CONFIGURATION GUIDE

# Extreme SLX-OS Network Packet Broker Configuration Guide, 18r.1.00

## Supporting the ExtremeRouting SLX 9850 Router

*9035462*
*31 August 2018*

# Contents

# Preface



# 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. |
| | Identifies text to enter in the GUI. |
| *italic* text | Identifies emphasis. |
| | Identifies variables. |
| | Identifies document titles. |
| `Courier font` | 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. |
| *italic* text | Identifies a variable. |
| [ ] | Syntax components displayed within square brackets are optional.<br>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, *member*[*member*...]. |
| \ | 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 https://www.extremenetworks.com/documentation-feedback/.
- Email us at documentation@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



# What's new in this document

This topic includes descriptions of changes in functionality for the current release.

On October 30, 2017, Extreme Networks, Inc. acquired the data center networking business from Brocade Communications Systems, Inc. This document has been updated to remove or replace references to Brocade Communications, Inc. with Extreme Networks, Inc., as appropriate.

# 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 Extreme Networks for the current 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
- ExtremeRouting SLX 9850-8

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 2** SLX 9850 interface modules capabilities

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

# Basics of Network Packet Broker



## Network Packet Broker overview

A Network Packet Broker (NPB) provides a collection of monitoring and analytics tools with access to traffic across the network.

An NPB is the part of a network visibility infrastructure responsible for aggregating network traffic and directing it to visibility applications. The following diagram represents a typical NPB context in data centers and enterprises for compliance, intrusion detection, network-performance monitoring, data recording, packet analysis, flow manipulation, forensics, and so forth.

