

Extreme SLX-OS SDN Configuration Guide, 17r.2.01

Supporting the ExtremeRouting SLX 9850 and ExtremeSwitching SLX 9540 devices

© 2018, Extreme Networks, Inc. All Rights Reserved.

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries. All other names are the property of their respective owners. For additional information on Extreme Networks Trademarks please see www.extremenetworks.com/company/legal/trademarks. Specifications and product availability are subject to change without notice.

© 2017, Brocade Communications Systems, Inc. All Rights Reserved.

Brocade, the B-wing symbol, and MyBrocade are registered trademarks of Brocade Communications Systems, Inc., in the United States and in other countries. Other brands, product names, or service names mentioned of Brocade Communications Systems, Inc. are listed at www.brocade.com/en/legal/brocade-Legal-intellectual-property/brocade-legal-trademarks.html. Other marks may belong to third parties.

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

The authors and Brocade Communications Systems, Inc. assume no liability or responsibility to any person or entity with respect to the accuracy of this document or any loss, cost, liability, or damages arising from the information contained herein or the computer programs that accompany it.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Brocade products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit <http://www.brocade.com/support/oscd>.

Contents

Preface	5
Document conventions.....	5
Notes, cautions, and warnings.....	5
Text formatting conventions.....	5
Command syntax conventions.....	6
Extreme resources.....	6
Document feedback.....	6
Contacting Extreme Technical Support.....	7
About This Document	9
Supported hardware and software.....	9
Interface module capabilities.....	9
OpenFlow 1.3	11
OpenFlow 1.3 protocol overview.....	11
Flow table entries.....	12
OpenFlow instructions.....	13
OpenFlow actions.....	14
Supported OpenFlow messages.....	15
Multiple controllers.....	16
Normal action support on hybrid OpenFlow ports.....	17
Rate limiting capabilities on OpenFlow enabled ports.....	17
OpenFlow TCAM profiles.....	18
OpenFlow configuration.....	20
Enabling OpenFlow on devices.....	20
Enabling OpenFlow on a specified interface.....	20
Configuring the OpenFlow Controller.....	20
OpenFlow hybrid port mode.....	21
OpenFlow hybrid port mode operation	22
Configuring OpenFlow hybrid port mode	23
Capabilities and prerequisites for hybrid ports.....	24
Enabling OpenFlow hybrid port mode.....	24
Configuring OpenFlow hybrid port on an interface.....	25
Adding or deleting protected VLANs.....	25
ACL and PBR support for OpenFlow hybrid ports.....	25
ACL and PBR action and behavior on the OpenFlow interface.....	25
OpenFlow logical port support on LAG ports.....	26
OpenFlow LAG port limitations.....	27
QinQ.....	27
QinQ action.....	28
Group table.....	28
Group messages.....	29
Displaying groups for the OpenFlow ports.....	30
Ingress replication for OpenFlow groups.....	30
Metering.....	31
Limitations.....	32
Meter messages.....	32

Show commands for OpenFlow.....	33
Clear commands for OpenFlow.....	33

Preface

- Document conventions..... 5
- Extreme resources..... 6
- Document feedback..... 6
- Contacting Extreme Technical Support..... 7

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Extreme technical documentation.

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

NOTE

A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION

An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.



CAUTION

A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.



DANGER

A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used to highlight specific words or phrases.

Format	Description
bold text	Identifies command names. Identifies keywords and operands. Identifies the names of GUI elements.
<i>italic text</i>	Identifies text to enter in the GUI. Identifies emphasis. Identifies variables.
Courier font	Identifies document titles. Identifies CLI output.

Format	Description
	Identifies command syntax examples.

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
<i>italic text</i>	Identifies a variable.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member[member...]</i> .
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Extreme resources

Visit the Extreme website to locate related documentation for your product and additional Extreme resources.

White papers, data sheets, and the most recent versions of Extreme software and hardware manuals are available at www.extremenetworks.com. Product documentation for all supported releases is available to registered users at www.extremenetworks.com/support/documentation.

Document feedback

Quality is our first concern at Extreme, and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you.

You can provide feedback in two ways:

- Use our short online feedback form at <http://www.extremenetworks.com/documentation-feedback-pdf/>
- Email us at internalinfodev@extremenetworks.com

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Contacting Extreme Technical Support

As an Extreme customer, you can contact Extreme Technical Support using one of the following methods: 24x7 online or by telephone. OEM customers should contact their OEM/solution provider.

If you require assistance, contact Extreme Networks using one of the following methods:

- [GTAC \(Global Technical Assistance Center\)](#) for immediate support
 - Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact.
 - Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- [GTAC Knowledge](#) - Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- [The Hub](#) - A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- [Support Portal](#) - Manage cases, downloads, service contracts, product licensing, and training and certifications.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

About This Document

- Supported hardware and software.....9

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some devices but not to others, this guide identifies exactly which devices are supported and which are not.

Although many different software and hardware configurations are tested and supported by this release, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by this release:

- ExtremeRouting SLX 9850-4 router
- ExtremeRouting SLX 9850-8 router
- ExtremeSwitching SLX 9540 switch

To obtain information about other releases, refer to the documentation specific to that release.

Interface module capabilities

The following table lists the supported capabilities for the following SLX 9850 interface modules:

- BR-SLX9850-10Gx72S-M
- BR-SLX9850-100Gx36CQ-M
- BR-SLX9850-10Gx72S-D
- BR-SLX9850-100Gx36CQ-D
- BR-SLX9850-100Gx12CQ-M

TABLE 1 SLX 9850 interface modules capabilities

Capability	Modular interface module
MPLS	Yes
Packet buffer memory per interface module	12GB (BR-SLX9850-10Gx72S-M) 36GB (BR-SLX9850-100Gx36CQ-M) 8GB (BR-SLX9850-10Gx72S-D) 24GB (BR-SLX9850-100Gx36CQ-D) 8GB (BR-SLX9850-100Gx12CQ-M)

OpenFlow 1.3

- OpenFlow 1.3 protocol overview..... 11
- OpenFlow configuration..... 20
- OpenFlow hybrid port mode..... 21
- ACL and PBR support for OpenFlow hybrid ports..... 25
- OpenFlow logical port support on LAG ports..... 26
- QinQ..... 27
- Group table..... 28
- Metering..... 31
- Show commands for OpenFlow..... 33
- Clear commands for OpenFlow..... 33


OpenFlow 1.3 protocol overview

An OpenFlow-enabled switch supports an OpenFlow Client (control plane software), which communicates with an OpenFlow Controller using the OpenFlow protocol. The OpenFlow Controller runs on a server or a server cluster. OpenFlow-enabled routers support the abstraction of a flow table, which is manipulated by the OpenFlow Controller. The flow table contains flow entries. Each flow entry represents a flow (that is, packets with a given MAC address, VLAN tag, IP address, or TCP/UDP port, and so on). The flow table is sorted by flow priority, which is defined by the OpenFlow Controller. The highest priority flows are at the top of the flow table.

Incoming packets on an OpenFlow-enabled port are matched (in order of priority) against the flow entries defined for that port by the OpenFlow Controller. If the packet matches a given flow entry, the flow-matching process stops, and the set of actions defined for that flow entry are performed. Packets that do not match any flow entry are dropped by default. The implementation of OpenFlow supports an option to send such packets to the OpenFlow Controller.

