32C 32p 100G Swch	1-32 (40G or 100G)

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:

```
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
1/1/2:1   --     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
1/1/2:1   --     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 8360 9300 10000	config-if	Administrators or local user group members with execution rights for this command.

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

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

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

SNMP support and AirWave

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

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

SNMP on the switch

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

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

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

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

```
!
...
```

Supported features with AirWave and the AOS-CX switch

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

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

Configuring the AOS-CX switch to be monitored by AirWave

Prerequisites

Aruba AirWave is active on the network.

Procedure

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

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

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

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

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

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

switch(config)# snmpv3 context Admin
```

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

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

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

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

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

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

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

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.

Parameter	Description
`trap version v3`	Specifies the trap notification type for SNMPv3.
`user <NAME>`	Specifies the SNMPv3 user name to be used in the SNMP trap notifications.
`community <STRING>`	Specifies the name of the community string to use when sending trap notifications. Range: 1 - 32 printable ASCII characters, excluding space and question mark. Default: `public`.
<UDP-PORT>	Specifies the UDP port on which notifications are sent. Range: 1 - 65535. Default: 162.
`vrf <VRF-NAME>`	Specifies the name of the VRF on which to send the notifications.

Examples

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

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

Command History

Release	Modification
10.07 or earlier	--

Command Information

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

snmp-server vrf

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

Description

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

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

Parameter	Description
<VRF-NAME>	Specifies the VRF on which the SNMP agent listens for incoming requests. The SNMP agent can listen on either the `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 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.

Accessing Aruba Support

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

Be sure to collect the following information before contacting Support:

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

Other useful sites

Other websites that can be used to find information:

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

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

Accessing Updates

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

Aruba Support Portal

https://asp.arubanetworks.com/downloads

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

My Networking

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

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

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

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

Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

To subscribe to eNewsletters and alerts:

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

Warranty Information

To view warranty information for your product, go to https://www.arubanetworks.com/support-services/product-warranties/.

Regulatory Information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at https://www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Aruba is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements, environmental data (company programs, product recycling, energy efficiency), and safety information and compliance data, (RoHS and WEEE). For more information, see https://www.arubanetworks.com/company/about-us/environmental-citizenship/.

Documentation Feedback

Aruba is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback-switching@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.