

Extreme SLX-OS Network Packet Broker Configuration Guide, 17s.1.02

Supporting the ExtremeSwitching SLX 9140 and SLX 9240 Switches

© 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 |
| What's new in this document..... | 9 |
| Supported hardware and software..... | 9 |
| Basics of Network Packet Broker..... | 11 |
| NPB overview..... | 11 |
| Route maps under NPB..... | 12 |
| NPB configuration guidelines..... | 12 |
| Forwarding priority guidelines..... | 13 |
| Stanza and ACL permit and deny keywords..... | 13 |
| Priority among Layer 2 and Layer 3 match acl statements | 14 |
| Setting NPB as the system mode | 14 |
| Creating ACLs for NPB..... | 15 |
| NPB show commands | 15 |
| Traffic Aggregation and Replication..... | 17 |
| Creating ACLs for NPB..... | 17 |
| Aggregating traffic from interfaces..... | 17 |
| Replication to multiple interfaces..... | 18 |
| Transparent VLAN flooding..... | 18 |
| Creating TVF domains..... | 19 |
| Assigning a TVF domain to a physical egress interface..... | 19 |
| Assigning a TVF domain to a port-channel egress interface..... | 20 |
| Replicating traffic to multiple interfaces..... | 21 |
| Load Balancing Under NPB..... | 23 |
| Load balancing overview..... | 23 |
| Configuring symmetric load balancing..... | 23 |
| Symmetric load-balancing options and examples..... | 24 |
| Forwarding traffic to a port-channel..... | 26 |
| Header Modification..... | 29 |
| Header-modification overview..... | 29 |
| VLAN header stripping..... | 29 |
| GTP-HTTPS encapsulated filtering..... | 29 |
| Enabling and disabling GTP-HTTPS filtering..... | 30 |

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. |
| { x y z } | Default responses to system prompts are enclosed in square brackets. 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

- [What's new in this document.....](#) 9
- [Supported hardware and software.....](#) 9

What's new in this document

This document is the first public version of the *Extreme SLX-OS Network Packet Broker Configuration Guide*, for Extreme SLX-OS 17s. 1.02.

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, Inc. for this SLX-OS release, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by this release:

- ExtremeSwitching SLX 9140
- ExtremeSwitching SLX 9240

NOTE

Some of the commands in this document use a slot/port designation. Because the SLX 9140 and the SLX 9240 do not contain line cards, the slot designation must always be "0" (for example, 0/1 for port 1).

Basics of Network Packet Broker

| | |
|--|----|
| • NPB overview..... | 11 |
| • NPB configuration guidelines..... | 12 |
| • Setting NPB as the system mode | 14 |
| • Creating ACLs for NPB..... | 15 |
| • NPB show commands | 15 |

NPB overview

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

When you configure and reboot an SLX-OS device into NPB mode, the device can function as a Network Packet Broker.

The SLX-OS NPB implementation includes the following features:

- Aggregation: Traffic on multiple ingress interfaces is aggregated and forwarded on a single egress interface ("many to one") to one monitoring tool.
- Replication: Traffic on a single ingress interface is replicated and forwarded on multiple egress interfaces ("one to many") to multiple monitoring tools.
- Load balancing: Aggregation traffic and replication traffic are load-balanced across the members of an egress port-channel interface. Symmetric load balancing is the only type of load balancing supported.
- Header modification: Header stripping can reduce packet overhead, increasing the efficiency of the security and monitoring tools. The dropping of GPRS Tunneling Protocol (GTP) frames that encapsulate HTTPs packets is also supported.
- Timestamping: With Precision Time Protocol (PTP) accuracy, ingress and egress timestamps aid analysis of congestion and security issues. For details, refer to the "Precision Time Protocol (PTP)" section of the *Extreme SLX-OS Management Configuration Guide for SLX 9140 and SLX 9240*.

The following table indicates which interface types are supported for the aggregation, replication, load balancing, and timestamping features:

TABLE 1 Ingress and egress interface types

| Feature | Ingress interfaces | Egress interfaces |
|----------------|--|--|
| Aggregation | Multiple physical, port-channel, or mixed interfaces | One physical or port-channel interface |
| Replication | One physical or port-channel interface | A TVF domain, including multiple physical, port-channel, or mixed interfaces |
| Load balancing | Not supported | One port-channel interface |
| Timestamping | Physical and port-channel interfaces | Physical and port-channel interfaces |

Route maps under NPB

A route map is a container for one or more numbered permit or deny stanzas; each stanza contains a sequence of statements.