An OpenFlow switch maintains one flow table, which is used for packet processing. The switch performs the actions listed in the table entry corresponding to the matched flow. The OpenFlow Controller manages the OpenFlow switch using the OpenFlow protocol. The OpenFlow Controller can add, delete, or modify flows and get statistics for ports and flows and other information using the OpenFlow protocol.

FIGURE 1 OpenFlow 1.3 architecture



OpenFlow 1.3 defines three types of tables:

- Flow table
- Group table
- Meter table

Flow table entries

Only one flow table is supported on an OpenFlow port. Each flow table entry contains the fields described in the following table.

TABLE 2 Flow table entries

Field	Description
Match fields	The match fields consist of ingress ports, packet header fields
Priority	Matching precedence of the entry
Counters	Statistics for matching packets
Instructions	Action set or pipeline processing
Cookie	Opaque data sent by the OpenFlow Controller

The following match fields are supported:

- All Layer 2 header fields
- All Layer 3 header fields
- All Layer 2 and IPv4 header fields jointly

TABLE 3 OpenFlow match fields supported is SLX-OS

Match field	Prerequisite	Description
OXM_OF_IN_PORT	None	Ingress port. Numerical representation of incoming port, starting at 1. This may be a physical or switch-defined logical port.
OXM_OF_ETH_SRC	None	Ethernet source MAC address
OXM_OF_ETH_DST	None	Ethernet destination MAC address
OXM_OF_ETH_TYPE	None	Ethernet type of the packet
OXM_OF_VLAN_VID	None	VLAN-ID from 802.1Q header
OXM_OF_VLAN_PCP	VLAN_VID = None	VLAN priority (VLAN-PCP) from 802.1Q header
OXM_OF_IPV4_SRC	Ether type = 0x0800	IPv4 source address
OXM_OF_IPV4_DST	Ether type = 0x0800	IPv4 destination address
OXM_OF_IPV6_SRC	Ether type = 0x86dd	IPv6 source address
OXM_OF_IPV6_DST	Ether type = 0x86dd	IPv6 destination address
OXM_OF_IP_PROTO	Ether type = 0x0800	IPv4 protocol number
OXM_OF_IP_DSCP	Ether type = 0x0800	IPv4 DSCP (IPv4 ToS bits)
OXM_OF_TCP_SRC	IP PROTO = 6	TCP source port
OXM_OF_TCP_DST	IP PROTO = 6	TCP destination port
OXM_OF_UDP_SRC	IP PROTO = 17	UDP source port
OXM_OF_UDP_DST	IP PROTO = 17	UDP destination port

OpenFlow instructions

Each flow entry has a set of instructions that are executed when the packet matches the entry.

The instruction set associated with each flow entry can have a maximum of one instruction of each type. The following table shows the actions supported on different devices.

NOTE

Only one flow table is supported on all platforms for OpenFlow.

TABLE 4 Actions for flow table instruction

Actions	Description	SLX 9850-4	SLX 9850-8	SLX 9450
Write-Action actions (Req)	Adds or overwrites specified actions to the action set.	Yes	Yes	Yes
Apply-Action actions	Applies the specified actions immediately.	Yes	Yes	Yes
Clear-Action actions	Clears all the actions in the action set.	No	No	No
Meter <i>meter-id</i>	Directs the packet to the specified meter.	Yes	Yes	Yes
Goto -Table <i>next-table-id</i> (Req)	Indicates the next table in pipeline processing.	No ¹	No ¹	No ¹
Write-Metadata metadata/mask	Writes the metadata field from the mask.	No ¹	No ¹	No ¹

¹ Not required with one flow table.

OpenFlow actions

Each flow has a set of instructions that are executed when the packet matches the flow as per OpenFlow 1.3 specifications. Each flow can have a maximum of one instruction of each type.

A switch can reject a flow entry if the switch is unable to execute the instructions associated with the flow entry. In this case, the switch returns an unsupported flow error. Flow table may not support every match, every instruction, or every action.

TABLE 5 Supported actions for devices

Actions	Description	SLX 9850-4	SLX 9850-8	SLX 9450
Output (Req)	Forwards the packet to a specified OpenFlow port. If out-port is Controller, then the packet will be sent as packet-in message.	Yes	Yes	Yes
Drop (Req)	No explicit drop action. Packet with empty action set should be dropped.	Yes	Yes	Yes
Group	Processes the packet through the specified group.	Yes	Yes	Yes
Set field	Modifies the values of the packet header based on the field type.	Yes (VLAN_VID, VLAN-PCP)	Yes (VLAN_VID, VLAN-PCP)	Yes (VLAN_VID, VLAN-PCP)
Push-Tag/ Pop-Tag	Adds and removes tag (newly inserted tags are always the outermost tags).	Yes	Yes	Yes
Set-Queue	Set the queue ID for the packet.	Yes	Yes	Yes
Change TTL	Modify the TTL value.	No	No	No

Set field is used to modify the packet content. Devices support VLAN ID and VLAN-PCP modifications only.

- OXM_OF_VLAN_VID modifies the outermost VLAN ID, when tag-type is 0x8100. In case of untagged packet, a new VLAN header is created.
- OXM_OF_VLAN_PCP modifies the PCP value in the outermost VLAN header. This action is ignored for untagged packet.

OpenFlow actions output

The OpenFlow flow supports the following Openflow ports: physical, logical, and reserved.

TABLE 6 Supported actions output on device

Type	On device
Hardware interface of the switch	Yes
LAG	No
Tunnel	No
All	No
Controller	Yes
Table	No
In_port	No
Local	No
Normal	No
Flood	No

Flow action supports more than one OpenFlow physical port in the action-list.

Limitations

- Combination of physical and reserved ports are not supported by flow action.
- Controller port is rate limited to 4 kbps from hardware to CPU.

OpenFlow actions push and pop VLAN

The following actions and limitations are supported for push and pop VLAN.

Push VLAN adds a VLAN tag to the existing packet. The flow configuration supports only 0x8100 tag-type, push VLAN with any other tag-types are not supported. Push VLAN is only valid on physical port.

Pop VLAN action removes the outermost VLAN header. If the action-list has both pop VLAN and set field greater than VLAN_VID or push VLAN, then the flow is rejected. Pop VLAN is used only with physical port out port. Multiple pop VLAN action in the same flow is not supported.

Supported OpenFlow messages

The following OpenFlow messages are supported on the devices.

