



Release Notes for XA1400 Series

Release 8.1.50
9035868-02 RevAA
January 2020

© 2019-2020, Extreme Networks
All Rights Reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

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 (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, please see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: www.extremenetworks.com/support/policies/software-licensing

Contents

Chapter 1: About this Document	5
Purpose.....	5
Conventions.....	5
Text Conventions.....	5
Documentation and Training.....	7
Getting Help.....	8
Providing Feedback.....	9
Chapter 2: New in this Release	10
VOSS 8.1.50.....	10
IPsec NAT-T.....	11
Configure the IPsec Initiator with the IPsec Responder Remote NAT IP Address.....	14
Configure an IPsec NAT-T Responder.....	15
ipsec remote-nat-ip.....	18
ipsec responder-only.....	18
show io.....	19
Configure IPsec NAT-T.....	20
Fabric Extend Considerations.....	22
ONA considerations.....	26
Configure Fabric Extend.....	30
Configure Fabric Extend Over IPsec.....	34
Configure Fabric Extend Tunnels.....	37
View Tunnel CoS Queue Statistics.....	39
Fabric Extend Tunnel MTU.....	39
logical-intf isis	40
show isis logical-interface.....	41
Configure Egress Tunnel Shaping.....	42
Fabric Extend configuration using EDM.....	47
Port-Rate Limiting, Policing, and Shaping.....	51
IEEE 802.3X Pause Frame Transmit.....	53
Configure IEEE 802.3X Pause Frame Transmit.....	55
Configure basic port parameters.....	58
Configure Boot Flags.....	64
Configure IEEE 802.3X Pause frame transmit.....	68
Filenames for this Release.....	69
Documentation Changes.....	69
Chapter 3: Upgrade and Downgrade Considerations	70
Chapter 4: XA1400 Series Hardware	71
Chapter 5: Software Scaling	72
Layer 2.....	72

Contents

IP Unicast.....	72
Layer 3 Route Table Size.....	74
IP Multicast.....	74
Filters, QoS, and Security.....	74
Fabric Scaling.....	75
OAM and Diagnostics.....	75
Chapter 6: Important Notices.....	76
Subscription Licensing for XA1400 Series.....	76
Supported Browsers.....	78
Chapter 7: Known Issues and Restrictions.....	79
Known Issues and Restrictions.....	79
Filter Restrictions and Expected Behaviors.....	81
Chapter 8: Resolved Issues.....	83
Chapter 9: Related Information.....	84
MIB Changes.....	84

Chapter 1: About this Document

This section discusses the purpose of this document, the conventions used, ways to provide feedback, additional help, and information regarding other Extreme Networks publications.

Purpose

This document describes important information about this release for supported VSP Operating System Software (VOSS) platforms.

This document includes the following information:

- supported hardware and software
- scaling capabilities
- known issues, including workarounds where appropriate
- known restrictions

Conventions

This section discusses the conventions used in this guide.

Text Conventions

The following tables list text conventions that can be used throughout this document.

Table 1: Notice Icons



Icon	Alerts you to...
 Important:	A situation that can cause serious inconvenience.
 Note:	Important features or instructions.

Table continues...





Icon	Alerts you to...
 Tip:	Helpful tips and notices for using the product.
 Danger:	Situations that will result in severe bodily injury; up to and including death.
 Warning:	Risk of severe personal injury or critical loss of data.
 Caution:	Risk of personal injury, system damage, or loss of data.

Table 2: Text Conventions

Convention	Description
Angle brackets (< >)	<p>Angle brackets (< >) indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when you enter the command.</p> <p>If the command syntax is <code>cfm maintenance-domain maintenance-level <0-7></code> , you can enter <code>cfm maintenance-domain maintenance-level 4</code>.</p>
Bold text	<p>Bold text indicates the GUI object name you must act upon.</p> <p>Examples:</p> <ul style="list-style-type: none"> • Click OK. • On the Tools menu, choose Options.
Braces ({ })	<p>Braces ({ }) indicate required elements in syntax descriptions. Do not type the braces when you enter the command.</p> <p>For example, if the command syntax is <code>ip address {A.B.C.D}</code>, you must enter the IP address in dotted, decimal notation.</p>
Brackets ([])	<p>Brackets ([]) indicate optional elements in syntax descriptions. Do not type the brackets when you enter the command.</p> <p>For example, if the command syntax is <code>show clock [detail]</code>, you can enter either <code>show clock</code> or <code>show clock detail</code>.</p>
Ellipses (...)	<p>An ellipsis (...) indicates that you repeat the last element of the command as needed.</p> <p>For example, if the command syntax is <code>ethernet/2/1 [<parameter></code></p>

Table continues...

Convention	Description
	<value>] . . . , you enter ethernet/2/1 and as many parameter-value pairs as you need.
<i>Italic Text</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles that are not active links.
Plain Courier Text	Plain Courier text indicates command names, options, and text that you must enter. Plain Courier text also indicates command syntax and system output, for example, prompts and system messages. Examples: <ul style="list-style-type: none"> • show ip route • Error: Invalid command syntax [Failed][2013-03-22 13:37:03.303 -04:00]
Separator (>)	A greater than sign (>) shows separation in menu paths. For example, in the Navigation tree, expand the Configuration > Edit folders.
Vertical Line ()	A vertical line () separates choices for command keywords and arguments. Enter only one choice. Do not type the vertical line when you enter the command. For example, if the command syntax is access-policy by-mac action { allow deny }, you enter either access-policy by-mac action allow or access-policy by-mac action deny, but not both.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware/software compatibility matrices](#) for Campus and Edge products

[Supported transceivers and cables](#) for Data Center products

[Other resources](#), like white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Getting Help

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

[Extreme Portal](#)

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

[The Hub](#)

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

[Call GTAC](#)

For immediate support: (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

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

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions 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

Subscribe to Service Notifications

You can subscribe to email notifications for product and software release announcements, Vulnerability Notices, and Service Notifications.

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form (all fields are required).
3. Select the products for which you would like to receive notifications.

 **Note:**

You can modify your product selections or unsubscribe at any time.

4. Select **Submit**.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the 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.

Chapter 2: New in this Release

The following sections describe what is new in VOSS 8.1.50.

Important:

The following platforms support VOSS 8.1.50:

- XA1400 Series

VOSS 8.1.50

IPsec NAT-T Support

Internet Protocol Security (IPsec) Network Address Translation-Traversal (NAT-T) support allows IPsec tunnel traffic to travel through a NAT router. In this release, a minimum of two XA1400 Series devices are required, one on each side of the IPsec tunnel. You must configure one side of the IPsec tunnel as an IPsec Responder device, and the other side of the IPsec tunnel as an IPsec Initiator device. The IPsec packets use Internet Key Exchange (IKE) protocol and are encapsulated in a User Datagram Protocol (UDP) wrapper, allowing IPsec packets to be sent and received by devices behind a NAT router. Port Address Translation (PAT) is also supported when multiple devices behind a NAT router are connecting to a single aggregator device.

For more information, see [IPsec NAT-T](#) on page 11.

Fabric Extend Tunnel VLAN Support

Fabric Extend (FE) enables the extension of Fabric Connect networking over Layer 2 or Layer 3 core IP networks. New in this release you can configure a VLAN IP interface as the FE tunnel source IP address on an XA1400 Series device. You must configure the VLAN in the same VRF as the ISIS tunnel source IP address. In VOSS Release 8.1 and earlier, the FE tunnel source IP address is limited to a Brouter port or a CLIP IP interface only.

For more information, see [Fabric Extend Considerations](#) on page 22.

Fabric Extend Tunnel Unique MTU Support

Fabric Extend (FE) enables the extension of Fabric Connect networking over Layer 2 or Layer 3 core IP networks. New in this release you can configure a unique MTU value for each FE tunnel on an XA1400 Series device. The IS-IS logical interface mode command `logical-interface isis <isis ID> dest-ip <A.B.C.D>` has a new `mtu` parameter with a supported MTU range of 750 to 9000 and default of 1950. Default MTU is applicable for both ingress and egress traffic if no tunnel MTU is configured. Also, the `ip-tunnel-source-address` command `mtu` parameter is

removed on XA1400 Series. In VOSS Release 8.1 and earlier, the MTU value is globally applied to all FE tunnels.

For more information, see [Fabric Extend Tunnel MTU](#) on page 39

Ingress Rate Limiting on Front Panel Ports

Ingress rate limiting can control the rate of traffic received on a front panel port. New in this release you can configure a rate limit for each front panel port on an XA1400 Series device. When the port receives more traffic than the rate limit value, the excess traffic is dropped. You can use the rate limiting feature to control the total ingress traffic forwarded by the switch. You can configure rate limiting with the command `qos if-rate-limiting [port {slot/port[/sub-port]} [-slot/port[/subport]] [, ...]] rate <1000-10000000>`.

Also new in this release, the egress queue shapers can dynamically adjust to the auto-negotiation link speed of the front panel ports. In VOSS Release 8.1 and earlier, the shaper value was fixed at the maximum speed of the port.

For more information, see [Port-Rate Limiting, Policing, and Shaping](#) on page 51.

Pause Frame IEEE 802.3x Flow Control Support

Pause Frames are a layer 2 flow control mechanism to provide congestion relief on full-duplex interfaces. Pause frame flow control can temporarily stop the transmission of data on Ethernet links to achieve dropleless flow in network congestion scenarios. A device can generate pause frames if transmit (TX) flow control is enabled. A device can accept a pause frame request if receive (RX) flow control is enabled. RX flow control is enabled on all front panel ports by default. You can configure RX flow control with the global configuration mode command `boot config flags flow-control-mode`. You can enable TX flow control with the interface configuration command `tx-flow-control-enable`.

For more information, see [IEEE 802.3X Pause Frame Transmit](#) on page 53.

IPsec NAT-T

IP security (IPsec) Network Address Translation Traversal (NAT-T) allows IPsec tunnel traffic through Network Address Translation (NAT) or Port Address Translation (PAT) points in the network. To enable IPsec NAT-T connectivity, you must deploy and configure an IPsec NAT-T supported IPsec logical interface on each side of the IPsec tunnel.

The following terms are specific to the IPsec NAT-T feature:

- IPsec Initiator - The initiator mode of an IPsec tunnel initiates IKE negotiation and rekeying. A supported IPsec logical interface configured as an IPsec Initiator initiates IKE message exchange and IPsec security association (SA) establishment, rekeys IKE, and initiates IPsec SAs. Configure the IPsec Initiator device on the branch side of the tunnel.



Note:

IPsec Initiator is the default mode of the logical interface IPsec tunnel.

- IPsec Responder - The responder-only mode of an IPsec tunnel responds to IKE negotiation and rekeying and only responds to IPsec negotiations from an IPsec Initiator. A VOSS device

with a logical interface configured as an IPsec Responder does not initiate IKE message exchange and IPsec SA establishment, and does not rekey IKE and IPsec SAs. Configure the IPsec Responder device on the aggregator side of the tunnel.

- IPsec Remote NAT IP – The remote-nat-ip is the public IP address of the NAT router connected to the IPsec Responder device. Only configure an IPsec Initiator with the IPsec Remote NAT IP address when NAT routers are present on both sides of the IPsec tunnel.

*** Note:**

Only configure the logical interface of the IPsec Initiator device with the Remote NAT IP address of the responder.