**FIGURE 1** A Network Packet Broker in context

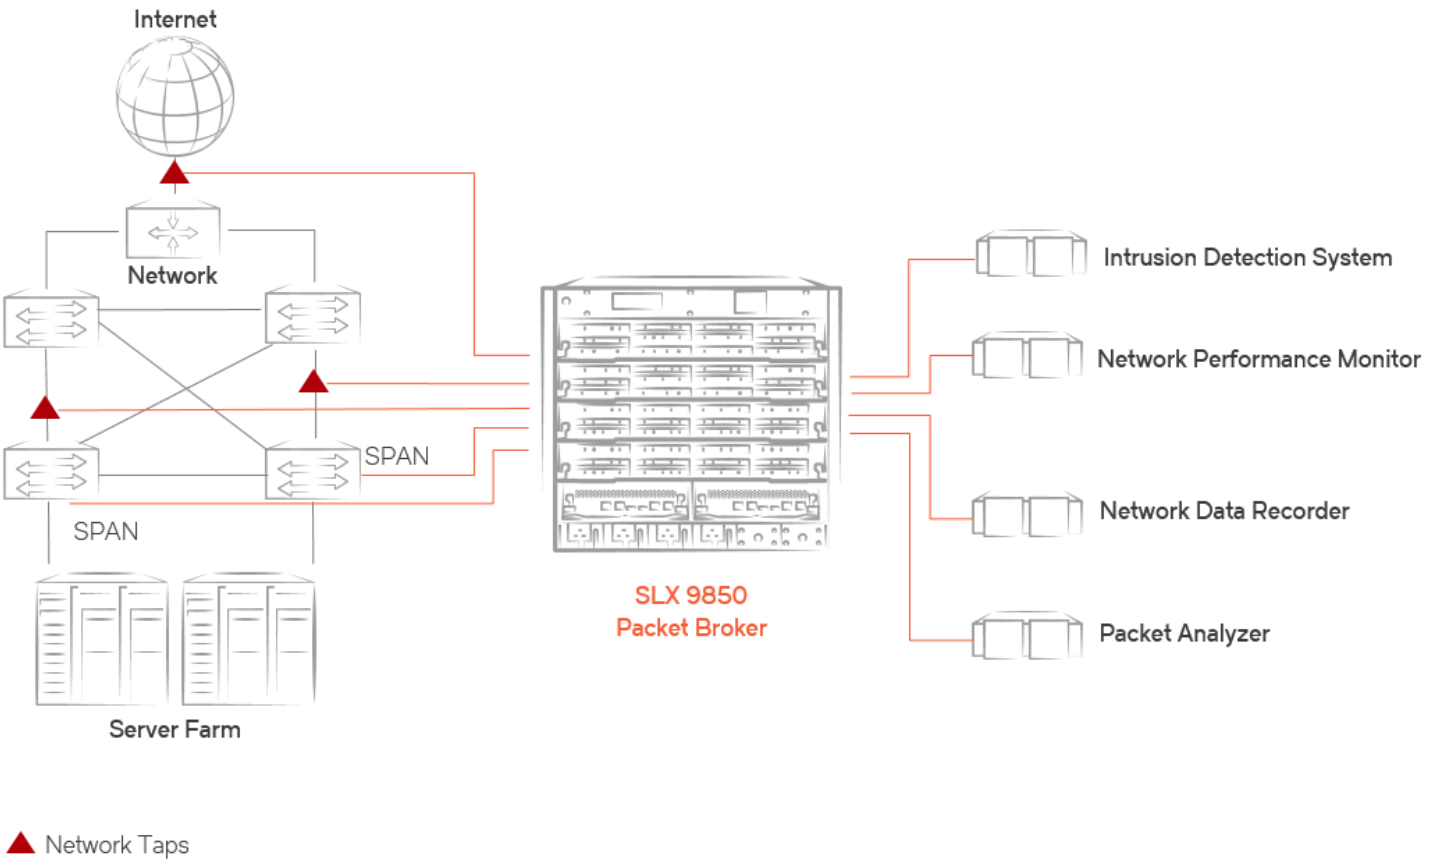


The basic requirements for using a SLX 9850 as an NPB are as follows:

- Manage all ingress and egress traffic under OpenFlow, as described in *Extreme SLX-OS SDN Configuration Guide*.
- Select the **npb-optimised-1** TCAM profile, as described in Configuring NPB hardware profiles on page 13.

For a sample implementation flow, refer to Header-stripping implementation flow on page 19.

## NPB features

SLX 9850 as a NPB supports the following features:

- Aggregation—the ability to aggregate traffic arriving from multiple TAPs or SPAN ports from upstream devices and direct the aggregated traffic to a single port or port-channel ("many to one").
- Replication—the ability to replicate network traffic to multiple ports and port-channels ("one to many").
- Load balancing—the ability to distribute network traffic among ports in a port-channel.
- ACL filtering (as described in the "ACLs" section of *Extreme SLX-OS Security Configuration Guide*:
  - Layer 2 and Layer 4 filtering—the ability to selectively direct network traffic based on fields of the Layer 2 and Layer 4 protocol headers.
  - User-defined ACLs (UDAs)—the ability to filter packets that Layer 2 and Layer 3 ACLs cannot filter.
- Encapsulation-header stripping—the ability to remove headers not supported by some visibility applications. Supported headers:
  - 802.1BR
  - VN-Tag
  - VXLAN

# NPB configuration guidelines

Follow these guidelines when implementing Network Packet Broker (NPB):

- A system reboot is required when moving to the required **profile tcam npb-optimised-1** TCAM profile.
- Most Layer 2 and Layer 3 configurations are not supported in NPB mode. Although implementing them does not generate error messages, do not configure them in NPB mode.
- Although port-channels are not supported, you can use OpenFlow Group concepts to manage sets of physical ports. For details, refer to *Extreme SLX-OS SDN Configuration Guide*.

# NPB TCAM-profiles

The only ternary content-addressable memory (TCAM) profile supported for Network Packet Broker (NPB) is **npb-optimised-1**.

Under **npb-optimised-1**, you can modify the default TCAM allocations for IPv4 and IPv6 flows.

TCAM for OpenFlow—12288 (12K) entries—is shared between IPv4 and IPv6 flows.

## Default and maximum NPB TCAM-profile settings

For the **npb-optimised-1** profile, the following table displays the default and configurable TCAM scaling limits per ASIC.