TABLE 7 OpenFlow messages

Message type	SLX 9850-4	SLX 9850-8	SLX 9450
OFPT_HELLO	Yes	Yes	Yes
OFPT_ERROR	Yes	Yes	Yes
OFPT_ECHO_REQUEST	Yes	Yes	Yes
OFPT_ECHO_REPLY	Yes	Yes	Yes
OFPT_EXPERIMENTER	No	No	No
OFPT_FEATURES_REQUEST	Yes	Yes	Yes
OFPT_FEATURES_REPLY	Yes	Yes	Yes
OFPT_GET_CONFIG_REQUEST	No	No	No
OFPT_GET_CONFIG_REPLY	No	No	No
OFPT_SET_CONFIG	No	No	No
OFPT_PACKET_IN	Yes	Yes	Yes
OFPT_FLOW_REMOVED	Yes	Yes	Yes
OFPT_PORT_STATUS	Yes	Yes	Yes
OFPT_PACKET_OUT	Yes	Yes	Yes
OFPT_FLOW_MOD	Yes	Yes	Yes
OFPT_GROUP_MOD	Yes	Yes	Yes
OFPT_PORT_MOD	No	No	No
OFPT_TABLE_MOD	No	No	No
OFPT_MULTIPART_REQUEST	Yes	Yes	Yes
OFPT_MULTIPART_REPLY	Yes	Yes	Yes
OFPT_BARRIER_REQUEST	Yes	Yes	Yes
OFPT_BARRIER_REPLY	Yes	Yes	Yes
OFPT_QUEUE_GET_CONFIG_REQUEST	No	No	No
OFPT_QUEUE_GET_CONFIG_REPLY	No	No	No
OFPT_ROLE_REQUEST	Yes	Yes	Yes

TABLE 7 OpenFlow messages (continued)

Message type	SLX 9850-4	SLX 9850-8	SLX 9450
OFPT_ROLE_REPLY	Yes	Yes	Yes
OFPT_GET_ASYNC_REQUEST	Yes	Yes	Yes
OFPT_GET_ASYNC_REPLY	Yes	Yes	Yes
OFPT_SET_ASYNC	Yes	Yes	Yes
OFPT_METER_MOD	Yes	Yes	Yes

Multiple controllers

An OpenFlow switch may be connected to multiple controllers for reliability, allowing the switch to continue to operate in OpenFlow mode if a controller or controller connection fails. The controllers coordinate the management of the switch amongst themselves to help synchronize controller handoffs.

Each controller can have one of the following roles:

- Equal (OFPCR_ROLE_EQUAL): The controller has full access to the switch. It receives all the asynchronous messages from the switch and sends commands to modify the state of the switch (add or delete flows).
- Slave (OFPCR_ROLE_SLAVE): The controller has a read-only access to the switch. It does not receive the asynchronous messages (apart from port status). It does not execute commands that modify the state of the switch: **packet-out**, **flow-mod**, **group-mod**, or **port-mod**. The switch must reply with an OFPT_ERROR message, if it receives one of those commands from a Slave controller. Other controller-to-switch messages are processed normally.
- Master (OFPCR_ROLE_MASTER): The controller has full access to the switch as in the Equal role. When the controller changes its role to Master, the switch changes the other controller in the Master role to have the Slave role. The role change does not affect controllers with the Equal role.

Regardless of the intended use of multiple controller connections, the device allows all the controller connections to concurrently manage the flow table. That is, flow entries in the flow table are not identified as belonging to any specific controller connection. In an active-standby controller deployment, controllers themselves must coordinate their actions and active-standby states. The device responds to all connected controllers without distinction.

The device supports the active controller connection (also called mode). The device initiates the TCP connection to a given OpenFlow Controller address. A counter Generation ID is assigned by the controller each time the mastership view changes. For role changing to Master or Slave, the switch validates Generation ID to check for stale messages.

Asynchronous configuration

Asynchronous messages may need to be sent to multiple controllers. An asynchronous message is duplicated for each eligible OpenFlow channel, and each message is sent when the respective controller connection allows it.

A controller can also control which types of switch asynchronous messages are sent over its OpenFlow channel. This is done using an asynchronous configuration message that has the filter setting for all the messages.

Each role for every controller may have its own set of asynchronous message setting. A controller in the Master role can selectively disable notifications, and a controller in the Slave role can enable notifications it wants to monitor.

Each controller configuration block for active connection maintains its own asynchronous configuration setting for every role. The default initial configuration is shown in the following table.

TABLE 8 Action for asynchronous configuration

Messages	Bit field	Master or Equal role	Slave role
Packet-in reasons	No_match	Enable	Disable
	Action	Enable	Disable
	Invalid_TTL	Disable	Disable
Port status reasons	Add	Enable	Enable
	Delete	Enable	Enable
	Modify	Enable	Enable
Flow removed reasons	Idle_timeout	Disable	Disable
	Hard_timeout	Disable	Disable
	Delete	Enable	Disable
	Group_delete	Enable	Disable

NOTE

The asynchronous messages, Action, Invalid_TTL, Idle_timeout, and Hard_timeout are not supported by devices. Controllers can set these bits in the filter setting and the device can accept the bits, but the messages are not sent out by the device.

Normal action support on hybrid OpenFlow ports

Normal action represents the traditional routing with traffic on configured protected and unprotected VLAN with the OpenFlow hybrid port. The OpenFlow 1.3 specification defines the following reserved ports for specifying generic forwarding action. Normal is a reserved port which represents the traditional non-OpenFlow pipeline of the switch. It can be used only as an output port and the switch processes the packet using the normal pipeline. When packets match an OpenFlow flow with action having the Normal port, then packet is forwarded based on the Layer 2 or Layer 3 forwarding entries. OpenFlow flows with Normal action is supported on hybrid OpenFlow ports.

Action supported with Normal action

The following actions are supported along with Normal action flows:

- Though an output port (physical port, drop, or controller) in the action cannot be specified along with the Normal action, a mirror port can be specified.
- A Normal action flow is supported only on an OpenFlow hybrid port.
- Only meter and mirror actions are supported for Normal action flows.

Rate limiting capabilities on OpenFlow enabled ports

Rate-limiting support on OpenFlow enabled ports:

- OpenFlow ports (non-hybrid port mode ports): Normal CLI configuration of port-based rate limiting is supported. Per VLAN rate limiting not supported.
- You must set the rate control for certain protocols at the global configuration level.

The following table describes the behavior, when rate-limiters are applied on OpenFlow interfaces:

TABLE 9 Behavior on OpenFlow ports with rate-limiters

Rate limiter	OpenFlow Actions	OpenFlow port type	Behavior
ACL based rate-limit	redirects to output port/group	OpenFlow Layer 2/ Layer 3/ Layer23	Traffic does rate-limit as configured in service policy, if OpenFlow rule is not applied. Traffic does not rate-limit if OpenFlow rule is matched.
ACL based rate-limit	redirects to output port/group and meter	OpenFlow Layer 2/ Layer 3/ Layer23	Traffic does rate-limit as configured in service policy, if OpenFlow rule is not applied. Traffic does rate-limit as configured via OpenFlow meter, if OpenFlow rule is matched.
Port based rate-limit	redirects to output port/group	OpenFlow Layer 2/ Layer 3/ Layer23	Traffic does rate-limit as configured port, if OpenFlow rule is not applied. Traffic does rate-limit if OpenFlow rule is matched.
Port based rate-limit	redirects to output port/group and meter	OpenFlow Layer 2/ Layer 3/ Layer23	Traffic does rate-limit as configured port, if OpenFlow rule is not applied. Traffic does rate-limit as configured via OpenFlow meter, if OpenFlow rule is matched.
BUM rate-limits + OpenFlow	redirects to output port/group	OpenFlow Layer 2/ Layer 3/ Layer23	Traffic does rate-limit in BUM, if OpenFlow rule is not applied. Traffic does rate-limit if OpenFlow rule is matched.
BUM rate-limits + OpenFlow	redirects to output port/group and meter	OpenFlow Layer 2/ Layer 3/ Layer23	Traffic does rate-limit in BUM configured port, if OpenFlow rule is not applied. Traffic does rate-limit as configured via OpenFlow meter, if OpenFlow rule is matched.