The following diagram illustrates an example of an IPsec NAT-T configuration with both sides of the connection behind NAT:

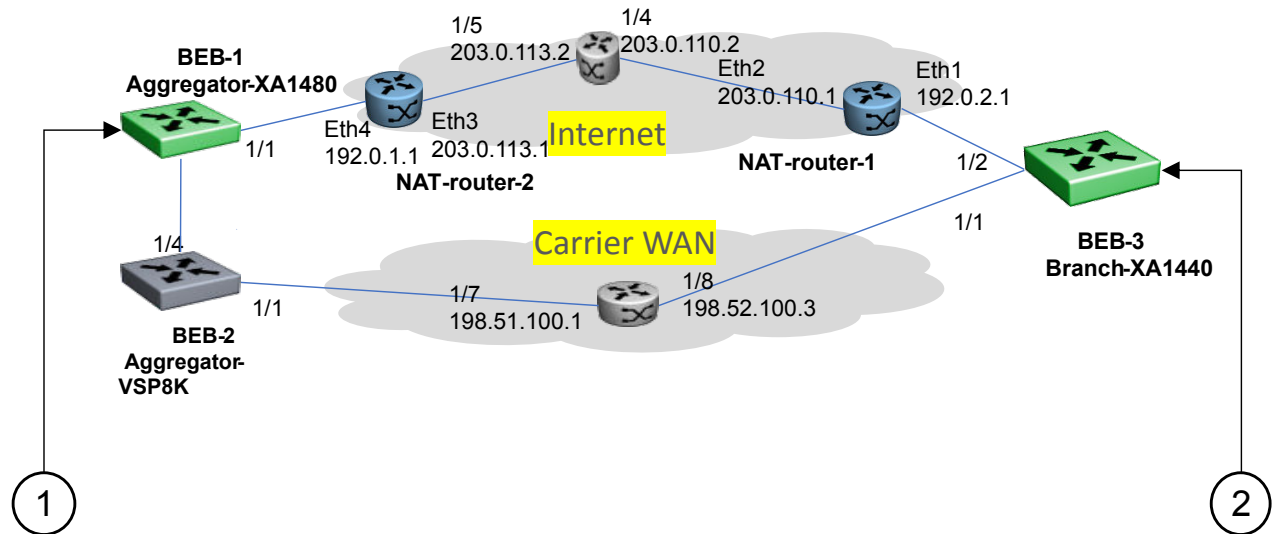


Figure 1: IPsec NAT-T with NAT router on both sides

1. BEB-1 Aggregator side IPsec Responder device configuration example:

```
logical-intf isis 2 dest-ip 192.0.2.2
mtu 1300
name "Tunnel-to-BEB3"
auth-key <key value>
ipsec responder-only
ipsec
```

2. BEB-3 Branch side IPsec Initiator device configuration example:

```
logical-intf isis 2 dest-ip 192.0.1.3
mtu 1300
name "Tunnel-to-BEB1"
auth-key <key value>
```

```
ipsec remote-nat-ip 203.0.113.1
ipsec
```

The following diagram illustrates an example of an IPsec NAT-T configuration with only 1 side of the connection behind NAT:

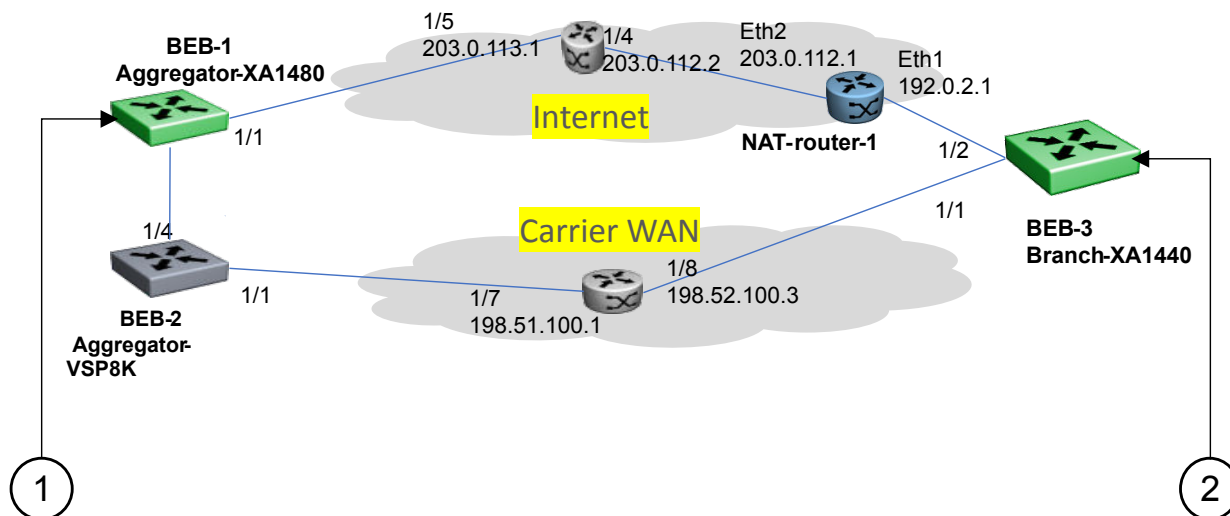


Figure 2: IPsec NAT-T with NAT router on one side

1. BEB-1 Aggregator side IPsec Responder device configuration example:

```
logical-intf isis 2 dest-ip 192.0.2.2
mtu 1300
name "Tunnel-to-BEB3"
auth-key <key value>
ipsec responder-only
ipsec
```

2. BEB-3 Branch side IPsec Initiator device configuration example:

```
logical-intf isis 2 dest-ip 192.0.1.2
mtu 1300
name "Tunnel-to-BEB1"
auth-key <key value>
ipsec
```

IPsec NAT-T Considerations

The following considerations apply to IPsec NAT-T:

- You must configure one side of the IPsec NAT-T tunnel as an IPsec responder. If IPsec is configured on the IPsec Initiator device and subsequently configured on the IPsec Responder device, IPsec must be restarted on the Initiator device. If IPsec is not restarted, it can take approximately 3 minutes for the adjacency to open.
- You must configure the aggregator device as the IPsec Responder device, and configure the branch device as the IPsec Initiator device.

- Among all the IPsec responders, the system uses the lowest configured maximum transmission unit (MTU) value of any responder IPsec tunnel as the MTU value for all IPsec responder-only tunnels. The system uses the lowest configured IPsec tunnel MTU value regardless of manually configured MTU tunnel values, and higher MTU values might be visible in the IPsec information for the logical interface. For non-responder IPsec tunnels or VXLAN tunnels, the configured and visible MTU value for the tunnel is used for fragmentation and reassembly.
- If both the IPsec Initiator device and the IPsec Responder device are behind NAT, you must configure the IPsec Initiator device with the public IP address of the NAT router connected to the IPsec Responder device.
- You must add route table entries on the IPsec Responder device with the public IP address and private IP address of the remote NAT for the IPsec Initiator device. A configured route table is required for IPsec NAT-T Fabric Extend (FE) connectivity.
- You must add a route table entry on the IPsec Initiator device with the public IPsec Remote NAT IP address for the IPsec Responder device. A configured route table is required for IPsec NAT-T Fabric Extend (FE) connectivity.

Configure the IPsec Initiator with the IPsec Responder Remote NAT IP Address

About this task

If both the Responder device and the Initiator device are behind Network Address Translation (NAT), you must configure the IPsec Initiator device with the public IP address of the NAT router connected to the IPsec Responder device.

- * **Note:**
Only perform this procedure on the IPsec Initiator device.

Procedure

1. Enter Layer 3 Logical IS-IS Interface Configuration mode:

```
enable
configure terminal
logical-intf isis <1-255> dest-ip {A.B.C.D} [name WORD<1-64>] [mtu
<mtu_value>]
```
2. Configure the public IP address of the NAT router connected to the IPsec Responder device:

```
ipsec remote-nat-ip {A.B.C.D}
```

Variable Definitions

The following table defines parameters for the `logical-intf isis` command.

Variable	Value
isis <1-255>	Specifies the Intermediate-System-to-Intermediate-System (IS-IS) logical interface ID.
dest-ip {A.B.C.D}	Specifies the destination IP address for the logical interface.
name WORD<1–64>	Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.
mtu <mtu_value> * Note: Exception: only supported on XA1400 Series.	Specifies the Maximum Transmission Unit (MTU) size for each packet. Different hardware platforms support different MTU ranges. Use the CLI Help to see the available range for your switch. The default value is 1950.

The following table defines parameters for the `ipsec remote-nat-ip` command.

Variable	Value
{A.B.C.D}	Specifies the public IP address of the NAT router connected to the Responder device in an IPsec Network Address Translation Traversal (NAT-T) connection.

Configure an IPsec NAT-T Responder

About this task

One side of an IPsec Network Address Translation Traversal (NAT-T) connection must be a responder device and the other side must be the Initiator device. By default, both sides of an IPsec NAT-T connection are Initiators. Use the following procedure to configure one side of an IPsec NAT-T connection as a Responder device.

Procedure

1. Enter Layer 3 Logical IS-IS Interface Configuration mode:

```
enable
```

```
configure terminal
```

```
logical-intf isis <1-255> dest-ip {A.B.C.D} [name WORD<1-64>] [mtu <mtu_value>]
```

2. Configure an IPsec Responder:

```
ipsec responder-only
```

Variable Definitions

The following table defines parameters for the `logical-intf isis` command.

Variable	Value
isis <1-255>	Specifies the Intermediate-System-to-Intermediate-System (IS-IS) logical interface ID.
dest-ip {A.B.C.D}	Specifies the destination IP address for the logical interface.
name WORD<1–64>	Specifies the administratively assigned name of this logical interface, which can be up to 64 characters.
mtu <mtu_value>	Specifies the Maximum Transmission Unit (MTU) size for each packet. Different hardware platforms support different MTU ranges. Use the CLI Help to see the available range for your switch. The default value is 1950.
<p>* Note:</p> <p>Exception: only supported on XA1400 Series.</p>	

Display IS-IS Logical Interfaces

Use the following procedure to display the Intermediate-System-to-Intermediate-System (IS-IS) logical interfaces configured on the switch.

Procedure

Display the IS-IS logical interfaces:

```
show isis logical-interface [name | ipsec | shaper | mtu]
```

Examples

Example of a Layer 2 Core

```
Switch:1# show isis logical-interface
=====
ISIS Logical Interfaces
=====
IFIDX  NAME      ENCAP      L2_INFO      TUNNEL      L3_TUNNEL_NEXT_HOP_INFO
      TYPE      PORT/MLT  VIDS (PRIMARY)  DEST-IP      PORT/MLT  VLAN  VRF
-----
1      --      L2-P2P-VID  Port2/40  101,201 (101)  --      --      --      --
2      --      L2-P2P-VID  Port1/3   102,202 (102)  --      --      --      --
-----
2 out of 2 Total Num of Logical ISIS interfaces
=====
```

Example of a Layer 3 Core

```
Switch:1# show isis logical-interface
=====
ISIS Logical Interfaces
=====
IFIDX  NAME      ENCAP      L2_INFO      TUNNEL      L3_TUNNEL_NEXT_HOP_INFO
      TYPE      PORT/MLT  VIDS (PRIMARY)  DEST-IP      PORT/MLT  VLAN  VRF
-----
1      SPBoIP_T1  IP      --      --      41.41.41.41  MLT10  2    vrf24
2      SPBoIP_T2  IP      --      --      42.42.42.42  MLT10  2    vrf24
3      SPBoIP_4K5 IP      --      --      187.187.187.187  MLT10  2    vrf24
-----
3 out of 3 Total Num of Logical ISIS interfaces
=====
```


Example showing the full IS-IS logical interface name

The command `show isis logical-interface` truncates the IS-IS logical interface name to the first 16 characters. To view the entire name (up to a maximum of 64 characters), use the command `show isis logical-interface name`.

```
Switch:1# show isis logical-interface name
=====
ISIS Logical Interface name
=====
ID      NAME
-----
1       SPBoIP_T1
2       SPBoIP_T2
3       SPBoIP_4K5
6       This_Is_A_50_Character_ISIS_Logical_Interface_Name
-----
4 out of 4 Total Num of Logical ISIS interfaces
-----
```

Example showing the authentication key

Display the IS-IS logical interface ID and IPsec authentication key. This command only displays IPsec-enabled interfaces with authentication keys configured.

* Note:

The Authentication-Key is obscured and not visible in plain text output.

```
Switch:1>show isis logical-interface ipsec
=====
ISIS Logical Interface IPsec
=====
ID      Authentication-Key      Responder-Only  Remote NAT IP
-----
1       *****                True            -
-----
1 out of 1 Total Num of Logical ISIS interfaces
-----
```

Example showing the IS-IS egress shaping rate values

Display the IS-IS logical interface egress shaping rate values. This command only displays interfaces with egress shaping rates configured.

```
Switch:1>show isis logical-interface shaper
=====
ISIS Logical Interface Egress Shaping Rate
=====
ID      NAME                                service-rate (Mbps)
-----
1       remotel                             135
2       remote2                             120
3       remote3                             178
-----
```

```
3 out of 3 Total Num of Logical ISIS interfaces
```

Example showing the IS-IS logical interfaces mtu values

This command displays the Maximum Transmission Unit (MTU) size for each logical interface.

```
Switch:1>show isis logical-interface mtu
```

```
=====
ISIS Logical Interface Mtu
=====
ID      NAME                MTU
-----
1       SPBoIP_T1           751
2       SPBoIP_T2           1000
3       SPBoIP_4K5          1950
=====
3 out of 3 Total Num of Logical ISIS interfaces
=====
```

ipsec remote-nat-ip

Configures the Network Address Translation Traversal (NAT-T) Responder device public IP address.

Syntax

- `ipsec remote-nat-ip {A.B.C.D}`
- `no ipsec remote-nat-ip`

Command Parameters

{A.B.C.D} Specifies the IP address of the Responder device in an IPsec NAT-T connection.