NOTE
For details of which ports share ASICs, refer to the "Association of ports to ASICs" tables in the *SLX 9850 Router Technical Specifications*.

**TABLE 3** Default and maximum OpenFlow TCAM entries per ASIC

| Configuration | IPv4 | IPv6 | Total |
|---|---|---|---|
| Default | 6144 (6K) | 4096 (4K) | 10240 (10K) |
| Modified | 0K–12K | 0K–12K | 12288(12K) |

For the **npb-optimised-1** profile, the following table displays total OpenFlow TCAM entries available on a SLX 9850.

**TABLE 4** Total maximum OpenFlow TCAM entries

| Type | ASIC | Line card | F4 chassis | F8 chassis |
|---|---|---|---|---|
| IPv4 and IPv6 flows | 12288 (12K) | 24576 (24K) | 100352 (98K) | 200704 (196K) |

# Configuring NPB hardware profiles

Use this task to select the counter and TCAM profiles required for Network Packet Broker (NPB), with an option to modify TCAM profile scaling limits.

1. Enter **configure terminal** to access global configuration mode.

```
device# configure terminal
```

2. Enter **hardware** to access hardware configuration mode.

```
device(config)# hardware
```

3. Enter **profile counters counter-profile-4**.

```
device(config-hardware)# profile counters counter-profile-4
```

4. Enter **profile tcam npb-optimised-1**.

```
device(config-hardware)# profile tcam npb-optimised-1
%INFO: For network packet broker (NPB) stats functionality select appropriate counter profile.
%Warning: To activate the new profile config, please run 'copy running-config startup-config'
followed by 'reload system'.
```

NOTE
By default, the IPv4 scaling limit is 6K and the IPv6 scaling limit is 4K.

5. To modify the default scaling limits, enter the **profile tcam limit** command, specifying the limits you require.

```
device(config-hardware)# profile tcam limit l2l3v4-app 10240 v6-app 2048
```

NOTE
To modify the default user-defined ACL (UDA) scaling limits, you can also include the **flex-acl** option.

6. Enter **end** to return to privileged EXEC mode.

```
device(config-hardware)# end
```

7. Enter **copy running-config startup-config**.

```
device# copy running-config startup-config
```

8. Enter **reload system** to reboot the system, effecting the TCAM profile change.

```
device# reload system
Warning: This operation will cause the chassis to reboot and requires all existing telnet,
secure telnet and SSH sessions to be restarted.
Unsaved configuration will be lost.
Please run `copy running-config startup-config` to save the current configuration if not done
already.
Are you sure you want to reboot the chassis [y/n]?
```

9. Press **y** and then **Enter**.

# Modifying NPB TCAM-profile scaling limits

Use this task to modify the IPv4, IPv6, and flex-ACL (UDA) scaling limits of the TCAM profile required for Network Packet Broker (NPB).

NOTE
This feature is available is available only if you first set the TCAM profile to **npb-optimised-1**, as described in Configuring NPB hardware profiles on page 13.

1. Enter **configure terminal** to access global configuration mode.

```
device# configure terminal
```

2. Enter **hardware** to access hardware configuration mode.

```
device(config)# hardware
```

3. Enter the **profile tcam limit l2l3v4-app** command, specifying the limits you require.

```
device(config-hardware)# profile tcam limit l2l3v4-app 10240 v6-app 2048
```

4. Enter **end** to return to privileged EXEC mode.

```
device(config-hardware)# end
```

5. Enter **copy running-config startup-config**.

```
device# copy running-config startup-config
```

6. Enter **reload system** to reboot the system, effecting the NPB TCAM-profile scaling change.

```
device# reload system
Warning: This operation will cause the chassis to reboot and requires all existing telnet,
secure telnet and SSH sessions to be restarted.
Unsaved configuration will be lost.
Please run `copy running-config startup-config` to save the current configuration if not done
already.
Are you sure you want to reboot the chassis [y/n]?
```

7. Press **y** and then **Enter**.

# Header Modification



## Header-modification overview

Protocol headers help packets reach their destinations, but are not needed by the security and monitoring tools to which NPB forwards traffic.