NOTE

There is no restriction on configuring port based BUM rate-limiting on a pure OpenFlow interface.

OpenFlow TCAM profiles

The following table illustrates the flow matches in the different OpenFlow TCAM profiles. To support Network Packet Broker (NPB), TCAM profile NPB-Profile is added. During scaling in NBP-Profile, the maximum number of flows per line card is 10,000.

TABLE 10 Flows and TCAM profiles

OpenFlow features		Default profile	Optimized profile 1	Optimized profile 2	Optimized profile 3	NPB-Profile
Hybrid mode		Yes	Yes	Yes	Yes	No
Match	In port	Yes	Yes	Yes	Yes	Yes
	VLAN ID	Yes	Yes	Yes	Yes	Yes
	VLAN-PCP	Yes	Yes	Yes	Yes	Yes
	Source MAC address	No	Yes	Yes	Yes ²	Yes ²
	Destination MAC address	No	Yes	Yes	Yes ²	Yes ²

² For IPv4 packets only

TABLE 10 Flows and TCAM profiles (continued)

OpenFlow features		Default profile	Optimized profile 1	Optimized profile 2	Optimized profile 3	NPB-Profile
	Ether type = 0x0800	Yes	Yes	Yes	Yes	Yes
	Ether type except 0x0800 and 0x86dd	No	Yes	Yes	Yes	Yes
	Source IPv4 address	Yes	Yes	Yes	Yes ²	Yes ²
	Destination IPv4 address	Yes	Yes	Yes	Yes ²	Yes ²
	DSCP	Yes	Yes	Yes	Yes	Yes
	Layer 4 source port	Yes	Yes	Yes	Yes	Yes
	Layer 4 destination port	Yes	Yes	Yes	Yes	Yes
	Ether type = 0x86dd	Yes	No	No	Yes	Yes
	Source IPv6 address	Yes	No	No	Yes	Yes
	Destination IPv6 address	No	No	No	Yes	Yes
	IPv6 TC	Yes	No	No	Yes	Yes
	IPv6 Next header	Yes	No	No	Yes	Yes
Action	Single out port	Yes	Yes	Yes	Yes	Yes
	Multiple out ports	Yes	Yes	Yes	Yes	Yes
	Send to controller	Yes	Yes	Yes	Yes	Yes
	Drop	Yes	Yes	Yes	Yes	Yes
	Normal action	Yes	Yes	Yes	Yes	Yes
	Normal action and mirror action	No	Yes	Yes	Yes	Yes
	VLAN modification (set/push/pop) with single out port	No	Yes	No	No	No
	VLAN modification (set/push/pop) with multiple out ports	Yes	Yes	Yes	Yes	Yes
	VLAN-PCP modification with single and multiple out ports	No	Yes ³	No	No	No
	Flow statistics	Yes ⁴	Yes ⁴	Yes ⁴	Yes ⁴	Yes ⁴
	Meter	No	No	Yes	Yes	No
	Enqueue	No	No	Yes	Yes	No
	Group (all group types)	Yes	Yes	Yes	Yes	Yes

² For IPv4 packets only³ For VLAN-PCP modification on multiple ports, all output ports must have the same VLAN-PCP.⁴ Collection of statistics depends on the availability of the counter engines.

OpenFlow configuration

You can enable OpenFlow on an interface with Layer23 flows in order to support Layer 2 and Layer 3 flows on that interface. Layer23 flows support the OpenFlow hybrid port mode also. Configured with Layer23, the controller can configure flows with Layer 2 and Layer 3 parameters together. A flow can contain the following fields: Ingress port, Destination MAC address, Source MAC address, Ether type, VLAN ID, P-bits, Source IP address, Destination IP address, IP protocol, and IP DSCP.

By default, OpenFlow is disabled on devices. You must first enable OpenFlow on the device before you can configure the parameters on the device.

Enabling OpenFlow on devices

After you enable OpenFlow on the device, you can enable OpenFlow on specific interfaces and configure additional OpenFlow parameters.

To enable OpenFlow, enter the following command:

```
device(config)# openflow enable ofv13
```

The **ofv13** keyword specifies the OpenFlow protocol version supported.

Use the **no openflow enable ofv13** command to disable OpenFlow on the device.

NOTE

You must disable OpenFlow on all interfaces individually before you can disable OpenFlow globally on the device.

Enabling OpenFlow on a specified interface

After you have enabled OpenFlow on the device, you can enable OpenFlow on specific interfaces.

NOTE

You can enable OpenFlow on an interface only after you have enabled OpenFlow globally on the device. In addition, you must use individual CLI commands to enable OpenFlow on each interface. You cannot specify a range of ports when enabling OpenFlow.

To enable OpenFlow on different interfaces, enter the following commands:

```
device(config)#interface Ethernet 3/1
device(config-if-eth-3/1)#openflow enable layer2
device(config)#interface Ethernet 3/2
device(config-if-eth-3/2)#openflow enable layer3
device(config)#interface Ethernet 3/3
device(config-if-eth-3/3)#openflow enable layer23
```

If you enable Layer 3 matching mode on the specified interface, only Layer 3 matching fields are supported on that interface.

Use the **no openflow enable** command to disable OpenFlow on the interface.

Configuring the OpenFlow Controller

To configure the OpenFlow controller, use the following steps.

By default no controller connection is present. The device supports up to three controller connections.

Set the IP address of the Controller.

1. From privileged EXEC mode, enter global configuration mode.

```
device# configure terminal
```

2. Enter **openflow enable** command to globally enable OpenFlow.

```
device(config)# openflow enable ofv13
```

The ofv13 option represents OpenFlow version 1.3.0.

3. Enter **openflow controller** command to name the controller and assign an IP address.

```
device(config)# openflow controller A1 ip-address 10.25.128.185 no-ssl
```

The OpenFlow controller is created to use TCP connectivity without SSL. By default, a controller role is Equal. The controller may change its role using OFPT_ROLL_REQUEST message.

4. Verify the OpenFlow controller configuration.

```
device(config)# show openflow controller
Controller      Mode      TCP/SSLIP-address      Port      Status      Role
A1              active    TCP 10.25.101.65       6653      CLOSE      Equal
```

Use the **no openflow controller A2** command to remove an OpenFlow controller.