Command Mode

Logical Interface Configuration

ipsec responder-only

Configure one side of an IPsec Network Address Translation Traversal (NAT-T) connection as a Responder device. By default, both sides of an IPsec NAT-T connection are initiators.

Syntax

- `ipsec responder-only`
- `no ipsec responder-only`

Default

The default is initiator.

Command Mode

Logical Interface Configuration

show io

Shows IO information.

Syntax

- `show io cpu-cosq-counters`
- `show io filter-tables`
- `show io ipsec logs`
- `show io ipsec stats`
- `show io l2-tables`
- `show io l3-tables`
- `show io logical-intf-ipsec`
- `show io logical-intf-tables`
- `show io nic-counters`
- `show io performance-vcpu`
- `show io spb-tables`
- `show io tunnel-stats`

Command Parameters

cpu-cosq-counters	Shows the CPU cosq counters.
filter-tables	Shows the filter tables.
ipsec logs	Shows the ipsec logs.
ipsec stats	Shows the ipsec statistics.
l2-tables	Shows the Layer 2 tables.
l3-tables	Shows the Layer 3 tables.
logical-intf-ipsec	Shows the logical interface ipsec status.
logical-intf-tables	Shows the logical interface tables.
nic-counters	Shows the network interface card counters.
performance-vcpu	Shows the CPU performance.
spb-tables	Shows the Shortest Path Bridging tables.
tunnel-stats	Shows the tunnel statistics.

Default

None

Command Mode

User EXEC

Configure IPsec NAT-T

About this task

By default, both sides of an IPsec connection are Initiator devices. IPsec Network Address Translation Traversal (NAT-T) connections require that one side of the connection is a Responder device and the other side of the connection is an Initiator device.

If the Responder device and the Initiator device are both behind NAT, the IPsec NAT-T Initiator device requires the public IP address of the Responder device.

Procedure

Perform the following steps to configure one side of an IPsec NAT-T connection as a Responder device:

1. In the navigation pane, expand **Configuration > IS-IS**.
2. Select **IS-IS**.
3. Select **Logical Interfaces**.
4. Select **Insert**.
5. Select **IpsecNatConfigResponderOnly**.

If required, perform the following steps on the IPsec NAT-T Initiator device to configure the public IP address of the Responder device:





6. In the navigation pane, expand **Configuration > IS-IS**.
7. Select **IS-IS**.
8. Select **Logical Interfaces**.
9. Select **Insert**.
10. For **IpsecNatConfigRemoteNatIPAddr**, enter the public IP address of the Responder device.

Logical Interfaces Field Descriptions

Use the data in the following table to use the **Logical Interfaces** tab and the Insert Logical Interfaces dialog. The available fields in the dialog differ depending on the type of core you select: **layer 2** or **ip**.

Name	Description
Id	Specifies the index number that uniquely identifies this logical interface. This field appears only on the Insert Logical Interfaces dialog.
IfIndex	Specifies the index number that uniquely identifies this logical interface. This field is read-only. This field appears only on the Logical Interfaces tab.
Name	Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.
Type * Note: Exception: type Layer 2 is not supported on XA1400 Series.	Specifies the type of logical interface to create: <ul style="list-style-type: none"> Specify layer 2 for a Layer 2 core network that the tunnel will be traversing. Specify ip for a Layer 3 core network that the tunnel will be traversing.
DestIPAddr	Specifies the destination IP address for the IP-type logical interface.
DestIfIndex * Note: Exception: not supported on XA1400 Series.	Specifies the physical port or MultiLink Trunking (MLT) that the Layer 2 logical interface is connected to.
Vids * Note: Exception: not supported on XA1400 Series.	Specifies the list of VLANs that are associated with this logical interface.
PrimaryVid * Note: Exception: not supported on XA1400 Series.	Specifies the primary tunnel VLAN ID associated with this L2 Intermediate-System-to-Intermediate-System (IS-IS) logical interface.
CircIndex * Note: Exception: not supported on XA1400 Series.	Identifies the IS-IS circuit created under the logical interface. This field appears only on the Logical Interfaces tab.
NextHopVrf * Note: Exception: not supported on XA1400 Series.	Identifies the next-hop VRF name to reach the logical tunnel destination IP. This field appears only on the Logical Interfaces tab.
IpsecEnable * Note: Exception: only supported on XA1400 Series.	Specifies whether the logical interace should use IPsec.

Table continues...

Name	Description
AuthenticationKey  Note: Exception: only supported on XA1400 Series.	Specifies the authentication key of this logical interface, which can be up to 32 characters.
IpssecNatConfigResponderOnly  Note: Exception: only supported on XA1400 Series.	Specifies whether the device is a Responder device in an IPsec Network Address Translation Traversal (NAT-T) connection.
IpssecNatConfigRemoteNatIPAddr  Note: Exception: only supported on XA1400 Series.	Specifies the public IP address of the NAT router connected to the Responder device in an IPsec NAT-T connection.
ShapingRate  Note: Exception: only supported on XA1400 Series.	Specifies the value, in Mbps, of the Egress Tunnel Shaper applied to the logical interface.
Mtu	Specifies the Maximum Transmission Unit (MTU) size for each logical interface. The default MTU value is 1950.

Fabric Extend Considerations

Review the following restrictions, limitations, and behavioral characteristics that are associated with Fabric Extend.

 **Note:**

If your Fabric Extend configuration includes a VSP 4000/ONA combination, see [ONA considerations](#) on page 26 for more information.

- **Tunnel source IP**

Fabric Extend supports the tunnel source IP address using a brouter port interface, a CLIP IP, or a VLAN IP.

- **Extreme Fabric Orchestrator (EFO)**—The Fabric Extend view within Extreme Fabric Orchestrator (EFO) is not required, but it is highly recommended.
- **Tunnel failover time**—With IS-IS interface default values, tunnel failure detection can take up to 27 seconds. You can reduce the IS-IS interface hello timers to speed up logical link failure detection, but be careful to avoid link flapping due to values that are too low.

 **Note:**

If the number of IS-IS interfaces on a node is greater than 100, it is a good practice to set the hello timer not lower than 5 seconds.

- **ACL Filters over VXLAN**—IP filters configured to match IP header fields in the headers of VXLAN encapsulated packets, work only when the switch acts as a transit router and does not participate in the initiation or termination of VXLAN traffic.
- **VLACP**—VLACP is not supported over logical IS-IS interfaces.
- **CFM CCM**—CFM Continuity Check Messages are not supported over logical IS-IS interfaces.
- **CFM traceroute and tracemroute**—If CFM packets transit over a layer 3 tunnel (that is the CFM packets ingress a Fabric Extend layer 3 core tunnel and egress through another layer 3 core tunnel), the transit SPBM nodes do not display as intermediate hops in the output for CFM `12 traceroute` and `12 tracemroute`.

This is because the CFM packets are encapsulated in the outer layer 3 header as part of VXLAN encapsulation, and the transit SPBM nodes cannot look into the payload of the VXLAN packet and send a copy of the CFM packet to local CPU for processing.

- **CFM L2 ping**—CFM L2 ping to MCoSPB source mac is not supported and may fail if they are reachable via Fabric Extend tunnel.
- **MACsec**—Switch-based MAC Security (MACsec) encryption is Layer 2 so it cannot be used with Fabric Extend IP, which is Layer 3.
- **MTU minimum in Layer 2 Pseudowire core networks**—Service provider Layer 2 connections must be at least 1544 bytes. In this type of deployment the tunnels are point-to-point VLAN connections that do not require VXLAN encapsulation. The default MTU value is 1950.
- **Logical IS-IS interfaces**—Layer 2 core and Layer 3 core logical IS-IS interfaces are not supported on the same switch at the same time.
- **Fragmentation/reassembly**—There is no fragmentation/reassembly support in Layer 2 core solutions.

If a tunnel was initially UP between a VSP 4000 and a VOSS switch with MTU 1950 and then the VSP 4000 was later configured for fragmentation, the following behavior occurs:

- If the ONA MTU is less than 1594, the tunnel to the VOSS switch will go DOWN.
- If the ONA MTU is 1594 and above, the tunnel will stay UP, but any fragmented packets received from the VSP 4000 will be lost at the VOSS switch site.

Fragmented traffic can only be sent with an XA1400 Series or VSP 4000/ONA combination on both ends with the same MTU configured on each end.

- **RFC4963 and RFC4459 considerations:**

The ONA 1101GT provides for the IP MTU of the Network port to be reduced from the default setting of 1950 bytes to 1500 bytes or lower. The MTU reduction feature with Fabric Extend is provided to facilitate the connection of two Fabric Connect networks over an IP network with any MTU without requiring end stations on the networks to reduce their MTU. The ONA 1101GT with the IP MTU of the network port set to 1500 bytes will fragment Fabric Extend VXLAN tunnel packets exceeding 1500 bytes. The ONA 1101GT will also reassemble fragmented Fabric Extend VXLAN tunnel packets at the tunnel termination point. The IP fragmentation and reassembly RFC 791 describes the procedure for IP fragmentation, and transmission and reassembly of datagrams and RFC4963 and RFC4459 detail limitations and

network design considerations when using fragmentation to avoid out of order packets and performance degradation.

Factors that can impact performance are —

- The link speed per VXLAN IP address should be slower than 1G to avoid reassembly context exhaustion.
- ECMP and link aggregation algorithms in the IP core should be configured not to use UDP port hashing that could send IP fragments after the first fragment on different paths causing out of order packets. This is due to the fact that subsequent fragments do not have UDP port information.

 **Important:**

Different MTU sizes on each end can result in traffic drops.

- **Layer 2 logical IS-IS interfaces**—Layer 2 logical IS-IS interfaces are created using VLANs. Different Layer 2 network Service Providers can share the same VLAN as long as they use different ports or MLT IDs.

 **Note:**

Exception: Layer 2 logical IS-IS interfaces are not supported on XA1400 Series.

- **MTU minimum in Layer 3 core networks**—Service provider IP connections must be at least 1594 bytes to establish IS-IS adjacency over FE tunnels. The 1594 bytes includes the actual maximum frame size with MAC-in-MAC and VXLAN headers. If this required MTU size is not available, a log message reports that the IS-IS adjacency was not established. MTU cannot be auto-discovered over an IP tunnel so the tunnel MTU will not be automatically set. The default MTU value is 1950.

If the maximum MTU size has to be fewer than 1594 bytes, then you require fragmentation and reassembly of packets. The XA1400 Series and VSP 4000/ONA combination supports fragmentation and reassembly, but you must have an XA1400 Series or VSP 4000s with ONAs at BOTH ends of the IP WAN connection.

- **IP Shortcuts**—The tunnel destination IP cannot be reachable through an IP Shortcuts route.

 **Important:**

If you enable IP Shortcuts and you are using the GRT as the tunnel source VRF, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

If you enable IP Shortcuts and you are using a VRF as the tunnel source VRF, this is not an issue.

- **Layer 3 over Layer 2 Limitation**—The VOSS switches require a single next hop (default gateway) for all tunnels.
 - Over a layer 3 core network, on a given outgoing port or MLT, there is no issue as the one router next hop can support multiple VXLAN tunnels to one or more remote sites.

- For layer 3 tunneling over a layer 2 core, the VOSS switch without any specific configuration supports only one Fabric Extend tunnel to one remote site. The workaround for this single next hop issue is to create an additional VRF, VLAN, and loopback interface.

For a configuration example of this workaround, see *Shortest Path Bridging (802.1aq) Technical Configuration Guide*.

 **Note:**

This limitation does not apply to VSP 4000.

- You cannot establish a Virtual IST (vIST) session over a logical IS-IS interface. IST hellos cannot be processed or sent over a logical IS-IS interface if that is the only interface to reach BEBs in vIST pairs.

Assume that vIST is established over a regular NNI interface and the NNI interface goes down. If the vIST pairs are reachable through a logical IS-IS interface, then the vIST session goes down in up to 240 seconds (based on the IST hold down timer). During this time, the error message `IST packets cannot be sent over Fabric Extend tunnels, vist session may go down` is logged.

 **Caution:**

Expect traffic loss when the vIST session is down or when the error message is being logged.

- **Port Mirroring Resources:**

Port mirroring resources are limited to four ports simultaneously (where each mirroring direction counts as one). For example, if two mirroring ports are designated to mirror both ingress and egress traffic then all four mirroring ports are consumed.

Port mirroring shares these four resources with other applications such as port mirroring RSPAN, Fabric Extend, Application Telemetry, IPFIX, and ACL with mirror action. Each one of these applications consumes at least one port mirroring resource. (port mirroring RSPAN consumes two if you configure both Ingress and Egress modes.)