Tagging and encapsulation techniques have long been a part of networking. However, the recent adoption of new encapsulation protocols—such as VXLAN, VN-Tag, and 802.1BR—can create visibility blind spots, because some visibility applications were not designed to interpret these new protocols.

By removing the encapsulation header, the NPB removes the burden of interpreting the various encapsulation protocols from the visibility applications. Therefore, the header-stripping feature enables network operators to deploy new encapsulation protocols without interfering with proper functioning of previously deployed visibility applications. Other header stripping benefits include:

- Reduced packet overhead and better visibility-application bandwidth utilization.
- Preservation of standard filtering and load balancing capabilities.

For the current SLX-OS version, 802.1BR, VN-Tag, and VXLAN header stripping are the only types of header modification supported.

NOTE
Although NPB is supported only on SLX 9850, header stripping is supported also on SLX 9540.

## Header-stripping configuration guidelines

Follow these guidelines when implementing header stripping.

- The primary use case for header stripping is Network Packet Broker (NPB).
- Header stripping applies to both IPv4 and IPv6 packets.

802.1BR header-stripping removes only 802.1BR headers. VN-Tag header-stripping removes only the VN-Tag headers. VXLAN header-stripping removes multiple headers, as described in the following table:

**TABLE 5** Header-stripping scope

| Type | Headers stripped |
|---|---|
| 802.1BR | 802.1BR headers |
| VN-Tag | VN-Tag headers |
| VXLAN | Layer 2, Layer 3, UDP, and VXLAN headers |

Among the supported types of header stripping, you can never enable more than one type on a specific interface. However, if one of two applied types is VXLAN header-stripping, you can apply one other type on other slot interfaces. These limitations are summarized in the following table.

**TABLE 6** Header-stripping compatibility matrix

| Type | 802.1BR | VN-Tag | VXLAN |
|---|---|---|---|
| 802.1BR | — | Not on the same slot. | Not on the same interface. |
| VN-Tag | Not on the same slot. | — | Not on the same interface. |
| VXLAN | Not on the same interface. | Not on the same interface. | — |

After you enable header stripping on an interface, there are cases for which you need to reboot the linecard:

```
device# show packet-encap-processing
Total number of packet-encap-processing interfaces: 2
------------------------------------------------------------
Port       Link     Encapsulation        Status
------------------------------------------------------------
eth1/1     Up       VN-Tag Stripping     Inactive
eth1/2     Up       VN-Tag Stripping     Inactive

device# power-off linecard 1
device# power-on linecard 1
```

**TABLE 7** Linecard rebooting following header-stripping enablement

| Type | Rebooting |
|---|---|
| VN-Tag | Rebooting is required only after transition from 802.1BR header-stripping. |
| VXLAN | Rebooting is not required. |
| 802.1BR | Rebooting is required only after transition from VN-Tag header-stripping. |

**NOTE**
If linecard rebooting is required for a header-stripping type, one rebooting is sufficient for all linecard interfaces enabled.

The traffic-redirection technology varies among the types of header stripping.

**TABLE 8** Traffic-redirection technologies

| 802.1BR | VN-Tag | VXLAN |
|---|---|---|
| OpenFlow | OpenFlow | Policy-based routing (PBR) |

# Configuring 802.1BR header stripping

Use this task to enable or disable 802.1BR header stripping on an Ethernet interface.

1. Enter **configure terminal** to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface ethernet** command to access Ethernet interface configuration mode.

```
device(config)# interface ethernet 1/2
```

3. To enable 802.1BR header stripping on that interface, enter the **strip-802-1br** command.

```
device(conf-if-eth-1/2)# strip-802-1br
```

4. To disable 802.1BR header stripping on that interface, enter the **no strip-802-1br** command.

```
device(conf-if-eth-1/2)# no strip-802-1br
```

5. NOTE
   (If you are moving from VN-Tag header-stripping) For each linecard, you need to perform the following step only for the first interface enabled for 802.1BR header-stripping:

   Enter the **power-off linecard** command and then the **power-on linecard** command.

```
device# power-off linecard 1
Linecard 1 is being powered-off
device# power-on linecard 1
Linecard 1 is being powered-on
```

The following example enables 802.1BR header stripping on an interface.

```
device# configure terminal
device(config)# interface ethernet 1/2
device(conf-if-eth-1/2)# strip-802-1br
```

# Configuring VN-Tag header stripping

Use this task to enable or disable Virtual NIC (VN)-Tag header stripping on an Ethernet interface.