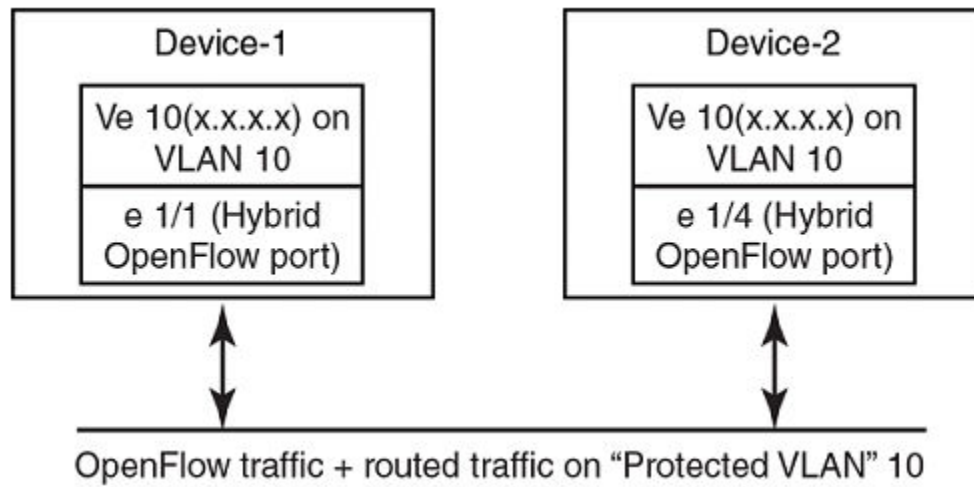
```
device(config)# do show run | inc openflow\ controller
openflow controller A1 ip-address 10.25.128.185 no-ssl port 9000
openflow controller A2 ip-address 10.25.128.185 no-ssl
device(config)# no openflow controller A2
device(config)# do show run | inc openflow\ controller
openflow controller A1 ip-address 10.25.128.185 no-ssl port 9000
```

OpenFlow hybrid port mode

OpenFlow hybrid-enabled ports support both OpenFlow traffic forwarding and normal routing traffic forwarding. OpenFlow hybrid-enabled ports support "protected VLANs" and "unprotected VLANs". Protected VLANs are not subject to defined OpenFlow flows on the OpenFlow hybrid-enabled ports. OpenFlow flows on a hybrid-enabled port will not match any traffic on protected VLANs. Unprotected VLANs are subject to defined OpenFlow flows on the OpenFlow hybrid-enabled port. OpenFlow flows on a hybrid-enabled port are allowed to match on the traffic of unprotected VLANs.

The following figure shows a topology in which port 1/1 on Device-1 and port 1/4 on Device-2 are hybrid-enabled OpenFlow ports with VLAN 10 as a configured protected VLAN. By configuring a virtual Ethernet (VE) interface on a protected VLAN 10 and assigning an address to route the traffic of the nodes, you are able to send protected VLAN traffic between the nodes and route the traffic as per the VE interface. Traffic flowing on other VEs created on top of other VLANs (the unprotected VLANs) is treated as unprotected VLAN traffic and is subject to OpenFlow rules lookup. OpenFlow traffic can be forwarded through this port.

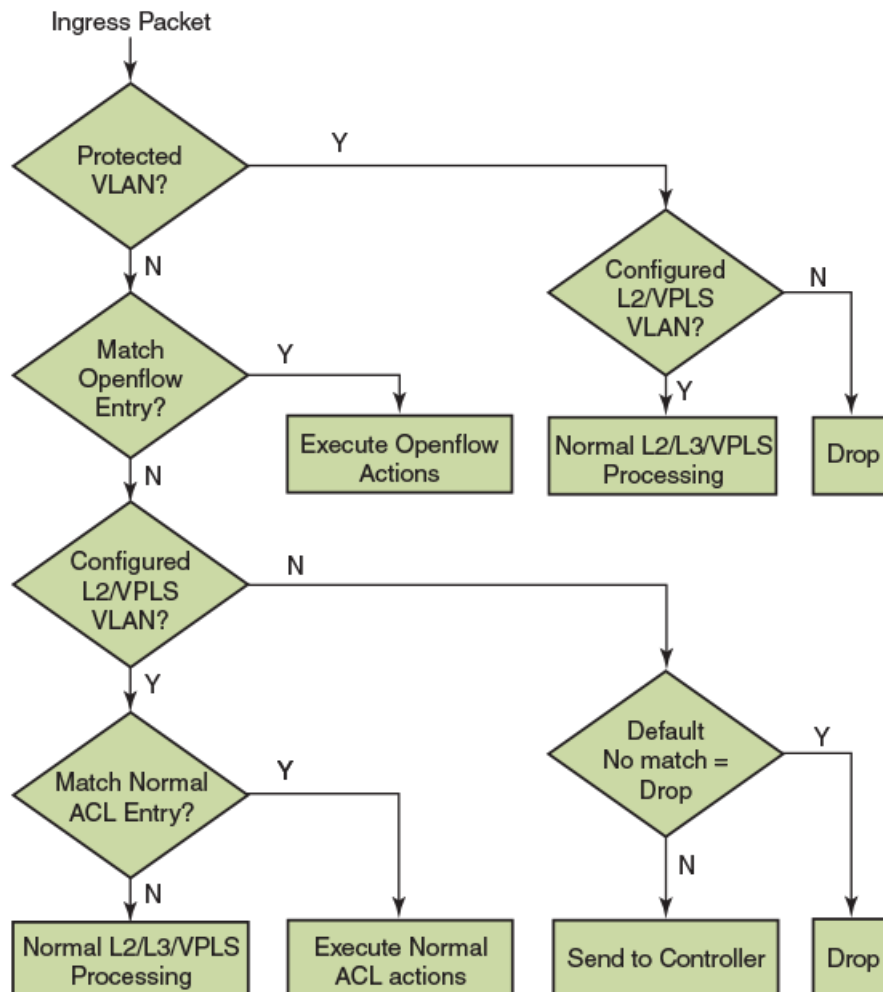
FIGURE 2 OpenFlow hybrid port mode topology



OpenFlow hybrid port mode operation

Ingress traffic on VLAN 10 on hybrid port 1/1 is processed for IPv4 and IPv6 unicast routing. Traffic on other VLANs is processed against OpenFlow flows on port 1/1 and switched accordingly. A preconfigured number of protected VLANs can be supported for normal routing. The following figure shows the flow for an ingress packet processed to reach the controller.

FIGURE 3 Packet flow diagram for OpenFlow hybrid port support



Configuring OpenFlow hybrid port mode

1. Enable OpenFlow at the global configuration level.
2. Configure the OpenFlow controller.

NOTE

Before being configured as an OpenFlow hybrid port, the port must be configured as a switch port mode.

3. Configure the maximum number of OpenFlow unprotected VLAN entries. (The default is 0.)
4. Configure protected VLANs on the port. A maximum of 40 protected VLANs can be configured on an OpenFlow port.
5. Enable OpenFlow hybrid port mode on the desired interfaces.
6. Configure a VE for the interface by specifying the protected VLANs and adding routing entries.

Capabilities and prerequisites for hybrid ports

- The system supports a maximum of 40 protected VLANs per port.
- The system supports up to 2000 protected VLANs.
- OpenFlow v1.3 protocol is supported.
- Layer 2 protocols and VPLS forwarding are supported on ports in hybrid-enabled ports.
- IPv4 and IPv6 unicast routing are supported on OpenFlow protected and unprotected VLANs.
- Packets tagged with a protected VLAN ID are forwarded by IPv4 and IPv6 unicast routing, if IPv4 or IPv6 routing is configured on the VLAN. If IPv4 or IPv6 routing is not configured on the VLAN (an unconfigured VLAN), packets are dropped.
- Packets tagged with an unprotected VLAN ID are subject first to OpenFlow flows. If there is a match on an OpenFlow flow, the packets are forwarded according to the flow actions. No further IPv4 or IPv6 routing is supported for packets that are forwarded by OpenFlow flows. If there is no match on any OpenFlow flow, the packets are forwarded by IPv4 or IPv6 unicast routing, if IPv4 or IPv6 routing is configured on the VLAN. If IPv4 or IPv6 routing is not configured on the VLAN, the packets are either dropped or sent to the controller, per the OpenFlow configuration.
- The BGP, OSPF, and IS-IS protocols are supported on protected VLANs.
- Layer 2 protocols such as MSTP, SSTP, or PVSTP are supported.
- Port-level sFlow sampling is supported.
- If a flow from the controller is added for a protected VLAN, or the VLAN is configured as protected on the port, the flow is rejected.
- If a flow belongs to a VLAN, then that VLAN is not allowed to be configured as a protected VLAN.
- IPv4 and IPv6 unicast routing with VRF Lite are supported.