Evaluation of route maps consists of a list scan, from the lowest-numbered stanza to the highest-numbered stanza. Within each stanza, statements are evaluated in order. The following technologies are some that use route maps:

- BGP and OSPF protocols. For details, refer to the *Extreme SLX-OS Layer 3 Routing Configuration Guide for SLX 9140 and SLX 9240*.
- Policy-based routing (PBR) (not supported)
- Network Packet Broker (NPB)

Route-map syntax and evaluation vary with the technology. However, "match" statements are common to all route-map implementations. Upon a match, the actions specified in the match statement and in the remaining statements of that stanza are implemented. The list scan ends without examining higher-numbered stanzas.

Route-map "set" statements are also supported for the mentioned technologies. Upon a match, set statements perform an action on the matched traffic. The action varies with the technology. The following table compares route-map features and functionality for the various technologies.

TABLE 2 Route-map comparison

| Feature | NPB | BGP | PBR (not supported) |
|--|---|---|---|
| Traffic routing or forwarding | Yes | Yes | Yes |
| Traffic redistribution | No | Yes | No |
| Route-attribute modification | No | Yes | No |
| Stanzas | One or more permit or deny stanzas | One or more permit or deny stanzas | One or more permit or deny stanzas |
| Match statements | One or more match { mac ip ipv4 } address acl statements | One or more supported match statements | One or more match { ip ipv4 } address acl statements |
| Set statements (supported only in permit stanzas) | set interface or set next-hop-tvf-domain | 0, 1, or multiple set statements | 0, 1, or multiple set statements |
| Continue statements | No | 0 or 1 continue statements | No |
| Interfaces | Physical and port-channel interfaces | Device-level | Layer 3 interfaces |

NPB configuration guidelines

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

- NPB requires a license. For details, refer to the *Extreme SLX-OS Software Licensing Guide for SLX 9140 and SLX 9240*.
- A system reboot is required when moving from default system mode to NPB mode and conversely.
- 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.

- Only one route map can be applied per interface as NPB policy. However, that route map can have multiple stanzas.
- You can apply a given route map on multiple interfaces.
- The only supported match statements are **match { mac | ip | ipv6 } address acl** statements. Other match statements are ignored.
- The only supported set statements are **set interface** and **set next-hop-tvf-domain**. Other set statements are ignored.
- If a stanza does not contain a match statement, "match any" is implied.
- The only supported egress interfaces are physical interfaces, port-channel interfaces, and TVF domains.
- If a physical interface is part of a port-channel, a **set interface** statement on that interface is ignored.

Forwarding priority guidelines

The following guidelines determine forwarding priority in an NPB route-map:

- In general, forwarding priority is determined by the first match in the lowest-numbered stanza. For exceptions, refer to [Priority among Layer 2 and Layer 3 match acl statements](#) on page 14.
- The order in which the egress interfaces are configured is also the order for configuring forwarding.
- If a route map has multiple egress interfaces, the first egress interface ready for forwarding is used.
 - A physical interface is considered ready for forwarding if its link is up.
 - Port-channels and TVF domains are considered ready if at least one member port link is up.
- If the current egress interface loses its Ready status, then the next configured egress interface that is ready for forwarding is used.
- If a Down egress interface becomes Ready and if it has higher priority than the currently selected egress interface, the higher priority egress interface replaces it.
- If none of the configured egress interfaces are ready, the traffic is dropped.

Stanza and ACL permit and deny keywords

Both route-map stanzas and access-control lists (ACLs) have **permit** and **deny** keywords.

Both NPB and PBR route-maps contain **match { mac | ip | ipv4 } address acl** statements. However, permit and deny rules in ACLs applied to route maps function differently than rules in the security ACLs discussed in the *Extreme SLX-OS Security Configuration Guide for SLX 9140 and SLX 9240*:

- In security ACLs, permit rules allow packets and deny rules drop packets.
- In ACLs applied to PBR route-maps, permit and deny rules specify criteria for route-map decisions.
- In ACLs applied to NPB route-maps, permit and deny rules are mechanisms to exclude certain traffic flows from set statements.