1. Enter **configure terminal** to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface ethernet** command to access Ethernet interface configuration mode.

```
device(config)# interface ethernet 1/2
```

3. To enable VN-Tag header stripping on that interface, enter the **strip-vn-tag** command.

```
device(conf-if-eth-1/2)# strip-vn-tag
%WARNING: Please reboot the linecard <> for stripping configuration to take effect.
```

4. To disable VN-Tag header stripping on that interface, enter the **no strip-vn-tag** command.

```
device(conf-if-eth-1/2)# no strip-vn-tag
```

5. Enter **end** to return to privileged EXEC mode.

```
device(config-hardware)# end
```

6. NOTE
   For each linecard, you need to perform the following step only for the first interface enabled for VN-Tag header-stripping:

   Enter the **power-off linecard** command and then the **power-on linecard** command.

```
device# power-off linecard 1
Linecard 1 is being powered-off
device# power-on linecard 1
Linecard 1 is being powered-on
```

The following example enables VN-Tag header stripping on an interface.

```
device# configure terminal
device(config)# interface ethernet 1/2
device(conf-if-eth-1/2)# strip-vn-tag
    %WARNING: Please reboot the linecard <> for stripping configuration to take effect.
device(config-hardware)# end
device# power-off linecard 1
Linecard 1 is being powered-off
device# power-on linecard 1
Linecard 1 is being powered-on
```

# Configuring VXLAN header stripping

Use this task to enable or disable Virtual Extensible LAN (VXLAN) header stripping on an Ethernet interface.

1. Enter **configure terminal** to access global configuration mode.

```
device# configure terminal
```

2. Enter the **interface ethernet** command to access Ethernet interface configuration mode.

```
device(config)# interface ethernet 1/2
```

3. To enable VXLAN header stripping on that interface, enter the **strip-vxlan** command.

```
device(conf-if-eth-1/2)# strip-vxlan
```

4. To disable VXLAN header stripping on that interface, enter the **no strip-vxlan** command.

```
device(conf-if-eth-1/2)# no strip-vxlan
```

The following example enables VXLAN header stripping on an interface.

```
device# configure terminal
device(config)# interface ethernet 1/2
device(conf-if-eth-1/2)# strip-vxlan
```

# Header-stripping implementation flow

This NPB implementation-flow includes 802.1BR or VN-Tag header stripping.