Limitations of OpenFlow hybrid ports

- The OpenFlow hybrid port cannot be an untagged member of any VLAN or VPLS VLAN except the default VLAN.
- OpenFlow hybrid port mode cannot be enabled on a port configured with switchport mode access.
- Protected VLAN traffic with no matching MAC or IP route entries is dropped.
- Policy Based Routing (PBR) is supported.
- The OpenFlow hybrid port cannot be configured as a pure Layer 3 port.
- Link-level Layer 2 protocols (UDLD, 802.1x, loop detection, link-OAM, LLDP, LACP) are not supported on OpenFlow interfaces.
- A LAG interface cannot be an OpenFlow hybrid port, and an OpenFlow hybrid port cannot be added to a LAG interface.
- An inbound normal ACL configuration is supported on the port in OpenFlow hybrid port mode.
- Wild card VLAN flows existing in the system do not affect the individual VLANs being configured as protected.

Enabling OpenFlow hybrid port mode

Use the **openflow enable hybrid-mode** command to enable OpenFlow hybrid port mode on the port for different interfaces. The **no openflow enable hybrid-mode** command disables the OpenFlow hybrid port mode on the port and the port becomes a normal port. For example, to enable OpenFlow hybrid mode on Layer 2, use the following command.

```
device(config-if-eth-3/1)# openflow enable layer2 hybrid-mode
```

To enable OpenFlow hybrid mode on Layer 23, use the following command.

```
device(config-if-eth-3/1)# openflow enable layer23 hybrid-mode
```


Configuring OpenFlow hybrid port on an interface

1. Enable OpenFlow at the global configuration level.
2. Configure OpenFlow controller configurations.
3. Configure the switchport on the interface.
4. Configure the switchport trunk mode on the interface.

The following example configures an OpenFlow hybrid port on an interface.

```
device(config)# interface Ethernet 3/1
device(config-if-eth-3/1)# switchport
device(config-if-eth-3/1)# switchport mode trunk
device(config-if-eth-3/1)# openflow enable layer23 hybrid-mode
```

Adding or deleting protected VLANs

Use the **openflow protected-vlans** command to add or delete protected VLANs on an OpenFlow hybrid port mode interface. The **no openflow protected-vlans** command deletes the configured protected VLANs from the hybrid-enabled port.

```
device(config)# interface Ethernet 3/1
device(config-if-eth-3/1)# openflow protected-vlans add 101-105, 110
device(config-if-eth-3/1)# openflow protected-vlans remove 102-104, 110
```

NOTE

1. Both range or individual VLAN are supported for addition and removal.
2. The **openflow protected-vlans** can be preconfigured, even when hybrid mode is not enabled on an interface though, but global **openflow enable** is the requirement on the switch.
3. The port does not be in trunk mode before pre-provisioning a protected VLAN.

ACL and PBR support for OpenFlow hybrid ports

Access Control Lists (ACLs) and Policy Based Routing (PBR) are supported on OpenFlow hybrid interfaces along with OpenFlow flows.

The configured protected VLAN traffic follows the OpenFlow rules as well as ACL and PBR rules. Unconfigured protected VLAN traffic does not follow the ACL or PBR rules and the traffic is dropped from the port.

The configured unprotected VLAN traffic follows OpenFlow rules. If there are no rules to forward the unprotected VLAN traffic, then it follows ACL and PBR rules. The unconfigured unprotected VLAN traffic on an OpenFlow hybrid port is dropped or sent to the OpenFlow controller.

The configured unprotected VLAN traffic follows the Normal action flow and OpenFlow rules. If there are no rules to forward the unprotected VLAN traffic, then it follows ACL and PBR rules.

ACL and PBR action and behavior on the OpenFlow interface

The following table describes the behavior, when ACL or PBR and OpenFlow flows with different actions are applied on the same interface for Layer 2, Layer 3, and Layer23 ports.

TABLE 11 ACL and PBR action on OpenFlow interface

ACL or PBR	Type	OpenFlow action	Behavior
ACL	MAC ACL	Redirect to output port or group	Packets matching the OpenFlow rules redirect to the port specified in the OpenFlow rule.

TABLE 11 ACL and PBR action on OpenFlow interface (continued)

ACL or PBR	Type	OpenFlow action	Behavior
ACL	MAC ACL	Send to controller	Packets matching the OpenFlow rules are sent to the controller.
ACL	IP ACL	Redirect to output port or group	Packets matching the OpenFlow rules redirect to the port specified in the OpenFlow rule.
ACL	IP ACL	Send to controller	Packets matching the OpenFlow rules are sent to the controller.
ACL	Receive ACL	Redirect to output port	Packets matching the OpenFlow rules redirect to the port specified in the OpenFlow rule.
ACL	Receive ACL	Send to controller	Packets matching the OpenFlow rules are sent to the controller.
PBR	Redirect to host	Redirect to output port	Packets matching the OpenFlow rules redirect to the port specified in the OpenFlow rule.
PBR	Redirect to host	Drop	Packets matching the OpenFlow rules are dropped.
PBR	Redirect to host	Send to controller	Packets matching the OpenFlow rules are sent to the controller.

NOTE

To bind MAC ACLs or IP ACLs on an interface, the interface must be configured as a switch port. On interfaces with the switch port configured, OpenFlow can be enabled only in hybrid mode. ACL and OpenFlow coexist only when the interface is configured in OpenFlow hybrid mode.

NOTE

OpenFlow cannot be enabled in the IP route map configured on the interface. For PBR and OpenFlow to coexist on the same interface, the route map is configured on a VE. If the packet uses the OpenFlow Normal action flow, then the packet does not take the route map to go forward. The packet goes forward based on the Layer 2 or Layer 3 forwarding table.

OpenFlow logical port support on LAG ports

This feature enables OpenFlow hybrid flows on LAG ports such that non OpenFlow traffic forwarding is also supported on a LAG port. The controller can configure OpenFlow and traditional flows on individual LAG ports.

Each LAG port has a unique OpenFlow port ID similar to a physical port.

In the Network Packet Broker (NPB) deployment, Transparent VLAN Flooding (TVF) is used to flood or to load balance the packet, which matches a route-map. TVF is supported in OpenFlow through groups.

TVF with physical ports as VLAN members

You must install OpenFlow group All to make VLAN members become the ports in the OpenFlow buckets.

TVF with only one LAG port as a member of the VLAN

You must install OpenFlow group Select to make the ports in the LAG become ports in the OpenFlow buckets using TVF with load-balancing.