The following table describes the interactions between route-map permit and deny stanzas; and permit and deny rules in ACLs applied to those stanzas by **match { mac | ip | ipv4 } address acl** statements.

TABLE 3 Route-map stanza and ACL permit and deny interactions

| Stanza | ACL rule | Resulting TCAM action |
|--------|----------|---|
| Permit | Permit | The set statement or statements are applied. |

TABLE 3 Route-map stanza and ACL permit and deny interactions (continued)

| Stanza | ACL rule | Resulting TCAM action |
|--------|----------|---|
| Permit | Deny | Packets that match a deny keyword are denied from using the stanza set statement: <ul style="list-style-type: none"> NPB: The packet is dropped. PBR: The packet is routed as normal. |
| Deny | Permit | No action is taken: <ul style="list-style-type: none"> NPB: The packet is dropped. PBR: The packet is routed as normal. |
| Deny | Deny | No action is taken: <ul style="list-style-type: none"> NPB: The packet is dropped. PBR: The packet is routed as normal. |

Priority among Layer 2 and Layer 3 match acl statements

In general, forwarding priority is determined by the first match in the lowest-numbered stanza. However, there are exceptions to this priority, if a route map contains match statements from both of the following groups:

- Layer 2: **match mac address acl**
- Layer 3: **match ip address acl** and **match ipv6 address acl**

Priority among Layer 2 and Layer 3 **match acl** statements is as follows:

- If the match within one ACL is on a permit rule and the match in the other ACL is on a deny rule, the permit match is accepted and the deny match is ignored.
- If both matches are on permit rules, the Layer 3 match is accepted and the Layer 2 match is ignored.
- If there are both IPv4 and IPv6 matches, the first match in the lowest-numbered stanza is accepted and other matches are ignored.

Setting NPB as the system mode

Before enabling Network Packet Broker (NPB), you must set NPB as the system mode.

Accessing NPB mode requires a license. For details, refer to the *Extreme SLX-OS Software Licensing Guide for SLX 9140 and SLX 9240*.

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 **system-mode npb** to change the system mode to NPB.

```
device(config-hardware)# system-mode npb
%Warning: To activate the new system-mode config, please reboot the system using 'reload system'.
```

To return to default system mode from NPB system mode, enter **system-mode default** rather than **system-mode npb**.

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

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

5. To reboot the system, effecting the system mode change, enter **reload system**.

```
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]?
```

6. Press **y** and then **Enter**.

Creating ACLs for NPB

For NPB traffic aggregation and traffic replication, an access-control list (ACL) is required.

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

```
device# configure terminal
```

2. Enter the **{ mac | ip | ipv6 } access-list** command to create an ACL.

```
device(conf)# ip access-list standard aclNPB_01
```

3. Create one or more permit or deny rules.

```
device(conf-ipacl-std)# permit host 192.1.1.1 count
```

NOTE

A deny rule specifies that a matching packet is denied from using the **set** statement. For details, refer to [NPB overview](#) on page 11.

NPB show commands

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

TABLE 4 NPB show commands in the *Command Reference*

| Command | Description |
|---------------------------------------|--|
| show inner-gtp-https | Displays a list of all interfaces on which dropping of GPRS Tunneling Protocol (GTP) frames that encapsulate HTTPs packets is enabled. |
| show route-map | Displays configuration details of all route maps, of a specific route map, or of the route map applied to an interface. |
| show running-config tvf-domain | Displays configuration details of one or all TVF domains. |

Traffic Aggregation and Replication

| | |
|--|----|
| • Creating ACLs for NPB..... | 17 |
| • Aggregating traffic from interfaces..... | 17 |
| • Replication to multiple interfaces..... | 18 |

Creating ACLs for NPB

For NPB traffic aggregation and traffic replication, an access-control list (ACL) is required.

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

```
device# configure terminal
```

2. Enter the { **mac | ip | ipv6** } **access-list** command to create an ACL.

```
device(config)# ip access-list standard aclNPB_01
```

3. Create one or more permit or deny rules.

```
device(config-ipacl-std)# permit host 192.1.1.1 count
```

NOTE

A deny rule specifies that a matching packet is denied from using the **set** statement. For details, refer to [NPB overview](#) on page 11.

Aggregating traffic from interfaces

This task aggregates and forwards traffic from multiple interfaces to a single, physical egress interface.