**FIGURE 2** NPB header-stripping flow

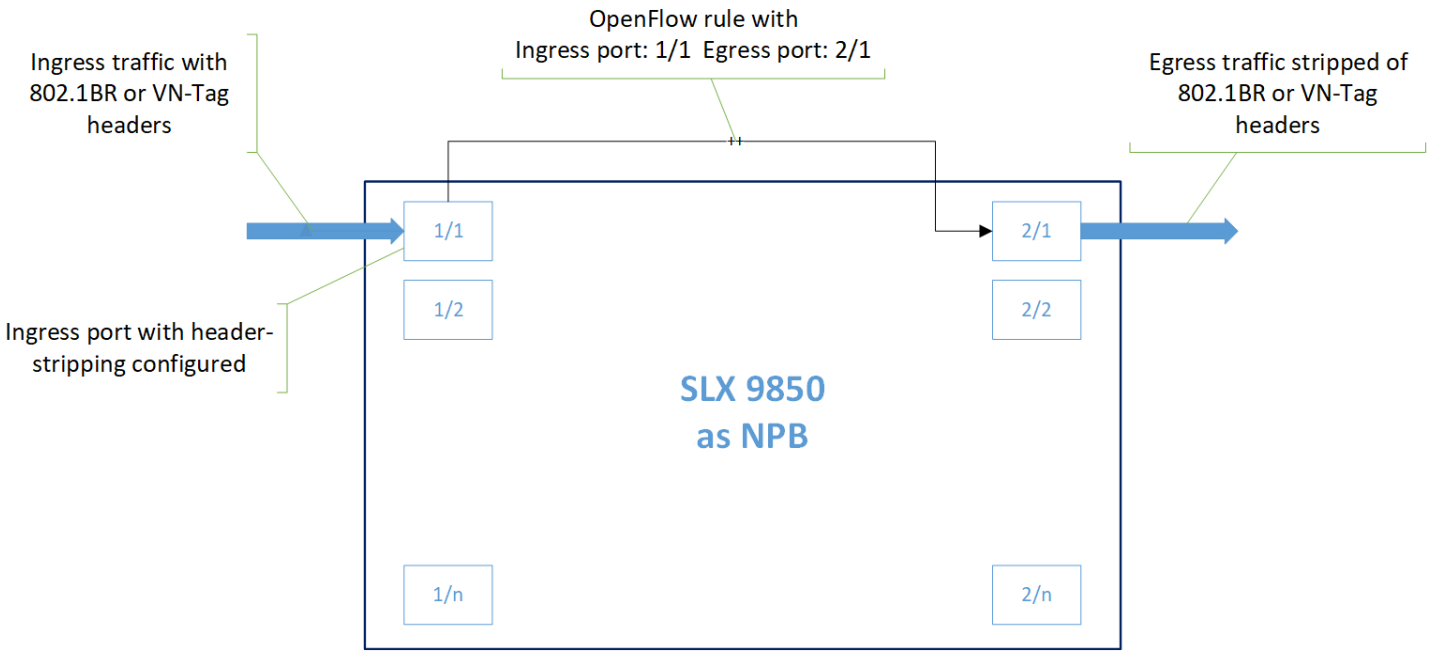


Port 1/1 is configured to strip one of the supported header types from incoming traffic and to forward the modified traffic—under OpenFlow—to port 2/1. Port 2/1 is configured to forward the modified traffic to the analytics tools.

> **NOTE**
> For VXLAN header-stripping, use policy-based routing (PBR) to redirect traffic.

1. Make sure that the **npb-optimised-1** TCAM profile is configured on the device, as described in Configuring NPB hardware profiles on page 13.
2. Enable OpenFlow on the ingress interface.

```
device# configure terminal
device(config)# interface ethernet 1/1
device(conf-if-eth-1/1)# openflow enable layer3
```

3. Enable the header-stripping feature on the ingress interface.

```
device(conf-if-eth-1/1)# strip-802-1br
device(conf-if-eth-1/1)# exit
```

4. Enable OpenFlow on the egress interface.

```
device(config)# interface ethernet 2/1
device(conf-if-eth-2/1)# openflow enable layer3
device(conf-if-eth-2/1)# end
```

5. Run the **show openflow interface** command to get needed details.

```
device# show openflow interface
Total number of Openflow interfaces: 2

Port       Link    Port-State   Speed   MAC               OF-Port-ID  Mode
Eth 1/1    Down    Forward      10G     748e.f88f.9a45    4           l3
Eth 2/1    Down    Forward      10G     748e.f88f.9a46    8           l3
```

6. Program the OpenFlow rule from the OpenFlow controller.

```
$ ovs-ofctl add-flow tcp:10.24.4.180:6633 "in_port=4 action=output:8"  --protocol Openflow13 –v
```

# NPB show commands

There are several show commands that display Network Packet Broker (NPB) information, as listed in the following table.

**TABLE 9** NPB show commands in the *Command Reference*

| Command | Description |
|---|---|
| **show interface ethernet** | Displays the detailed configuration and capabilities of a specific interface, including loopback status. |
| **show packet-encap-processing** | Displays information about interfaces on which header processing is enabled. |
| **show running-config interface ethernet** | Displays the contents of the running configuration—including header-stripping settings—on a specific Ethernet interface. |