- **Important:**

- To enable any one of the above applications, you must have at least one free mirroring resource. If all four port mirroring resources are already in use, the switch displays a `Resource not available` error message when you try to enable the application.
- The VSP 8600 uses the four reserved resources for port mirroring and ACLs that have a mirroring action. For the other applications, this restriction does not apply because the VSP 8600 uses mirroring resources that do not come out of the four reserved port mirroring resources.

- **Fabric Extend over IPsec limitations**

- Fabric Extend over IPsec is only supported on XA1400 Series devices.
- Only pre-shared authentication key IPsec parameters are user configurable. Other, third-party solutions are not configurable.

- IKEv2 protocol key exchange only.
- IPsec support is only added for Fabric Extend tunnels.
- IPsec is not supported for regular layer 3 routed packets.

ONA considerations

Note:

Review the following restrictions, limitations, and behavioral characteristics that are associated with the ONA.

ONA Network port requirements

The following are **Network** port mandatory requirements for configuring Fabric Extend on the VSP 4000:

- The ONA Network port should not be part of any static/LACP MLT configurations.
- The ONA Network port should be part of a VLAN that belongs to the GRT.
- The ONA Network port that is configured on the switch cannot be tagged. It must be an Access port.

ONA Device port requirements

The following are **Device** port mandatory requirements for configuring Fabric Extend on the VSP 4000:

- The ONA Device port should not be part of any static/LACP MLT, VLAN, or brouter configurations.
- The ONA Device port should not be configured as an access port. It is automatically configured as a trunk port when the `ip-tunnel-source-address` command is configured.
- The ONA Device port has to be connected directly to the VSP 4000 node where the FE tunnels originate.

Layer 3 and Layer 2 ONA requirements

An ONA is required for Fabric Extend Layer 3 core solutions. An ONA is *not* required in Layer 2 core solutions because the tunnels are point-to-point VLAN connections, not VXLAN. Therefore, there is no need for an ONA to encapsulate a VXLAN header to SPB packets.

DHCP server

ONAs require access to a local DHCP server to automatically configure IP addresses. Configure an untagged ONA management VLAN to where the ONA is connected with its network side interface. If DHCP is used, a DHCP relay configuration needs to be added to the ONA network side port in order for the ONA to get an IP address assigned from a DHCP server. Alternatively, you can manually configure its IP address and other required settings with the ONA Manual Configuration menu.

IP tunnel source address

Before the ONA can get an IP tunnel source address from the VSP 4000, the following steps must be taken:

- Connect the Device and Network ports on the ONA to the VSP 4000.
- Make sure that the ONA is connected to a DHCP server. If a DHCP server is unavailable, statically configure an IP tunnel source address on the ONA.
- Create a Management VLAN on the ONA that includes the Network port.
- Designate the Device port for the IP tunnel source address in the configuration file.

The syntax for the IP tunnel source address is: `ip-tunnel-source-address <A.B.C.D> port <slot/port> [mtu <mtu_value>] [vrf WORD<1-16>]`.

Automatic routing of VXLAN packets on the VSP 4000

If you configure an IP tunnel source address in a VRF instead of a GRT, then the VSP 4000 automatically routes VXLAN packets from the ONA network port into the VRF configured as part of the IP tunnel source. Although the ONA network port is a part of the management VLAN that is in the GRT, for VXLAN encapsulated packets, the VSP 4000 automatically routes the packets into the VRF in which the tunnel source IP address is configured. This is done using a filter rule that the VSP 4000 software automatically sets up that filters based on whether the incoming port is equal to the ONA network port and the packet has a VXLAN header.

The Management VLAN on the VSP 4000 that is used to communicate with the ONA must always be in a GRT and must not be a part of the IP tunnel source VRF.

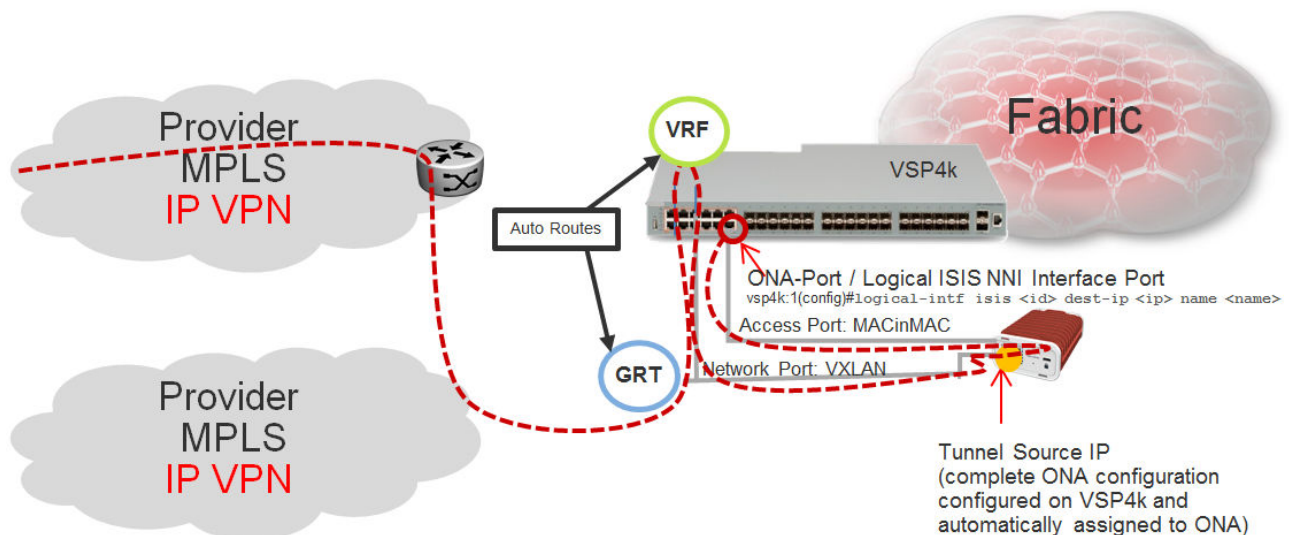


Figure 3: Autorouting between GRT and VRF

ONA Gateway

The ONA gateway has to be a local IP address on the ONA Management VLAN. The ONA gateway IP address must be the same as the local IP address of the VSP 4000 connected to the ONA.

*** Note:**

Extreme does not support ONA gateway IP addresses that are not local to the VSP 4000. For example, you cannot use a VRRP IP address configured in a switch cluster for the ONA gateway.

Maximum MTU

The ONA supports a maximum transmission unit (MTU) size of 1950 bytes. For the VSP 4000 to work with a switch that supports Fabric Extend natively, the MTU size must be left at the default setting of 1950. If the core network does not support jumbo frames, the VSP 4000 with ONA must be used on all sites.

Fragmentation and reassembly

If the maximum MTU size has to be fewer than 1594 bytes, then you require fragmentation and reassembly of packets. The VSP 4000s with ONAs support fragmentation and reassembly, but you must have VSP 4000s with ONAs at BOTH ends of the IP WAN connection.

QoS priority queues

The ONA 1101GT implements both Layer 2 and Layer 3 QoS. Specifically, it implements IEEE 802.1Q VLAN TCI PCP (Priority Code Point) and IETF IPv4 DSCP (Differentiated Services Code Point). These are implemented in hardware with the limitation that there are four Weighted Random Early Detection (WRED) priority queues, numbered 4 (highest) to 7 (lowest). The following tables show the mappings from the PCP and DSCP values in the packet to the priority queue.

The hardware puts each packet in 1 of the 4 HW queues in the following order:

1. If a packet is a tagged VLAN packet, the PCP field determines the priority queue. (Ethertypes 0x8100 and 0x88a8 identify tagged VLAN packets.)
2. If the packet is an IPv4 packet, the DSCP field determines the priority queue.
3. Use the highest priority queue (4).

The HW QoS is always enabled, and the CP to priority queue mappings are static.

The following table defines the 3 bit VLAN PCP value to queue number mapping. The queues are numbered 4..7 with 4 being the highest priority and 7 the lowest priority.

Table 3: VLAN PCP to queue mapping

VLAN PCP	Queue Number
0	7
1	7
2	6
3	6
4	5
5	5
6	4
7	4

The following table defines the 6 bit IPv4 DSCP value to queue number mapping. The queues are numbered 4..7 with 4 being the highest priority and 7 the lowest.

Table 4: IPv4 DSCP to queue mapping

IPv4 DSCP	VLAN PCP	Queue Number
0	1	7
1	1	7
2	1	7
3	1	7
4	1	7
5	1	7
6	1	7
7	1	7
8	2	6
9	1	7
10	2	6
11	1	7
12	2	6
13	1	7
14	2	6
15	1	7
16	3	6
17	1	7
18	3	6
19	1	7
20	3	6
21	1	7
22	3	6
23	1	7
24	4	5
25	1	7
26	4	5
27	4	5
28	4	5
29	1	7
30	4	5
31	1	7
32	5	5
33	1	7

Table continues...

IPv4 DSCP	VLAN PCP	Queue Number
34	5	5
35	5	5
36	5	5
37	1	7
38	5	5
39	1	7
40	6	4
41	5	5
42	1	7
43	1	7
44	1	7
45	1	7
46	6	4
47	6	4
48	7	4
49	1	7
50	1	7
51	1	7
52	1	7
53	1	7
54	1	7
55	1	7
56	7	4
57	1	7
58	1	7
59	1	7
60	1	7
61	1	7
62	1	7
63	1	7

Configure Fabric Extend

Use the following procedure to configure Fabric Extend (FE) between a Main office to a Branch office. This is a typical deployment. However, if your deployment creates tunnels between two

switches that support Fabric Extend natively, then repeat those steps and ignore the steps for switches that require an ONA.

*** Note:**

VRF is an optional parameter. If a VRF is not configured, then FE uses the GRT.

Before you begin

The tunnel source IP address can be a brouter port IP, a CLIP IP, or a VLAN IP.

If using the tunnel originating address on the **GRT**, Fabric Extend has the following requirements:

- The tunnel source IP address must be on the GRT, not on a VRF.

*** Note:**

A Best Practice is to use separate IP addresses for the SPBM IP Shortcuts **ip-source-address** command and the Fabric Extend **ip-tunnel-source-address** command. However, if you want these IP addresses to be the same, you **MUST** exclude the **ip-source-address** address with an IS-IS accept policy. You cannot use the redistribute command with a route map exclusion.

Specify a CLIP interface to use as the source address for SPBM IP shortcuts.

- If IP Shortcuts is enabled, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

If you are using the tunnel originating address on a **VRF**, Fabric Extend has the following requirements:

- Configure a CLIP and tunnel source IP address on the VRF.
- Remote management of the VSP 4000 is only possible after establishing IP Shortcut over IS-IS. (Alternatively, you can enable GRT-VRF redistribution locally.)

About this task

Configuring Fabric Extend consists of two primary tasks: configuring the tunnel source address and configuring the logical interface. These tasks must be completed on both ends of the tunnel.

Note that the VSP 4000 source address command is different than other platforms. Also note that the logical interface commands are different between Layer 2 and Layer 3 networks.

Procedure

The following steps are for platforms that support FE natively:

1. Enter IS-IS Router Configuration mode:

```
enable
configure terminal
router isis
```

2. Configure the IP tunnel source address:

```
ip-tunnel-source-address <A.B.C.D> [vrf WORD<1-16>]
```

3. Enter Global Configuration mode:

```
exit
```

4. Use one of the following commands to create a logical IS-IS interface:

- In a network with a Layer 3 Core, enter `logical-intf isis <1-255> dest-ip <A.B.C.D> [name WORD<1-64>] [mtu <750-9000>]`
- In a network with a Layer 2 Core, enter `logical-intf isis <1-255> vid <list of vids> primary-vid <2-4059> port <slot/port> mlt <mltId> [name WORD<1-64>] [mtu <750-9000>]`

*** Note:**

The primary VLAN ID (**primary-vid**) must be one of the VLANs in the **vid <list of vids>**.

The following steps are for platforms that require an ONA to support FE:

*** Note:**

The interface VLAN connecting to the ONA network port is always in the GRT and the member port that the VLAN is part of is always an access port.

5. Enter IS-IS Router Configuration mode:

```
enable  
configure terminal  
router isis
```

6. Configure the IP tunnel source address on the port that connects to the Device side of the ONA:

```
ip-tunnel-source-address <A.B.C.D> port <slot/port> [mtu  
<mtu_value>] [vrf WORD<1-16>]
```

7. Exit back into Global Configuration mode:

```
exit
```

8. Use one of the following commands to create a logical IS-IS interface:

- In a network with a Layer 3 Core, enter:
`logical-intf isis <1-255> dest-ip <A.B.C.D> [name WORD<1-64>]`
- In a network with a Layer 2 Core, enter:
`logical-intf isis <1-255> vid <list of vids> primary-vid <2-4059>
port <slot/port> mlt <mltId> [name WORD<1-64>]`

*** Note:**

The primary VLAN ID (**primary-vid**) must be one of the VLANs in the **vid <list of vids>**.