1. Configure an ACL, as described in [Creating ACLs for NPB](#) on page 15.

```
device(config)# ip access-list standard aclNPB_01
device(config-ipacl-std)# permit host 192.1.1.1 count
device(config-ipacl-std)# exit
device(config)#
```

2. Configure a route map that contains the relevant **match { mac | ip | ipv6 } address acl** command.

```
device(config)# route-map npb_map1 permit 1
device(config-route-map-npb_map1/permit/1)# match ip address acl acl_2
```

NOTE

For a MAC ACL, use the **match mac address acl** command. For an IPv6 ACL, use the **match ipv6 address acl** command.

3. Specify the route-map egress physical interface.

- Without stripping 802.1q VLAN tags:

```
device(config-route-map-npb_map1/permit/1)# set interface ethernet 0/5
```

- Stripping 802.1q VLAN tags:

```
device(config-route-map-npb_map1/permit/1)# set interface ethernet 0/5 strip-vlan outer
```

4. Exit to global configuration mode.

```
device(config-route-map-npb_map1/permit/1)# exit
```

5. Apply the route map to the ingress interfaces.

- Physical interfaces

```
device(config)# interface ethernet 0/2
device(conf-if-eth-0/2)# npb policy route-map npb_map1
```

- Port-channel interfaces

```
device(config)# interface port-channel 10
device(config-Port-channel-10)# npb policy route-map npb_map1
```

The following example configures ingress traffic from Ethernet 0/1 and port-channel 100 to egress Ethernet 0/5.

```
device# configure terminal
device(config)# route-map npb_map permit 10
device(config-route-map-npb_map/permit/10)# match ip address acl acl_2
device(config-route-map-npb_map/permit/10)# set interface ethernet 0/5
device(config-route-map-npb_map/permit/10)# exit
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# npb policy route-map npb_map
device(conf-if-eth-0/1)# exit
device(config)# interface port-channel 100
device(config-Port-channel-100)# npb policy route-map npb_map
```

Replication to multiple interfaces

Traffic on an ingress interface is replicated and forwarded on multiple egress interfaces, usually to multiple monitoring tools.

The replication ingress-interface is one physical or port-channel interface.

Multiple physical and port-channel egress interfaces are supported. However, you first need to associate such interfaces with a Transparent VLAN flooding (TVF) domain. You then use a route map to designate that TVF domain as egress. Ingress traffic is replicated, and forwarded by way of the TVF to the egress interfaces that you specified.

Transparent VLAN flooding

Transparent VLAN flooding (TVF) forwards packets without any form of CPU intervention, including MAC learning and MAC destination lookups.

This implementation of TVF has the following attributes:

- Traffic is distributed in hardware to all members of the TVF domain.
- Because this feature does not use any MAC address entries in the CAM, it is useful when MAC address entries need to be conserved.

- You can create as many as 4096 TVF domains.
- Domain members can be tagged and untagged ports. There is no software limitation on the number of member ports.
- You can mix and match ports with different speeds.
- At the TVF domain level, load balancing does not occur. If you need load balancing, associate a port-channel with the TVF domain, which balances the loads among the interfaces included in the port-channel.
- CPU intervention is not required, enabling line-rate traffic forwarding.

Creating TVF domains

This task creates one or more Transparent VLAN flooding (TVF) domains.

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

```
device# configure terminal
```

2. Enter the **tvf-domain** command, specifying one of the valid formats.

- To create one TVF domain, specify an integer from 1 through 4096.

```
device(config)# tvf-domain 10
```

- To specify a range of TVF domains, insert a hyphen (-) between the beginning and ending integers.

```
device(config)# tvf-domain 20-30
```

- To specify individual domains and ranges of domains, separate them with commas (for example, 1,5-7,55).

```
device(config)# tvf-domain 1,5-7,55
```

3. Enter **exit** to return to global configuration mode.

```
device(config-tvf-domain-10)# exit
device(config)#
```

Assigning a TVF domain to a physical egress interface

This task assigns a TVF domain to a physical egress interface.