TVF with physical ports and LAG ports as members

You must install OpenFlow group All to make the VLAN members become the ports in the OpenFlow buckets using TVF with load-balancing and flooding.

OpenFlow port status messages are sent to the controller in the following case:

- OpenFlow is enabled on LAG (Add).
- OpenFlow is disabled on LAG (Delete).
- LAG operational status change (Modify).
- A port is added or deleted in the LAG.

OpenFlow LAG port limitations

- If OpenFlow is enabled on a physical port then the port cannot be added as a LAG port member. OpenFlow must be disabled on the physical port before adding the port as a LAG member port.
- OpenFlow cannot be enabled or disabled on the individual member ports of a LAG port.
- LAG ports cannot be used as in-port in OpenFlow flows.
- LAG ports can be used only in OpenFlow group All. LAG ports are not supported in Select, Fast failover, and Indirect groups.
- Packet-in or packet-out are not supported on LAG ports.
- OpenFlow is supported only on static LAG port.
- In case a logical port is used as a watch_port in any OpenFlow group, then no validation checks are performed and the group gets programmed successfully. This limitation is not supported on SLX 9450 devices.

QinQ

You can push, pop or set VLAN tags in the outgoing packets of an OpenFlow flow with QinQ support. The ingress packet can be untagged or tagged. You can use QinQ to transport multiple customer segments or VLANs across Layer 2 infrastructures.

A OpenFlow flow matches on VLAN and it does one of the following.

- Push VLAN
- Pop VLAN
- Set field (VLAN ID or VLAN-PCP)

A packet is identified as tagged if the tag-type is 0x8100, packet with any other tag type is considered as untagged. For example, packet received with tag-type 0x9100 is an untagged packet.

NOTE

QinQ supports only one tag modification.

QinQ action

The following table illustrates the behavior of push or set field VLAN ID for different tags.

TABLE 12 OpenFlow QinQ actions

OpenFlow action	Input traffic	
	Untagged packet	Tagged packet
Push VLAN	New VLAN header is added with the given tag-type and VLAN ID.	New VLAN header is added with the given tag-type and VLAN ID.
Set VLAN	New VLAN header is added with default tag-type 0x8100 and set field VLAN ID.	Outermost VLAN header is modified.
Pop VLAN	Pop action is ignored.	Outermost VLAN header is removed.

OpenFlow flow is rejected, if there is mismatch in tag-type of port and the push action.

The following table illustrates the behavior of set field VLAN-PCP along with other set or push VLAN action.

TABLE 13 OpenFlow QinQ actions for set field VLAN-PCP

Incoming traffic	Redirect action	VLAN action		
		Push VLAN + Set VLAN-PCP	Set VLAN + Set VLAN-PCP	Set VLAN-PCP
Untagged packet	One output port	VLAN pushed along with new VLAN-PCP.	Both VLAN and VLAN-PCP are added as the outer VLAN header.	Not applicable
	Multiple output port	Not supported.	VLAN header is added without VLAN-PCP.	Not applicable
Tagged packet	One output port	VLAN pushed along with new VLAN-PCP.	Both VLAN and VLAN-PCP are modified.	VLAN-PCP is modified.
	Multiple output port	Not supported.	Both VLAN and VLAN-PCP are modified.	VLAN-PCP is modified.

NOTE

OpenFlow flow is rejected, when it is not supported. OpenFlow action is silently ignored, when it is not applicable.

Group table

The group table introduces the ability to add support for port group abstraction for multi-pathing. This enables OpenFlow to represent a set of ports as a single entity for forwarding packets.

The group table supports the following group types:

- All: Executes all the buckets in the group; mostly used for flooding and multicasting.
- Indirect: Executes one defined bucket in the group. The action taken by this group type is sending packets to the next hop.
- Select: Executes one bucket in the group. The action bucket is chosen by a switch-defined algorithm, such as round robin or hashing (for example, load sharing).
- Fast failover: Executes the first live bucket, used in cases such as redundancy.

A group table consists of group entries. The counters in the following table are available in a group entry.

TABLE 14 Group entry counters

Counter	Description
Group Identifier	A 32-bit unsigned integer uniquely identifying the group
Group type	Determines group semantics
Counter	Number of packets processed by a group
Action bucket	Ordered list of action buckets, where each action bucket contains a set of actions to execute and associated parameters

Group messages

The following table describes the processing of group messages.

TABLE 15 Group messages

Group message type	Entry exists	Entry does not exist	Notes
Add (OFPGC_ADD)	Deny Add. Return error message to controller	Add is processed	Subject to constraints below
Mod (OFPGC_MODIFY)	Group parameters and action buckets are updated	Deny Mod. Return error message to controller	Update or modify is implemented, as delete followed by add in the driver.
Del (OFPGC_DELETE)	Group entry is deleted. Flows which are associated with this group are also removed.	No Error. Message ignored	If a Del comes in, that has flows associated with it, then delete those flows from the system.

Error conditions and messages

This table lists the error conditions and the error opcodes sent to the controller. The error type is always OFPET_GROUP_MOD_FAILED.

TABLE 16 Group messages

Error condition	Opcode
Adding group, if group already exists	OFPGMFC_GROUP_EXISTS
When group allocation exceeds memory or system limit	OFPGMFC_OUT_OF_GROUPS
Group type is not supported	OFPGMFC_BAD_TYPE
In case of group modification or deletion, if group does not exist	OFPGMFC_UNKNOWN_GROUP
Number of buckets in a group is greater than 8 in all group types except Indirect	OFPGMFC_OUT_OF_BUCKETS
Number of output ports in a bucket is greater than 1	OFPGMFC_BAD_BUCKET
(For all group types) Not an output port action or set field VLAN_VID or OFPAT_PUSH_VLAN or OFPAT_POP_VLAN	OFPGMFC_BAD_BUCKET
Modify group for Select type to non Select type and vice versa.	OFPGMFC_EPERM
Fast Failover: The watch port is absent	OFPGMFC_BAD_WATCH
Fast Failover: The watch group is present	OFPGMFC_BAD_BUCKET

Displaying groups for the OpenFlow ports

To show the number of groups for the OpenFlow ports or for given group ID.

1. From privileged EXEC mode, enter global configuration mode.

```
device # show openflow group
```

2. To show groups for a specific group ID, enter the following command.

```
device # show openflow group 12
```

The following example shows the output of the **show openflow group** command.

```
device # show openflow group
Max number of total groups           : 512
Max number of buckets per group      : 8

TOTAL number of groups(Type:ALL) in the system : 1
TOTAL number of groups(Type:SELECT) in the system : 0
TOTAL number of groups(Type:Indirect) in the system : 0
TOTAL number of groups(Type:Fast Failover) in the system : 0

TOTAL number of groups in the system : 1

Group id 1
  Transaction id      4043243760 (0xf0ff00f0)
  Type                ALL
  Packet Count        0
  Byte Count          0
  Flow Count          0
  Number of buckets   2

  bucket #1
    Weight            1
    out port: Eth 3/25

  bucket #2
    Weight            2
    out port: Eth 3/26
-----
```

Ingress replication for OpenFlow groups

Ingress Replication reduces the packet losses replicating multicast packets at the ingress slot and not at the egress.