Variable Definitions

Use the data in the following tables to use the `ip-tunnel-source-address` command.

To delete an IS-IS IP tunnel source address, use the `no ip-tunnel-source-address` option.

*** Note:**

The `port` parameter is for the VSP 4000 only.

*** Note:**

Variable	Value
<A.B.C.D>	Specifies the IS-IS IPv4 tunnel source address, which can be a brouter interface IP, a CLIP IP, or a VLAN IP.
port <slot/port>	Specifies the port that is connected to the ONA's Device port.
vrf WORD<1–16>	Specifies the VRF name associated with the IP tunnel.
mtu <mtu_value>	Specifies the Maximum Transmission Unit (MTU) size for each packet. Different hardware platforms support different MTU ranges. Use the CLI Help to see the available range for your switch. This parameter only applies to an ONA configuration.

Use the data in one of the following tables to use the `logical-intf isis` command, depending on whether you have a Layer 2 or Layer 3 core.

To delete a logical IS-IS interface, use the `no logical-intf isis` option.

Table 5: Layer 2 core

Variable	Value
<1–255>	Specifies the index number that uniquely identifies this logical interface.
port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}	Specifies the physical port that the logical interface is connected to in a Layer 2 network.
vid <list of vids>	Specifies the list of VLANs that are associated with this logical interface.
primary-vid <2–4059>	Specifies the primary tunnel VLAN ID associated with this Layer 2 IS-IS logical interface.
mlt <mltid>	Specifies the MLT ID that the logical interface is connected to in a Layer 2 network.
name WORD<1–64>	Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.

Table continues...



Variable	Value
mtu<750-9000>  Note: Exception: only supported on XA1400 Series.	Specifies the Maximum Transmission Unit (MTU) size of each packet. The default MTU value is 1950.

Table 6: Layer 3 core

Variable	Value
<1-255>	Specifies the index number that uniquely identifies this logical interface.
dest-ip <A.B.C.D>	Specifies the tunnel destination IP address of the remote BEB.
name WORD<1-64>	Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.
mtu<750-9000>  Note: Exception: only supported on XA1400 Series.	Specifies the Maximum Transmission Unit (MTU) size of each packet. The default MTU value is 1950.

Configure Fabric Extend Over IPsec

Use the following procedure to configure Fabric Extend (FE) over IPsec.

Before you begin

The tunnel source IP address can be a brouter port IP, a CLIP IP, or a VLAN IP.

About this task

Configuring Fabric Extend over IPsec consists of two primary tasks: configuring the tunnel source address and configuring the logical interface. These tasks must be completed on both ends of the tunnel.

Procedure

Switch A Steps

- Enter IS-IS Router Configuration mode:


```
enable
configure terminal
router isis
```
- Configure the IP tunnel source address:


```
ip-tunnel-source-address <A.B.C.D> [vrf WORD<1-16>]
```
- Enter Global Configuration mode:

```
exit
```

4. Use one of the following commands to create a logical IS-IS interface:

- In a network with a Layer 3 Core, enter `logical-intf isis <1-255> dest-ip <A.B.C.D> [name WORD<1-64>] [mtu <750-9000>]`

5. Configure an IS-IS interface on the selected ports or MLTs:

- a. Create an IS-IS circuit and interface on the selected ports or MLTs:

```
isis
```

- b. Enable the SPBM instance on the IS-IS interfaces:

```
isis spbm <1-100>
```

- c. Enable the IS-IS circuit/interface on the selected ports or MLTs:

```
isis enable
```

6. Create the authentication key:

```
auth-key WORD<1-32>
```

7. Enable IPsec on the logical interface:

```
ipsec
```

8. Exit interface configuration mode:

```
exit
```

Switch B Steps

9. Enter IS-IS Router Configuration mode:

```
enable
```

```
configure terminal
```

```
router isis
```

10. Configure the IP tunnel source address:

```
ip-tunnel-source-address <A.B.C.D> [vrf WORD<1-16>]
```

11. Enter Global Configuration mode:

```
exit
```

12. Use one of the following commands to create a logical IS-IS interface:

- In a network with a Layer 3 Core, enter `logical-intf isis <1-255> dest-ip <A.B.C.D> [name WORD<1-64>] [mtu <750-9000>]`

13. Configure an IS-IS interface on the selected ports or MLTs:

- a. Create an IS-IS circuit and interface on the selected ports or MLTs:

```
isis
```

- b. Enable the SPBM instance on the IS-IS interfaces:

```
isis spbm <1-100>
```

- c. Enable the IS-IS circuit/interface on the selected ports or MLTs:

```
isis enable
```

14. Create the authentication key:

```
auth-key WORD<1-32>
```

15. Enable IPsec on the logical interface:

```
ipsec
```

16. Exit interface configuration mode:

```
exit
```


Variable definitions

Use the data in the following tables to set up Fabric Extend over IPsec on a device.

Use the data in the following table to use the ip-tunnel-source-address command.

Variable	Value
<A.B.C.D>	Specifies the IS-IS IPv4 tunnel source address, which can be a brouter IP, a CLIP IP, or a VLAN IP.
vrf WORD<1-16>	Specifies the VRF name associated with the IP tunnel.

Table 7: Layer 3 core

Variable	Value
<1-255>	Specifies the index number that uniquely identifies this logical interface.
<A.B.C.D>	Specifies the IS-IS IPv4 tunnel source address, which can be either a brouter interface IP or a CLIP IP.
name WORD<1-64>	Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.
mtu <750-900>	Specifies the Maximum Transmission Unit (MTU) size of each packet. The default MTU value is 1950.
 Note: Exception: only supported on XA1400 Series.	

Use the data in the following table to use the `isis` command.

Variable	Value
enable	Enables or disables the IS-IS circuit/interface on the specified port or MLT. The default is disabled. Use the no option to disable IS-IS on the specified interface.
spbm <1–100>	Enable the SPBM instance on the IS-IS interfaces.

Use the data in the following table to use the **auth-key** command.

Variable	Value
WORD<1–32>	Specifies the authentication key on the assigned logical interface, which can be up to 32 characters. Use the no option to disable the authentication key on the specified interface.

Configure Fabric Extend Tunnels

Use the following procedure to configure Fabric Extend (FE) between a Main office to a Branch office. This is a typical deployment. However, if your deployment creates tunnels between two switches that support Fabric Extend natively, then repeat those steps and ignore the steps for switches that require an ONA.

* Note:

VRF is an optional parameter. If a VRF is not configured, then FE uses the GRT.

Before you begin

The tunnel source IP address can be either a brouter port IP, a CLIP IP, or a VLAN IP.

If using the tunnel originating address on the **GRT**, Fabric Extend has the following requirements:

- The tunnel source IP address must be on the GRT, not on a VRF.

* Note:

A Best Practice is to use separate IP addresses for the SPBM IP Shortcuts **ip-source-address** command and the Fabric Extend **ip-tunnel-source-address** command. However, if you want these IP addresses to be the same, you **MUST** exclude the **ip-source-address** address with an IS-IS accept policy. You cannot use the redistribute command with a route map exclusion.

Specify a CLIP interface to use as the source address for SPBM IP shortcuts.

- If IP Shortcuts is enabled, you must configure an IS-IS accept policy or exclude route-map to ensure that tunnel destination IP addresses are not learned through IS-IS.

If you are using the tunnel originating address on a **VRF**, Fabric Extend has the following requirements:

- Configure a CLIP and tunnel source IP address on the VRF.

- Remote management of the VSP 4000 is only possible after establishing IP Shortcut over IS-IS. (Alternatively, you can enable GRT-VRF redistribution locally.)

About this task

Configuring Fabric Extend consists of two primary tasks: configuring the tunnel source address and configuring the logical interface. These tasks must be completed on both ends of the tunnel.

Note that the VSP 4000 source address command is different than other platforms. Also note that the logical interface commands are different between Layer 2 and Layer 3 networks.

Procedure

The following steps are for platforms that support FE natively:

1. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
2. Click the **Globals** tab.
3. In the **IpTunnelSourceAddress** field, enter the IP tunnel source address.
4. If you are using a VRF, select its name from the drop down menu in the **IpTunnelVrf** field.
5. Click **Apply**.

The following steps are for platforms that require an ONA to support FE:

*** Note:**

The interface VLAN connecting to the ONA network port is always in the GRT and the member port that the VLAN is part of is always an access port.

6. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
7. Click the **Globals** tab.
8. In the **IpTunnelSourceAddress** field, enter the IP tunnel source address.
9. In the **IpTunnelPort** field, select from the drop down menu the physical port that the logical interface is connected to in an L2 network.
10. If you are using a VRF, select its name from the drop down menu in the **IpTunnelVrf** field.
11. In the **IpTunnelMtu** field, enter a value between 750 and 1950 to specify the size of the maximum transmission unit (MTU). The default is 1950. This parameter only applies to an ONA configuration.
12. Click **Apply**.

Fabric Extend Field Descriptions

Use the data in the following table to use the **Globals** tab.

Name	Description
IpTunnelSourceAddress	Specifies the IS-IS IPv4 tunnel source address.

Table continues...

Name	Description
IpTunnelPort	Specifies the physical port that the logical interface is connected to in an L2 network. The parameter is for the VSP 4000 only.
IpTunnelVrf	Specifies the VRF name associated with the IP tunnel.
IpTunnelMtu	Specifies the size of the maximum transmission unit (MTU). The default is 1950. This parameter only applies to an ONA configuration.

View Tunnel CoS Queue Statistics

Use the following procedure to retrieve the tunnel CoS queue statistics. The system opens the statistics of the forwarded packets and bytes and the dropped packets and bytes.

Procedure

1. In the navigation tree, expand: **Configuration > QOS**.
2. Click **CoS Queue Stats**.
3. Select the **Tunnel** tab.

Tunnel Field Descriptions

The following table describes the fields from the CoS Queue Stats Tunnel tab.

Name	Description
Index	Indicates the loopback port number from 192(1/1) to 241(1/50).
Que<0–7>OutPackets	Indicates the out packets by CoS queue number 0–7.
Que<0–7>OutBytes	Indicates the out bytes by CoS queue number 0–7.
Que<0–7>DropPackets	Indicates the drop packets by CoS queue number 0–7.
Que<0–7>DropBytes	Indicates the drop bytes by CoS queue number 0–7.

Fabric Extend Tunnel MTU

You can configure a unique MTU value for each Fabric Extend (FE) tunnel on a XA1400 Series device. You can configure each ISIS logical interface with a unique MTU value for each FE tunnel in the VXLAN interface to improve fragmentation and reassembly in WAN connectivity over MPLS IP VPN and internet-based connections through a NAT router.

Fragmentation and reassembly is based on the MTU value configured for each FE tunnel. You can change the MTU configuration at any time for each FE tunnel. The supported MTU range is 750 to 9000, and the default MTU value is 1950.

*** Note:**

FE Tunnel MTU is an optional configuration.

For example, if you configure an FE tunnel with an MTU of 900, and a packet size of 1950 is received on UNI with the destination on the FE tunnel, the system fragments the original 1950-sized packet into the three packets (900, 900, 150) with a packet size equal to or less than 900. The system transmits the three fragmented packets over the ISIS logical interface of the FE tunnel. After the packets are received at the destination, the system performs the packet reassembly (900, 900, 150) into the 1950-sized packet.

Fabric Extend Tunnel MTU Considerations

Consider the following interactions between route MTU and FE Tunnel MTU configurations:

- If route MTU is not configured, the MTU value for each FE tunnel is applicable to ingress and egress traffic on the tunnel.
- If route MTU is configured, the MTU value for each FE tunnel is applicable for ingress traffic on the tunnel. The route MTU value applies to all egress traffic.

*** Note:**

System MTU maximum is a separate configuration. You can configure a system maximum MTU size of 1522, 1950, or 9022. The default value is 1950.

logical-intf isis

Create a logical IS-IS interface.