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

```
device# configure terminal
```

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

```
device(config)# interface ethernet 0/5
```

3. Enter the **tvf-domain** command, specifying one of the valid formats.

- To assign one TVF domain to the interface, specify its integer ID.

```
device(conf-if-eth-0/5)# tvf-domain add 10
```

- To assign a range of TVF domains to the interface, insert a hyphen (-) between the beginning and ending integers.

```
device(conf-if-eth-0/5)# tvf-domain add 20-30
```

- To assign individual domains and ranges of domains, separate them with commas (for example, 1,5-7,55).

```
device(conf-if-eth-0/5)# tvf-domain add 1,5-7,55
```

- To assign all defined TVF domains to the interface, enter **tvf-domain all**.

```
device(conf-if-eth-0/5)# tvf-domain all
```

- To assign all defined TVF domains to the interface—except for those specified—enter the **tvf-domain except** option.

```
device(conf-if-eth-0/5)# tvf-domain except 1,2,4-7
```

Assigning a TVF domain to a port-channel egress interface

This task assigns a TVF domain to a port-channel egress interface.

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

```
device# configure terminal
```

2. Enter the **interface port-channel** command to access port-channel configuration mode.

```
device(config)# interface port-channel 10
```

3. Enter the **tvf-domain** command, specifying one of the valid formats.

- To assign one TVF domain to the interface, specify its integer ID.

```
device(config-Port-channel-10)# tvf-domain add 10
```

- To assign a range of TVF domains to the interface, insert a hyphen (-) between the beginning and ending integers.

```
device(config-Port-channel-10)# tvf-domain add 20-30
```

- To assign individual domains and ranges of domains, separate them with commas (for example, 1,5-7,55).

```
device(config-Port-channel-10)# tvf-domain add 1,5-7,55
```

- To assign all defined TVF domains to the interface, enter **tvf-domain all**.

```
device(config-Port-channel-10)# tvf-domain all
```

- To assign all defined TVF domains to the interface—except for those specified—enter the **tvf-domain except** option.

```
device(config-Port-channel-10)# tvf-domain except 1,2,4-7
```

Replicating traffic to multiple interfaces

This task replicates traffic entering an interface to multiple egress interfaces.

1. Create needed TVF domains, as described in [Creating TVF domains](#) on page 19.

```
device# configure terminal
device(config)# tvf-domain 10
device(config-tvf-domain-10)# exit
```

2. Assign TVF domains to the egress interfaces.

- Physical interfaces (For details, refer to [Assigning a TVF domain to a physical egress interface](#) on page 19.)

```
device(config)# interface ethernet 0/5
device(config-if-eth-0/5)# tvf-domain add 10
```

- Port-channel interfaces

```
device(config)# interface port-channel 10
device(config-Port-channel-10)# tvf-domain add 10
```

3. Configure an ACL, as described in [Creating ACLs for NPB](#) on page 15.

```
device(config)# ip access-list standard aclNPB_01
device(config-ipacl-std)# permit host 192.1.1.1 count
device(config-ipacl-std)# exit
```

4. Configure a route map that contains the relevant **match { mac | ip | ipv6 } address acl** command.

```
device(config)# route-map npb_map1 permit 1
device(config-route-map-npb_map1/permit/1)# match ip address acl acl_2
```

5. Specify the TVF domain that contains the egress interfaces for the replicated traffic.

- Without stripping 802.1q VLAN tags:

```
device(config-route-map-npb_map1/permit/1)# set next-hop-tvf-domain 5
```

- Stripping 802.1q VLAN tags:

```
device(config-route-map-npb_map1/permit/1)# set next-hop-tvf-domain 5 strip-vlan outer
```

6. Exit to global configuration mode.

```
device(config-route-map-npb_map1/permit/1)# exit
```

7. Apply the route map to the ingress interface.

- Physical interface

```
device(config)# interface ethernet 0/2
device(config-if-eth-0/2)# npb policy route-map npb_map1
```

- Port-channel interface

```
device(config)# interface port-channel 10
device(config-Port-channel-10)# npb policy route-map npb_map1
```

The following example replicates traffic entering an interface to multiple egress interfaces.

```
device# configure terminal
device(config)# tvf-domain 10
device(config-tvf-domain-10)# exit
device(config)# interface ethernet 0/5
device(conf-if-eth-0/5)# tvf-domain add 10
device(conf-if-eth-0/5)# exit
device(config)# interface port-channel 10
device(config-Port-channel-10)# tvf-domain add 10
device(config-Port-channel-10)# exit
device(config)# route-map npb_map1 permit 1
device(config-route-map-npb_map1/permit/1)# match ip address acl acl_2
device(config-route-map-npb_map1/permit/1)# set next-hop-tvf-domain 5
device(config-route-map-npb_map1/permit/1)# exit
device(config)# interface ethernet 0/2
device(conf-if-eth-0/2)# npb policy route-map npb_map1
```

Load Balancing Under NPB

| | |
|---|----|
| • Load balancing overview..... | 23 |
| • Configuring symmetric load balancing..... | 23 |
| • Forwarding traffic to a port-channel..... | 26 |

Load balancing overview

The only type of link-aggregation load-balancing supported under Network Packet Broker (NPB) is symmetric load balancing.