The group supports the following rules for ingress replication:

- The packet stream replication is one packet stream for each egress device.
- The streams go to the fabric element and are forwarded to the destination slot.
- If the egress device for a multicast stream is the same as the ingress device then the stream is reflected internally in the device. The stream does not go to the fabric element.
- If the egress device is in the same slot as the ingress device but a different device then the stream goes to the fabric element and is reflected back to the same slot but to the different egress device.
- Ingress replication is limited to 4 egress devices.

The default range of groups to use ingress replication is from 65535 to 131071. If the controller installs group All within this range the device enables Ingress replication for such groups.

Use the following command to show ingress replication:

```

device(config)# openflow ingress-replication group-id-range 100 1000
device(config)# do show openflow
Administrative Status:      ENABLED
SSL Status:                 ENABLED
Datapath-ID:               8ff88e740000
Number of Controllers:     2

Controller:                 passive
Controller Type:            OFV130
Connection Mode:           passive TCP
Listening Address:         0.0.0.0
Connection Port:           6633
VRF Name:                  mgmt-vrf
Source IP used:            NA
Connection Status:        TCP_LISTENING
Role:                       Equal
Controller:                 active
Controller Type:            OFV130
Connection Mode:           active TCP
Controller Address:        10.30.5.142
Connection Port:           7777
VRF Name:                  mgmt-vrf
Source IP used:            NA
Connection Status:        CLOSE
Role:                       Equal
Asynchronous Configuration:
                            Packet-in (no-match|action)
                            Port-status (add|delete|modify)
                            Flow-removed (idle-timeout|hard-timeout|delete|group-delete)

Match Capability:
L2:                         Port, Source MAC, Destination MAC, Ether type, Vlan, Vlan PCP
L3:                         Port, Vlan, Vlan PCP, Ethertype, Source IP(V4 and V6), Destination
IP(V4 and V6), IP Protocol, IP TOS,TCP/UDP Src Port, TCP/UDP Dst Port
L23:                        All

Openflow Enabled Ports:     Eth 3/1, Eth 3/2, Eth 3/3, Eth 3/4, Eth 3/5, Eth 3/6, Eth 3/7, Eth
3/8, Eth 3/9, Eth 3/10, Eth 3/11, Eth 3/12, Eth 3/13, Eth 3/14, Eth 3/15, Eth 3/16, Eth 3/17, Eth 3/18, Eth
3/19, Eth 3/20, Eth 3/21, Eth 3/22, Eth 3/23, Eth 3/24, Eth 3/26, Eth 3/27, Eth 3/28, Eth 3/29, Eth 3/30,
Eth 3/31, Eth 3/32, Eth 3/33, Eth 3/34, Eth 3/35, Eth 3/36, Eth 3/37, Eth 3/38, Eth 3/39, Eth 3/40, Eth
3/41, Eth 3/42, Eth 3/43, Eth 3/44, Eth 3/45, Eth 3/46, Eth 3/47, Eth 3/48, Eth 3/49, Eth 3/50, Eth 3/51,
Eth 3/52, Eth 3/53, Eth 3/54, Eth 3/55, Eth 3/56, Eth 3/57, Eth 3/58, Eth 3/59, Eth 3/60,

Default action:             DROP
Maximum number of flows allowed: 40960
Active flow:                 0

Ingress Replication Group ID range:      100 to 1000

```

Metering

Per-flow metering measures and controls the rate of packets for each flow entry. Per-flow meters enable OpenFlow to implement simple QoS operations, such as rate-limiting, and can be combined with per-port queues to implement complex QoS frameworks, such as DiffServ.

Meters are attached directly to flow entries. Each meter can have one or more meter bands. Each meter band specifies the rate of the band applies and the way packets are processed (DROP or DIFFSERV). OpenFlow metering operation is similar to ingress rate limiting in a QoS operation.

A meter table consists of meter entries. The counters in the following table are available in the meter entry.

TABLE 17 Meter entry

Counter	Description
Meter Identifier	A 32-bit unsigned integer uniquely identifying the meter
Meter band	Each meter band specifies the rate of the band and the way to process the packet. Rate and burst size are based on the line rate of the data traffic in contrast to the information rate.
Counter	Number of packets processed by a meter

Packets are processed by a single meter band based on the current measured meter rate. The meter applies the meter band with the highest configured rate that is lower than the current measured rate. If the current rate is lower than any specified meter band rate, no meter band is applied.

TABLE 18 Meter band supported on devices

Meter bands	Supported
DROP	Yes

Each band type contains the following meter configuration parameters from the controller:

- Rate value in kbps
- Rate value in packets per second
- Burst size
- Statistics collection

The metering system supports the features in the following table.

TABLE 19 Metering capabilities supported for metering features

Feature	SLX 9850-4	SLX 9850-8
Max meter	1K	1K
Band types (bitmap)	DROP	DROP
Capabilities (bitmap)	KBPS, BURST	KBPS, BURST
Maximum color value	2 (RED, GREEN)	2 (RED, GREEN)

Limitations

- Minimum burst size for DROP band is 10 kbits.
- Maximum burst size supported in hardware is 33292 kbits.

Meter messages

The following table describes the processing of the meter messages.

TABLE 20 Meter messages

Meter message type	Entry exists	Entry does not exist	Notes
Add (OFPMC_ADD)	Deny Add. Return error message to controller.	Add is processed.	Deny Add for the lack of memory or internal error or due to restrictions (hardware or otherwise) limiting the number

TABLE 20 Meter messages (continued)

Meter message type	Entry exists	Entry does not exist	Notes
			of bands. Return error message to controller as per standard.
Mod (OFPMC_MODIFY)	Meter parameters and bands are updated. New bands replace the existing bands.	Deny Mod. Return error message to controller.	Meter Mod updates meter parameters and new bands replacing the existing bands. Forward layers apply these changes instantly or make before break mechanism.
Del (OFPMC_DELETE)	Meter entry and bands removed. Flows associated with this meter are also removed.	No Error. Message ignored.	Only the meter identifier is specified for the delete request.

Show commands for OpenFlow

Show commands for OpenFlow are included in the following table and described in detail in the *Extreme SLX-OS Command Reference*.

TABLE 21 Show commands for OpenFlow

Command	Description
show openflow	Shows all the OpenFlow configuration.
show openflow flow	Shows all the flows configured in the system flow table.
show openflow controller	Shows the status of all the controllers.
show openflow interface	Displays the ports with OpenFlow.
show openflow resources	OpenFlow usage of the resources.
show openflow queues	Shows the queue entries for the interface.
show openflow group	Shows all the groups in a flow.
show openflow meter	Shows all the meters in a flow.

Clear commands for OpenFlow

Clear commands for OpenFlow are included in the following table and described in detail in the *Extreme SLX-OS Command Reference*.

When an individual OpenFlow rule or all flows in the flow table need to be deleted, you can use the **clear openflow** command. Use this command to delete a single OpenFlow rule based on a Flow ID or delete all flows/groups/meters configured in the system.

You can clear the flow statistics for all flows or for a specified flow.

TABLE 22 Clear commands for OpenFlow

Command	Description
clear openflow	Clears a single OpenFlow rule based on a Flow ID or deletes all flows/groups/meters configured in the system.
clear statistics openflow	Clears statistics for the controller, flow, all groups, or all meters.