Syntax

- `logical-intf isis <1-255> dest-ip {A.B.C.D} name WORD<1-64> mtu <mtu_value>`
- `logical-intf isis <1-255> vid {vlan-id[-vlan-id][, ...]} primary-vid <2-4059> mlt PT_MLT<1-512> mtu <mtu_value>`
- `logical-intf isis <1-255> vid {vlan-id[-vlan-id][, ...]} primary-vid <2-4059> port {slot/port[/sub-port]} name WORD<1-64> mtu <mtu_value>`
- `no logical-intf isis <1-255>`

Command Parameters

- | | |
|-----------------------------------|---|
| <1-255> | Specifies the ISIS logical interface ID. |
| dest-ip {A.B.C.D} | Specifies the destination IP address for the logical interface. |
| mlt PT_MLT<750-9000> | Specifies the MLT ID that the logical interface is connected to in an L2 network. |

- name WORD<1-64>** Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.
- mtu <mtu_value>** Specifies the Maximum Transmission Unit (MTU) size for each packet. Different hardware platforms support different MTU ranges. Use the CLI Help to see the available range for your switch.

*** Note:**

Exception: only supported on XA1400 Series.

Identifies a single slot and port. If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

- primary-vid <2-4059>** Specifies the primary tunnel VLAN ID associated with this L2 IS-IS logical interface.
- Specifies the VLAN ID in the range of 2 to 4059. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. By default, the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998.

- vid {vlan-id [-vlan-id][,...]}** Specifies the list of VLANs that are associated with this logical interface.
- The VLAN ID is in one of the following formats: A single VLAN ID (vlan-id), a range of VLAN IDs [(vlan-id)-(vlan-id)] or a series of VLAN IDs (vlan-id, vlan-id, vlan-id).

Default

Default MTU value is 1950.

Command Mode

Global Configuration

show isis logical-interface

Display IS-IS logical interfaces.

Syntax

- **show isis logical-interface**
- **show isis logical-interface [name]**
- **show isis logical-interface [ipsec]**
- **show isis logical-interface [shaper]**
- **show isis logical-interface [mtu]**

Command Parameters

name Displays IS-IS logical interface name.

ipsec Displays IS-IS logical interface IDs with authentication key (**auth-key**) values.

shaper Displays IS-IS logical interface IDs, names, and egress shaping rate values in Mbps.

*** Note:**

Displays only interfaces with egress shaping rate values configured.

mtu Displays IS-IS logical interface IDs, names, and the Maximum Transmission Unit (MTU) values.

*** Note:**

Exception: only supported on XA1400 Series.

Default

none

Command Mode

User EXEC

Configure Egress Tunnel Shaping

About this task

Perform this procedure to configure Egress Tunnel Shaping on a logical interface.

Procedure

1. To enter User EXEC mode, log on to the switch.

2. Enter Global Configuration mode:

```
enable
```

```
configure terminal
```

3. Enable SPBM globally:

```
spbm
```

4. Enter IS-IS Router Configuration mode:

```
router isis
```

5. Create a global SPBM instance:

```
spbm <1-100>
```

6. Create a nickname for the global SPBM instance:

- ```
spbm <1-100> nick-name x.xx.xx
```
7. Add the backbone VLANs to the SPBM instance and set the primary VLAN:
 

```
spbm <1-100> b-vid {vlan-id[-vlan-id][,...]} primary <1-4059>
```
  8. Exit IS-IS Router Configuration mode:
 

```
exit
```
  9. Enter IS-IS Router Configuration mode:
 

```
router isis
```
  10. Configure the system name:
 

```
sys-name WORD<0-255>
```
  11. Configure the global router type:
 

```
is-type l1
```
  12. Configure the manual area:
 

```
manual-area xx.xxx.xxx...xxxx
```
  13. Exit IS-IS Router Configuration mode:
 

```
exit
```
  14. Create two SPBM Backbone VLANs that correspond to those configured in the previous step using the following command twice:
 

```
vlan create <1-4059> type spbm-bvlan
```
  15. Enable IS-IS globally:
 

```
router isis enable
```
  16. Remove the brouter port from all VLANs:
 

```
vlan members remove <1-4059> {slot/port[/sub-port] [-slot/port[/sub-port]][,...]} {
```
  17. Enter Interface Configuration mode:
 

```
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]][,...]}
```
  18. Configure a brouter port:
 

```
brouter port {slot/port[/sub-port]} vlan <1-4059> subnet {A.B.C.D/X}
mac-offset <MAC-offset>
```
  19. Exit Interface Configuration mode:
 

```
exit
```
  20. Enter IS-IS Router Configuration mode:
 

```
router isis
```

## New in this Release

21. Configure the IP tunnel source address:

```
ip-tunnel-source-address {A.B.C.D}
```

22. Exit IS-IS Router Configuration mode:

```
exit
```

23. Create a logical IS-IS interface and enter Logical Interface Configuration mode:

```
logical-intf isis <1-255> dest-ip {A.B.C.D} name WORD <1-64> [mtu <750-9000>]
```

24. Create an IS-IS circuit and interface:

```
isis
```

25. Enable the SPBM instance:

```
isis spbm <1-100>
```

26. Enable the IS-IS circuit and interface:

```
isis enable
```

27. Configure the Egress Tunnel Shaper:

```
egress-shaping-rate <1-1000>
```

28. Exit Logical Interface Configuration mode:

```
exit
```

## Variable Definitions

Use the data in the following table to use the **vlan members** command.

| Variable                                             | Value                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| remove <1-4059>                                      | Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. |
| {slot/port[/sub-port][/-slot/port[/sub-port]][,...]} | Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.                                               |

Use the data in the following table to use the **spbm** command.

| Variable | Value                                                                                                                        |
|----------|------------------------------------------------------------------------------------------------------------------------------|
| <1-100>  | Specifies the Shortest Path Bridging MAC (SPBM) instance ID. Creates the SPBM instance. Only one SPBM instance is supported. |

Use the data in the following table to use the **spbm** command to create a system nickname for an SPBM instance.

| Variable                 | Value                                                                                                                        |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------|
| <1-100>                  | Specifies the Shortest Path Bridging MAC (SPBM) instance ID. Creates the SPBM instance. Only one SPBM instance is supported. |
| <i>nick-name</i> x.xx.xx | Specifies the system nickname (2.5 bytes in the format ).                                                                    |

Use the data in the following table to use the **spbm** command to assign Backbone VLANs to the SPBM instance.

| Variable                               | Value                                                                                                                                              |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <1-100>                                | Specifies the Shortest Path Bridging MAC (SPBM) instance ID. Creates the SPBM instance. Only one SPBM instance is supported.                       |
| <i>b-vid</i> {vlan-id[-vlan-id][,...]} | Specifies the VLANs to add to the Shortest Path Bridging MAC (SPBM) instance as Backbone VLANs (B-VLANs). Sets the IS-IS SPBM instance data VLANs. |
| <i>primary</i> <1-4059>                | Specifies the primary BVLAN by VLAN ID.                                                                                                            |

Use the data in the following table to use the **sys-name** command.

| Variable            | Value                      |
|---------------------|----------------------------|
| <i>WORD</i> <0-255> | Specifies the system name. |

Use the data in the following table to use the **is-type** command.

| Variable  | Value                                                                                     |
|-----------|-------------------------------------------------------------------------------------------|
| <i>l1</i> | Configures the router type as Level 1 Intermediate-System-to-Intermediate-System (IS-IS). |

Use the data in the following table to use the **manual-area** command.

| Variable                   | Value                                                                                                                                                                                                 |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>xx.xxxx.xxxx...xxxx</i> | Configures the manual area in a size up to 13 octets. The current release supports one area. For Intermediate-System-to-Intermediate-System (IS-IS) to operate, you must configure at least one area. |

Use the data in the following table to use the **vlan create** command.

| Variable               | Value                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <1-4059>               | Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. |
| <i>type spbm-bvlan</i> | Specifies the VLAN type as the backbone VLAN (B-VLAN) for Shortest Path Bridging MAC (SPBM).                                                                                                                                                                                                                                                                                                                 |

Use the data in the following table to use the **interface GigabitEthernet** command.

| Variable                                             | Value                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| {slot/port[/sub-port][/-slot/port[/sub-port]][,...]} | Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. |


Use the data in the following table to use the **brouter port** command.

| Variable                             | Value                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| {slot/port[/sub-port]}               | Identifies a single slot and port. If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.                                                                                                                                                                                                                                |
| vlan <1-4059>                        | Specifies the VLAN ID in the range of 1 to 4059. By default, VLAN IDs 1 to 4059 are configurable and the system reserves VLAN IDs 4060 to 4094 for internal use. On switches that support the vrf-scaling and spbm-config-mode boot configuration flags, if you enable these flags, the system also reserves VLAN IDs 3500 to 3998. VLAN ID 1 is the default VLAN and you cannot create or delete VLAN ID 1. |
| <i>subnet &lt;A.B.C.D/X&gt;</i>      | Assigns an IP address and mask for the management port.                                                                                                                                                                                                                                                                                                                                                      |
| <i>mac-offset &lt;MAC-offset&gt;</i> | Specifies a number by which to offset the MAC address from the chassis MAC address. This ensures that each IP address has a different MAC address. If you omit this variable, a unique MAC offset is automatically generated. Different hardware platforms support different ranges. To see which range is available on your switch, use the CLI command completion Help.                                    |

Use the data in the following table to use the `ip-tunnel-source-address` command.

| Variable               | Value                                           |
|------------------------|-------------------------------------------------|
| <code>{A.B.C.D}</code> | Specifies the IS-IS IPv4 tunnel source address. |

Use the following table to use the `logical-intf` command.

| Variable                                                                                                                                     | Value                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <code>isis &lt;1–255&gt;</code>                                                                                                              | Specifies the ISIS logical interface ID.                                                                  |
| <code>dest-ip {A.B.C.D}</code>                                                                                                               | Specifies the destination IP address for the logical interface.                                           |
| <code>name WORD &lt;1–64&gt;</code>                                                                                                          | Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters. |
| <code>mtu &lt;750–9000&gt;</code>                                                                                                            | Specifies the Maximum Transmission Unit (MTU) size for each packet. Default mtu value is 1950.            |
|  <b>Note:</b><br>Exception: only supported on XA1400 Series |                                                                                                           |

Use the data in the following table to use the `egress-shaping-rate` command.

| Variable                    | Value                                   |
|-----------------------------|-----------------------------------------|
| <code>&lt;1-1000&gt;</code> | Specifies the shaper bandwidth in Mbps. |

---

## Fabric Extend configuration using EDM

The following sections provide procedural information you can use to configure Fabric Extend (FE) using Enterprise Device Manager (EDM).

### Configure Fabric Extend Logical Interfaces

Use the following procedure to configure Fabric Extend (FE) between a Main office to a Branch office. This is a typical deployment. However, if your deployment creates tunnels between two switches that support Fabric Extend natively, then repeat those steps and ignore the steps for switches that require an ONA.

 **Note:**

VRF is an optional parameter. If a VRF is not configured, then FE uses the GRT.

#### About this task

Configuring Fabric Extend consists of two primary tasks: configuring the tunnel source address and configuring the logical interface. These tasks must be completed on both ends of the tunnel.

Note that the VSP 4000 Series source address command is different than other platforms. Also note that the logical interface commands are different between Layer 2 and Layer 3 networks.

#### Procedure

**The following steps are for platforms that support FE natively:**

1. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
2. Click the Logical Interfaces tab.
3. Click Insert.
4. In the **Id** field, enter the index number that uniquely identifies this logical interface.
5. In the **Name** field, enter the name of this logical interface.
6. In the **Type** field, select the type of core network that the tunnel will be traversing. If it's a Layer 2 Core, select **layer2**. If it's a Layer 3 Core, select **ip**.

**\* Note:**

Different fields will be available depending on which type of core network you select.

7. For a Layer 2 Core, complete the following fields:
  - a. In the **DestIfIndex** field, click the ellipsis button (...) to select the physical port that the logical interface is connected to or enter the name of the MLT.
  - b. In the **Vids** field, enter the list of VLANs for this logical interface.
  - c. In the **PrimaryVid** field, enter the primary tunnel VLAN ID.

**\* Note:**

The primary VLAN ID must be one of the VLANs listed in the **Vids** field.

8. For a Layer 3 Core, complete the following field:

In the **DestIPAddr** field, enter the destination IP address for the logical interface.
9. In the **IpsecEnable** field, select whether you want to enable a Fabric Extend over IPsec connection for the logical interface.
10. In the **AuthenticationKey** field, enter the authentication key that will be used to secure your Fabric Extend over IPsec connection for the logical interface. The key may be up to 32 characters in length.
11. In the **ShapingRate** field, enter the value in Mbps of the shaper used for Egress Tunnel Shaping.
12. In the **Mtu** field, enter a value to specify the size of the maximum transmission unit (MTU). The default is 1950.
13. Click **Insert**.

**The following steps are for platforms that require an ONA to support FE:**

**\* Note:**

The interface VLAN connecting to the ONA network port is always in the GRT and the member port that the VLAN is part of is always an access port.