NOTE

For details on link aggregation, including link-aggregation groups (LAGs), refer to the "Link Aggregation" section of the *Extreme SLX-OS Layer 2 Switching Configuration Guide*. LAGs are also referred to as *port-channels*.

Symmetric load-balancing interchanges source and destination addresses to ensure that bidirectional traffic between a source and destination pair flows through one LAG member. Under NPB, symmetric load-balancing is enabled by default and cannot be disabled. However, you can modify the load-balancing parameters.

For network telemetry applications, network traffic is tapped and sent to a device that hashes selected traffic to the application servers downstream. For many monitoring and security applications, bidirectional conversations flowing through the system must be carried on the same port of a port-channel (LAG). In addition, firewalls between devices can be configured to allow such bidirectional conversations.

FIGURE 1 Symmetric load balancing



Configuring symmetric load balancing

Use this task to change the load-balancing mode from the default **src-dst-ip-mac-vid-port** mode or another current mode to the load-balancing mode that you require.

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

```
device# configure terminal
```

2. Enter the **load-balance** command to specify the mode that you require.

```
device(conf)# load-balance src-dst-ip
```

Symmetric load-balancing options and examples

The following tables display examples of the NPB symmetric load-balancing hashing options.

src-dst-ip

The following table displays inputs for a **load-balance src-dst-ip** example. The distribution is based on source and destination IPv4 or IPv6 addresses.

TABLE 5 Symmetric load-balancing **src-dst-ip** option

| Flow | Source IP | Destination IP | Source MAC | Destination MAC | VLAN |
|------|-----------|----------------|------------|-----------------|------|
| 1 | IPA | IPB | MACA | MACB | — |
| 2 | IPB | IPA | MACD | MACE | — |
| 3 | IPD | IPE | MACA | MACB | VID2 |
| 4 | IPA | IPC | MACA | MACD | VID2 |

Because flows 1 and 2 share the same source and destination IP addresses (the other fields are not considered), they are load balanced on the same port-channel port.

src-dst-ip-mac-vid

The following table displays inputs for a **load-balance src-dst-ip-mac-vid** example. The distribution is based on source and destination IPv4 or IPv6 and MAC addresses; and outer VLAN ID (VID).

TABLE 6 Symmetric load-balancing **src-dst-ip-mac-vid** option

| Flow | Source IP | Destination IP | Source MAC | Destination MAC | VLAN |
|------|-----------|----------------|------------|-----------------|------|
| 1 | IPA | IPB | MACA | MACB | VID1 |
| 2 | IPB | IPA | MACB | MACA | VID1 |
| 3 | IPA | IPB | MACA | MACB | VID2 |
| 4 | IPA | IPB | MACD | MACE | VID2 |

Because flows 1 and 2 share the same source and destination IP and MAC addresses and the same VLAN, they are load-balanced on the same port-channel port.

src-dst-ip-mac-vid-port

The following table displays inputs for a **load-balance src-dst-ip-mac-vid-port** example. The distribution is based on source and destination IPv4 or IPv6 and MAC addresses, VID, and TCP or UDP destination port.

NOTE

This is the default **load-balance** option.

TABLE 7 Symmetric load-balancing **src-dst-ip-mac-vid-port** option

| Flow | Source IP | Destination IP | Source MAC | Destination MAC | Source L4 Port | Destination L4 Port | VLAN |
|------|-----------|----------------|------------|-----------------|----------------|---------------------|------|
| 1 | IPA | IPB | MACA | MACB | P1 | P2 | VID1 |
| 2 | IPB | IPA | MACB | MACA | P2 | P1 | VID1 |
| 3 | IPA | IPB | MACA | MACB | P2 | P1 | VID2 |
| 4 | IPA | IPB | MACD | MACE | P3 | P4 | VID2 |

Because flows 1 and 2 share the same source and destination IP and MAC addresses, the same TCP or UDP ports, and the same VLAN, they are load-balanced on the same port-channel port.

src-dst-ip-port

The following table displays inputs for a **load-balance src-dst-ip-port** example. The distribution is based on source and destination IPv4 or IPv6 address and TCP or UDP destination port.