14. In the navigation pane, expand the **Configuration > IS-IS > IS-IS** folders.
15. Click the Logical Interfaces tab.



16. Click **Insert**.
17. In the **Id** field, enter the index number that uniquely identifies this logical interface.
18. In the **Name** field, enter the name of this logical interface.
19. In the **Type** field, select the type of core network that the tunnel will be traversing. If it's a Layer 2 Core, select **layer2**. If it's a Layer 3 Core, select **ip**.

**\* Note:**

Different fields will be available depending on which type of core network you select.

20. For a Layer 2 Core, complete the following fields:
  - a. In the **DestIfIndex** field, click the ellipsis button (...) to select the physical port that the logical interface is connected to or enter the name of the MLT.
  - b. In the **Vids** field, enter the list of VLANs for this logical interface.
  - c. In the **PrimaryVid** field, enter the primary tunnel VLAN ID.

**\* Note:**

The primary VLAN ID must be one of the VIDs listed in the **Vids** field.

21. For a Layer 3 Core, complete the following field:
 









In the **DestIPAddr** field, enter the destination IP address for the logical interface.
22. In the **IpsecEnable** field, select whether you want to enable a Fabric Extend over IPsec connection for the logical interface.
23. In the **AuthenticationKey** field, enter the authentication key that will be used to secure your Fabric Extend over IPsec connection for the logical interface. The key may be up to 32 characters in length.
24. In the **ShapingRate** field, enter the value in Mbps of the shaper used for Egress Tunnel Shaping.
25. Click **Insert**.

## Logical Interfaces Field Descriptions

Use the data in the following table to use the **Logical Interfaces** tab and the Insert Logical Interfaces dialog. The available fields in the dialog differ depending on the type of core you select: **layer 2** or **ip**.

| Name           | Description                                                                                                                                         |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Id</b>      | Specifies the index number that uniquely identifies this logical interface.<br><br>This field appears only on the Insert Logical Interfaces dialog. |
| <b>IfIndex</b> | Specifies the index number that uniquely identifies this logical interface. This field is read-only.                                                |

*Table continues...*

| Name                                                                                                                                                                        | Description                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                                                                                                             | This field appears only on the Logical Interfaces tab.                                                                                                                                                                                                                                      |
| <b>Name</b>                                                                                                                                                                 | Specifies the administratively-assigned name of this logical interface, which can be up to 64 characters.                                                                                                                                                                                   |
| <b>Type</b><br> <b>Note:</b><br>Exception: type Layer 2 is not supported on XA1400 Series. | Specifies the type of logical interface to create:<br><ul style="list-style-type: none"> <li>• Specify <b>layer 2</b> for a Layer 2 core network that the tunnel will be traversing.</li> <li>• Specify <b>ip</b> for a Layer 3 core network that the tunnel will be traversing.</li> </ul> |
| <b>DestIPAddr</b>                                                                                                                                                           | Specifies the destination IP address for the IP-type logical interface.                                                                                                                                                                                                                     |
| <b>DestIfIndex</b><br> <b>Note:</b><br>Exception: not supported on XA1400 Series.          | Specifies the physical port or MultiLink Trunking (MLT) that the Layer 2 logical interface is connected to.                                                                                                                                                                                 |
| <b>Vids</b><br> <b>Note:</b><br>Exception: not supported on XA1400 Series.                 | Specifies the list of VLANs that are associated with this logical interface.                                                                                                                                                                                                                |
| <b>PrimaryVid</b><br> <b>Note:</b><br>Exception: not supported on XA1400 Series.          | Specifies the primary tunnel VLAN ID associated with this L2 Intermediate-System-to-Intermediate-System (IS-IS) logical interface.                                                                                                                                                          |
| <b>CircIndex</b><br> <b>Note:</b><br>Exception: not supported on XA1400 Series.          | Identifies the IS-IS circuit created under the logical interface.<br><br>This field appears only on the Logical Interfaces tab.                                                                                                                                                             |
| <b>NextHopVrf</b><br> <b>Note:</b><br>Exception: not supported on XA1400 Series.         | Identifies the next-hop VRF name to reach the logical tunnel destination IP.<br><br>This field appears only on the Logical Interfaces tab.                                                                                                                                                  |
| <b>IpssecEnable</b><br> <b>Note:</b><br>Exception: only supported on XA1400 Series.      | Specifies whether the logical interace should use IPsec.                                                                                                                                                                                                                                    |
| <b>AuthenticationKey</b><br> <b>Note:</b><br>Exception: only supported on XA1400 Series. | Specifies the authentication key of this logical interface, which can be up to 32 characters.                                                                                                                                                                                               |
| <b>IpssecNatConfigResponderOnly</b>                                                                                                                                         | Specifies whether the device is a Responder device in an IPsec Network Address Translation Traversal (NAT-T) connection.                                                                                                                                                                    |

*Table continues...*

| Name                                                                                                               | Description                                                                                                       |
|--------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| <p><b>* Note:</b><br/>Exception: only supported on XA1400 Series.</p>                                              |                                                                                                                   |
| <p><b>IpssecNatConfigRemoteNatIPAddr</b></p> <p><b>* Note:</b><br/>Exception: only supported on XA1400 Series.</p> | Specifies the public IP address of the NAT router connected to the Responder device in an IPsec NAT-T connection. |
| <p><b>ShapingRate</b></p> <p><b>* Note:</b><br/>Exception: only supported on XA1400 Series.</p>                    | Specifies the value, in Mbps, of the Egress Tunnel Shaper applied to the logical interface.                       |
| <b>Mtu</b>                                                                                                         | Specifies the Maximum Transmission Unit (MTU) size for each logical interface. The default MTU value is 1950.     |

## Port-Rate Limiting, Policing, and Shaping

**Table 8: Rate Limiting, Policing, and Shaping product support**

| Feature                                                                                                   | Product         | Release introduced |
|-----------------------------------------------------------------------------------------------------------|-----------------|--------------------|
| For configuration details, see <a href="#">Configuring QoS and ACL-Based Traffic Filtering for VOSS</a> . |                 |                    |
| Egress port shaper                                                                                        | VSP 4450 Series | VSP 4000 4.0       |
|                                                                                                           | VSP 4900 Series | VOSS 8.1           |
|                                                                                                           | VSP 7200 Series | VOSS 4.2.1         |
|                                                                                                           | VSP 7400 Series | VOSS 8.0           |
|                                                                                                           | VSP 8200 Series | VSP 8200 4.0       |
|                                                                                                           | VSP 8400 Series | VOSS 4.2           |
|                                                                                                           | VSP 8600 Series | VSP 8600 4.5       |
|                                                                                                           | XA1400 Series   | VOSS 8.0.50        |
| Ingress dual rate port policers                                                                           | VSP 4450 Series | VSP 4000 4.0       |
|                                                                                                           | VSP 4900 Series | Not Supported      |
|                                                                                                           | VSP 7200 Series | Not Supported      |
|                                                                                                           | VSP 7400 Series | Not Supported      |
|                                                                                                           | VSP 8200 Series | Not Supported      |
|                                                                                                           | VSP 8400 Series | Not Supported      |
|                                                                                                           | VSP 8600 Series | VSP 8600 4.5       |
|                                                                                                           | XA1400 Series   | Not Supported      |

*Table continues...*

| Feature                       | Product         | Release introduced |
|-------------------------------|-----------------|--------------------|
| QoS ingress port rate limiter | VSP 4450 Series | Not Supported      |
|                               | VSP 4900 Series | VOSS 8.1           |
|                               | VSP 7200 Series | VOSS 4.2.1         |
|                               | VSP 7400 Series | Not Supported      |
|                               | VSP 8200 Series | VSP 8200 4.0       |
|                               | VSP 8400 Series | VOSS 4.2           |
|                               | VSP 8600 Series | Not Supported      |
|                               | XA1400 Series   | VOSS 8.1.50        |

The switch QoS implementation supports the following two features for bandwidth management and traffic control:

- ingress port-rate limiting—a mechanism to limit the traffic rate accepted by the specified ingress port

**\* Note:**

The VSP 4900 Series, VSP 7400 Series, and XA1400 Series do not support ingress policers. The VSP 7400 Series does not support port-based rate limiting.

- egress port-rate shaping—the process by which the system delays and transmits packets to produce an even and predictable flow rate

Each port has eight unicast and multicast queues, Class of Service (CoS) 0 to CoS 7. Traffic shaping exists on the egress CoS 6 and CoS 7, but you cannot change the configuration. CoS 6 and CoS 7 are strict priority queues, with traffic shaping for CoS 6 at 50 percent and CoS 7 to five percent of line rate.

Some VOSS hardware platforms allow you to configure an egress shaping rate for each port manually. For XA1400 Series, the egress shaping rate for each front panel port dynamically adjusts to the auto-negotiated link speed, up to the maximum link speed of the port.

The VSP 4000 Series switch QoS implementation supports the following two features for bandwidth management and traffic control:

- ingress traffic policing—a mechanism to limit the number of packets in a stream that matches a particular classification
- egress traffic shaping—the process by which the system delays (or drops) and transmits packets to produce an even and predictable flow rate

Each feature is important to deliver DiffServ within a QoS network domain.

**Token buckets**

Tokens are a key concept in traffic control. A port-rate limiter, policer, or shaper calculates the number of packets that passed, and at what data rate. Each packet corresponds to a token, and the port-rate limiter, policer, or shaper transmits or passes the packet if the token is available. For more information, see [Figure 4: Token flow](#) on page 53.

The token container is like a bucket. In this view, the bucket represents both the number of tokens that a port-rate limiter, policer, or shaper can use instantaneously (the depth of the bucket) and the rate at which the tokens replenish (how fast the bucket refills).

Each policer has two token buckets: one for the peak rate and the other for the service rate. The following figure shows the flow of tokens.

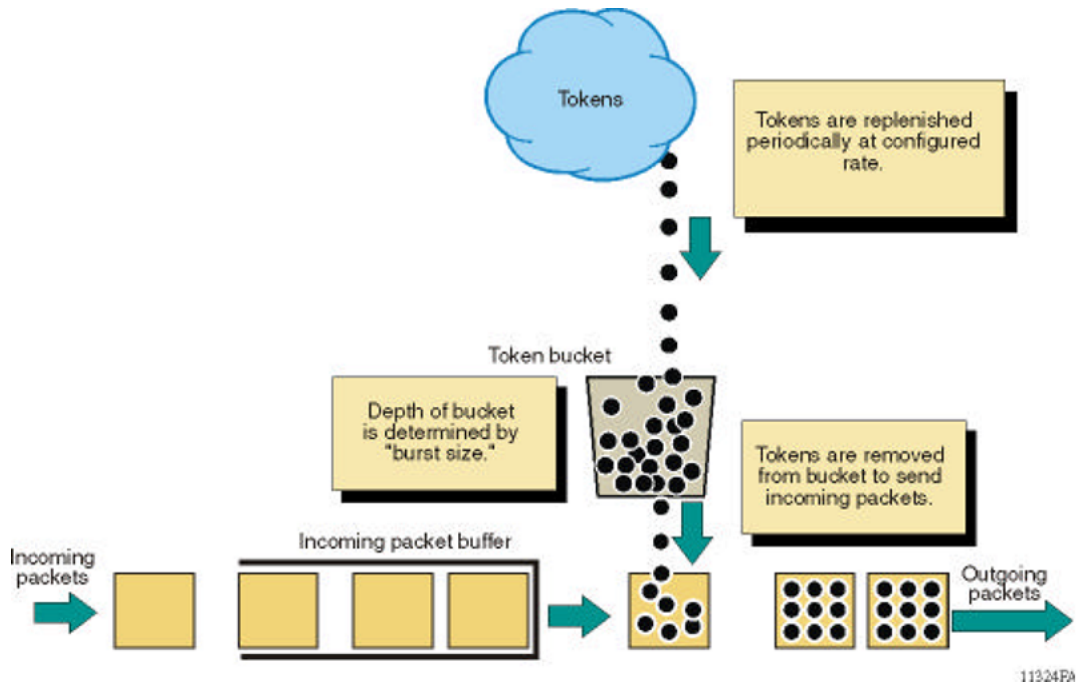


Figure 4: Token flow

## IEEE 802.3X Pause Frame Transmit

Table 9: IEEE 802.3X Pause Frame Transmit product support

| Feature                                                             | Product         | Release introduced |
|---------------------------------------------------------------------|-----------------|--------------------|
| For configuration details, see <a href="#">Administering VOSS</a> . |                 |                    |
| IEEE 802.3X Pause frame transmit                                    | VSP 4450 Series | VOSS 6.0           |
|                                                                     | VSP 4900 Series | VOSS 8.1           |
|                                                                     | VSP 7200 Series | VOSS 6.0           |
|                                                                     | VSP 7400 Series | VOSS 8.0           |
|                                                                     | VSP 8200 Series | VOSS 6.0           |
|                                                                     | VSP 8400 Series | VOSS 6.0           |
|                                                                     | VSP 8600 Series | Not Supported      |
|                                                                     | XA1400 Series   | VOSS 8.1.50        |

The switch uses MAC pause frames to provide congestion relief on full-duplex interfaces.