TABLE 8 Symmetric load-balancing **src-dst-ip-port** option

| Flow | Source IP | Destination IP | Source L4 Port | Destination L4 Port |
|------|-----------|----------------|----------------|---------------------|
| 1 | IPA | IPB | P1 | P2 |
| 2 | IPB | IPA | P2 | P1 |
| 3 | IPA | IPB | P2 | P1 |
| 4 | IPA | IPB | P3 | P4 |

Because flows 1 and 2 share the same source and destination IP addresses and the same TCP or UDP ports, they are load-balanced on the same port-channel port.

src-dst-mac-vid

The following table displays inputs for a **load-balance src-dst-mac-vid** example. The distribution is based on the source and destination MAC addresses and VID.

TABLE 9 Symmetric load-balancing **src-dst-mac-vid** option

| Flow | Source MAC | Destination MAC | VLAN |
|------|------------|-----------------|------|
| 1 | MACA | MACB | VID1 |
| 2 | MACB | MACA | VID1 |

TABLE 9 Symmetric load-balancing **src-dst-mac-vid** option (continued)

| Flow | Source MAC | Destination MAC | VLAN |
|------|------------|-----------------|------|
| 3 | MACA | MACB | VID2 |
| 4 | MACD | MACE | VID2 |

Because flows 1 and 2 share the same source and destination MAC addresses and the same VLAN, they are load-balanced on the same port-channel port.

Forwarding traffic to a port-channel

For load balancing, this task forwards traffic entering an interface to a port-channel.

1. Configure an ACL, as described in [Creating ACLs for NPB](#) on page 15.

```
device# configure terminal
device(config)# ip access-list standard aclNPB_01
device(conf-ipacl-std)# permit host 192.1.1.1 count
device(conf-ipacl-std)# exit
```

2. Configure a route map that contains the relevant **match { mac | ip | ipv6 } address acl** command.

```
device(config)# route-map npb_map1 permit 1
device(config-route-map-npb_map1/permit/1)# match ip address acl aclNPB_01
```

3. Specify the egress port-channel interface.

- Without stripping 802.1q VLAN tags:

```
device(config-route-map-npb_map1/permit/1)# set interface port-channel 10
```

- Stripping 802.1q VLAN tags:

```
device(config-route-map-npb_map1/permit/1)# set interface port-channel 10 strip-vlan outer
```

4. Exit to global configuration mode.

```
device(config-route-map-npb_map1/permit/1)# exit
```

5. Apply the route map to a single ingress interface:

- Physical interface

```
device(config)# interface ethernet 0/2
device(conf-if-eth-0/2)# npb policy route-map npb_map1
```

- Port-channel interface

```
device(config)# port-channel 5
device(config-Port-channel-10)# npb policy route-map npb_map1
```

The following example forwards traffic entering a physical interface to a port-channel.

```
device# configure terminal
device(config)# ip access-list standard aclNPB_01
device(conf-ipacl-std)# permit host 192.1.1.1 count
device(conf-ipacl-std)# exit
device(config)# route-map npb_map1 permit 1
device(config-route-map-npb_map1/permit/1)# match ip address acl aclNPB_01
device(config-route-map-npb_map1/permit/1)# set interface port-channel 10
device(config-route-map-npb_map1/permit/1)# exit
device(config)# interface ethernet 0/2
device(conf-if-eth-0/2)# npb policy route-map npb_map1
```


Header Modification

| | |
|---|----|
| • Header-modification overview..... | 29 |
| • VLAN header stripping..... | 29 |
| • GTP-HTTPS encapsulated filtering..... | 29 |

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.

Header stripping can reduce packet overhead, increasing the efficiency of the security and monitoring tools. Also, header stripping helps to avoid the following problems:

- Unstripped protocol headers can interfere with filtering and load balancing.
- Some protocol headers prevent a tool from accessing the information it needs.

Dropping GPRS Tunneling Protocol (GTP) frames that encapsulate HTTPs packets is also supported.

VLAN header stripping

NPB supports 802.1q VLAN outer-header stripping for aggregation, replication, and load balancing, as described in the following table. For implementation details, refer to the linked tasks.

TABLE 10 VLAN header-stripping commands and tasks

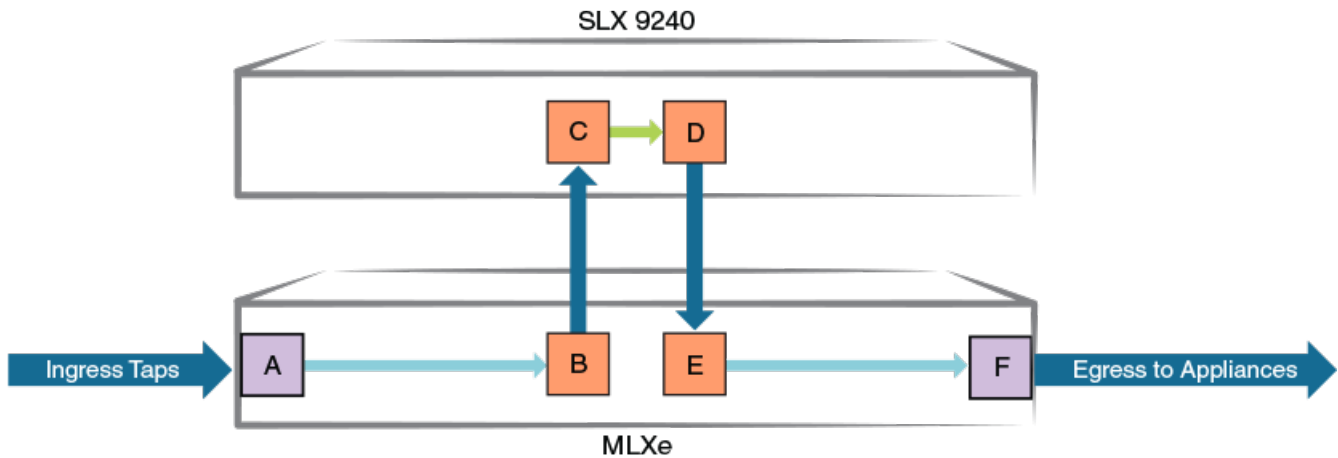
| Feature | set command | Task |
|----------------|--------------------------------|---|
| Aggregation | set interface | Aggregating traffic from interfaces on page 17 |
| Replication | set next-hop-tvf-domain | Replicating traffic to multiple interfaces on page 21 |
| Load balancing | set interface | Forwarding traffic to a port-channel on page 26 |

GTP-HTTPS encapsulated filtering

This feature enables you to drop—from ingress traffic—GPRS Tunneling Protocol (GTP) frames that encapsulate HTTPS packets.

NOTE

This feature applies to GTP v.1 frames.

FIGURE 2 Example of dropping GTP frames that encapsulate HTTPS packets

In the example flow, GTP frames that encapsulate HTTPS packets are removed, as follows:

1. Traffic enters an MLXe through port A.
2. The traffic exits the MLXe from port B and enters the SLX 9240 through port C. Because **deny inner-gtp-http** is enabled on port C, GTP frames that encapsulate HTTPS packets are removed from the flow.
3. Traffic is forwarded from SLX 9240 port C to port D.
4. Traffic is forwarded back to MLXe port E, using an NPB route map. (For details, refer to [Route maps under NPB](#) on page 12.)
5. Traffic is forwarded from MLXe port E to port F, using an NPB route map.
6. Traffic is forwarded from MLXe port F to monitoring tools.

Enabling and disabling GTP-HTTPS filtering

Use this task to enable and disable dropping GPRS Tunneling Protocol (GTP) frames that encapsulate HTTPS packets.

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

```
device# configure terminal
```

2. Access an interface on which you need to enable this feature.

- Physical interface:

```
device(conf)# interface ethernet 0/5
```

- Port-channel interface:

```
device(conf)# interface port-channel 10
```

3. Enable or disable this feature.

- Enable:

```
device(conf-if-eth-0/5)# deny inner-gtp-https
```

- Disable:

```
device(config-Port-channel-10)# no deny inner-gtp-https
```