## Overview

When congestion occurs on a port, the system can send or receive pause frames, also known as flow control, to temporarily pause the packet flow. The system uses flow control if the rate at which one or more ports receives or sends packets is greater than the rate the switch can process or accept the packets.

The switch can generate pause frames to tell the sending device to stop sending additional packets for a specified time period. After the time period expires, the sending device can resume sending packets. During the specified time period, if the switch determines the congestion is reduced, it can send pause frames to the sending device to instruct it to begin sending packets immediately.

## Flow control mode and pause frames

If you enable flow control mode, the switch drops packets on ingress when congestion occurs. If the switch is not in flow control mode, it drops packets at egress when congestion occurs.

Configure an interface to send pause frames when congestion occurs to alleviate packet drops due to flow control mode.

## Auto-Negotiation

Interfaces that support auto-negotiation advertise and exchange their flow control capability to agree on a pause frame configuration. IEEE 802.3 annex 28b defines the auto-negotiation ability fields and the pause resolution. The switch advertises only two capabilities. The following table shows the software bit settings based on the flow control configuration.

### \* Note:

Not all interfaces support Auto-Negotiation. For more information, see your hardware documentation.

**Table 10: Advertised abilities**

| Interface configuration | Pause | ASM | Capability advertised                     |
|-------------------------|-------|-----|-------------------------------------------|
| Flow control enabled    | 1     | 0   | Symmetric pause                           |
| Flow control disabled   | 1     | 1   | Both Symmetric pause and asymmetric pause |

The following tables identifies the pause resolution.

**Table 11: Pause resolution**

| Local device pause | Local device ASM | Peer device pause | Peer device ASM | Local device resolution             | Peer device resolution              |
|--------------------|------------------|-------------------|-----------------|-------------------------------------|-------------------------------------|
| 0                  | 0                | Do not care       | Do not care     | Disable pause transmit and receive. | Disable pause transmit and receive. |
| 0                  | 1                | 0                 | Do not care     | Disable pause transmit and receive. | Disable pause transmit and receive. |

*Table continues...*

| Local device pause | Local device ASM | Peer device pause | Peer device ASM | Local device resolution                          | Peer device resolution                           |
|--------------------|------------------|-------------------|-----------------|--------------------------------------------------|--------------------------------------------------|
| 0                  | 1                | 1                 | 0               | Disable pause transmit and receive.              | Disable pause transmit and receive.              |
| 0                  | 1                | 1                 | 1               | Enable pause transmit.<br>Disable pause receive. | Disable pause transmit. Enable pause receive.    |
| 1                  | 0                | 0                 | Do not care     | Disable pause transmit and receive.              | Disable pause transmit and receive.              |
| 1                  | Do not care      | 1                 | Do not care     | Enable pause transmit and receive.               | Enable pause transmit and receive.               |
| 1                  | 1                | 0                 | 0               | Disable pause transmit and receive.              | Disable pause transmit and receive.              |
| 1                  | 1                | 0                 | 1               | Disable pause transmit. Enable pause receive.    | Enable pause transmit.<br>Disable pause receive. |

The following list identifies the type of interfaces that support auto-negotiated flow control:

- 10 Mbps/100 Mbps/1 Gbps copper
- 100 Mbps/1 Gbps/10 Gbps copper
- 1 Gbps fiber (in both SFP and SFP+ ports)

---

## Configure IEEE 802.3X Pause Frame Transmit

Configure IEEE 802.3X Pause frame transmit to eliminate or minimize packet loss.

### About this task

By default, flow control mode is disabled. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames.

By default, an interface does not send pause frames.

#### Note:

If you enable MACsec on an interface and you send small packet size traffic near line rate, the **In FlowCtrl** frame might increment in the output of the **show interface gigabitEthernet statistics** command because of the processing overhead caused by adding the MACsec header of 32 bytes. This is part of the expected over-subscription footprint.

**Procedure**

1. Enter Global Configuration mode:

```
enable
configure terminal
```

2. Enable flow control mode:

```
boot config flags flow-control-mode
```

3. Save the configuration.

4. Exit Privileged EXEC mode:

```
exit
```

5. Reboot the chassis.

```
boot
```

6. Enter GigabitEthernet Interface Configuration mode:

```
enable
configure terminal
interface GigabitEthernet {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

**\* Note:**

If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port.

7. Configure the interface to generate pause frames:

```
tx-flow-control [enable]
```

8. **(Optional)** Configure other interfaces to generate pause frames:

```
tx-flow-control port {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]} enable
```

9. Verify the boot flag configuration:

```
show boot config flags
```

10. Verify the interface configuration:

```
show interfaces gigabitEthernet l1-config {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```

11. View the pause-frame packet count:

```
show interfaces gigabitEthernet statistics {slot/port[/sub-port] [-slot/port[/sub-port]] [,...]}
```



## Example

Enable flow control on the system and configure slot 1, port 10 to send pause frames. Verify the configuration.

```
Switch:1>enable
Switch:1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch:1(config)#boot config flags flow-control-mode
Warning: Please save the configuration and reboot the switch
 for this configuration to take effect.
Switch:1<config>#save config
CP-1: Save config to file /intflash/config.cfg successful.
CP-1: Save license to file /intflash/license.xml successful.
Switch:1<config>#exit
Switch:1#boot
Are you sure you want to re-boot the switch (y/n) ?y
```

### \* Note:

Flag support can vary across hardware models.

```
Switch:1#show boot config flags
flags advanced-feature-bandwidth-reservation low
flags block-snmp false
flags debug-config false
flags debugmode false
flags dvr-leaf-mode false
flags enhancedsecure-mode false
flags factorydefaults false
flags flow-control-mode true
flags ftpd true
flags ha-cpu true
flags hsecure false
flags insight-port-connect-type vtd
flags ipv6-egress-filter true
flags ipv6-mode false
flags linerate-directed-broadcast false
flags logging true
flags nni-mstp false
flags reboot true
flags rlogind false
flags savetostandby true
flags spanning-tree-mode mstp
flags spbm-config-mode true
flags sshd true
flags syslog-rfc5424-format true
flags telnetd true
flags tftpd true
flags trace-logging false
flags urpf-mode true
flags verify-config true
flags vrf-scaling true
flags vxlan-gw-full-interworking-mode false
```

```
Switch:1(config-if)#show interfaces gigabitEthernet 1/10
```

```
=====
 Port Config L1
=====
```

| PORT<br>NUM | AUTO<br>NEG. | OPERATE<br>AUTO-NEG | CUSTOM<br>ADVERTISEMENTS | AUTO<br>NEGOTIATION | ADMIN<br>DPLX SPD | OPERATE<br>DPLX SPD | ADMIN<br>TX-FLW-CTRL | OPERATE<br>TX-FLW-CTRL |
|-------------|--------------|---------------------|--------------------------|---------------------|-------------------|---------------------|----------------------|------------------------|
| 1/10        | true         | true                | Not Configured           |                     | full 10000        | 0                   | enable               | enable                 |

```
=====
```

View the pause-frame packet count for slot 1, port 10.

```
Switch:1(config-if)#show interfaces gigabitEthernet statistics 1/10
=====
Port Stats Interface
=====
PORT IN OUT IN OUT
NUM OCTETS OCTETS PACKET PACKET

1/1 29964704384 22788614528 234106526 178034166

PORT IN OUT IN OUT OUTLOSS
NUM FLOWCTRL FLOWCTRL PFC PFC PACKETS

1/1 0 11014 0 0 0
```

## Variable Definitions

Use the data in the following table to use the `tx-flow-control` command.

| Variable                                                  | Value                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| enable                                                    | Configures the interface to send pause frames. By default, flow control is disabled.                                                                                                                                                                                                                                                                           |
| port {slot/port[/sub-port] [-slot/port[/sub-port]] [...]} | Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. |

Use the data in the following table to use the `show interfaces gigabitEthernet l1-config` and `show interfaces gigabitEthernet statistics` commands.

| Variable                                             | Value                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| {slot/port[/sub-port] [-slot/port[/sub-port]] [...]} | Identifies the slot and port in one of the following formats: a single slot and port (slot/port), a range of slots and ports (slot/port-slot/port), or a series of slots and ports (slot/port,slot/port,slot/port). If your platform supports channelization and the port is channelized, you must also specify the sub-port in the format slot/port/sub-port. |

## Configure basic port parameters

Configure options for port operations.

### About this task

If you select more than one port, the format of the tab changes to a table-based tab.

**\* Note:**

When you use 1 Gigabit Ethernet SFP transceivers on VSP 7254XSQ, the software disables auto-negotiation on the port:

- If you use 1 Gbps fiber SFP transceivers, the remote end must also have auto-negotiation disabled. Otherwise this is not a supported configuration with VSP 7254XSQ.
- If you use 1 Gbps copper SFP transceivers, the remote end must have auto-negotiation enabled. If not, the link will not be established.

## Procedure

1. In the Device Physical View tab, select one or more ports.
2. In the navigation pane, expand the **Configuration > Edit > Port** folders.
3. Click **General**.
4. Click the **Interface** tab.
5. Configure the fields as required.

10/100BASE-TX ports do not consistently auto-negotiate with older 10/100BASE-TX equipment. You can sometimes upgrade the older devices with new firmware or driver revisions. If an upgrade does not allow auto-negotiation to correctly identify the link speed and duplex settings, you can manually configure the settings for the link in question.

Check the Extreme Networks Web site for the latest compatibility information.

6. Click **Apply**.

## Interface Field Descriptions


Use the data in the following table to use the Interface tab.

| Name               | Description                                                                                                                                                                                                                                                                                                           |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Index</b>       | Displays the index of the port, written in the slot/port[/sub-port] format.                                                                                                                                                                                                                                           |
| <b>Name</b>        | Configures the name of the port.                                                                                                                                                                                                                                                                                      |
| <b>Descr</b>       | Displays the description of the port. A textual string containing information about the interface.                                                                                                                                                                                                                    |
| <b>Type</b>        | Displays the type of connector plugged in the port.                                                                                                                                                                                                                                                                   |
| <b>Mtu</b>         | Displays the Maximum Transmission Unit (MTU) for the port. The size of the largest datagram which can be sent or received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface. |
| <b>PhysAddress</b> | Displays the physical address of the port. The address of the interface at the protocol layer immediately `below' the network layer in the protocol                                                                                                                                                                   |

*Table continues...*

| Name                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                             | stack. For interfaces which do not have such an address (e.g., a serial line), this object should contain an octet string of zero length.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>VendorDescr</b>          | Displays the vendor of the connector plugged in the port.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>DisplayFormat</b>        | Identifies the slot and port numbers (slot/port). If the port is channelized, the format also includes the sub-port in the format slot/port/sub-port                                                                                                                                                                                                                                                                                                                                                                    |
| <b>AdminStatus</b>          | Configures the port as enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>OperStatus</b>           | Displays the current status of the port. The status includes enabled (up) or disabled (down) or testing. The testing state indicates that no operational packets can be passed.                                                                                                                                                                                                                                                                                                                                         |
| <b>LicenseControlStatus</b> | Shows the port license status. This field only applies to VSP 7200 Series.                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>ShutdownReason</b>       | Indicates the reason for a port state change.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>LastChange</b>           | Displays the timestamp of the last change.                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>LinkTrap</b>             | Enable or disable link trapping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>AutoNegotiate</b>        | <p>Enables or disables Auto-Negotiation for this port.</p> <p>The default varies depending on the platform:</p> <ul style="list-style-type: none"> <li>• VSP 4000 Series - Enabled</li> <li>• VSP 4900 Series - Enabled</li> <li>• VSP 7200 Series - Disabled</li> <li>• VSP 7400 Series - Enabled</li> <li>• VSP 8200 Series - Enabled</li> <li>• VSP 8400 Series - Enabled</li> <li>• VSP 8600 Series - Enabled (except 10 Gbps SFP+ ports)</li> <li>• XA1400 Series - Enabled (except 10 Gbps SFP+ ports)</li> </ul> |
| <b>AutoNegAd</b>            | <p>Specifies the port speed and duplex abilities to advertise during link negotiation.</p> <p>Supported speeds and duplex modes vary, depending on your hardware.</p> <p>The abilities specified in this object are only used when auto-negotiation is enabled on the port. If all bits in this object are disabled, and auto-negotiation</p>                                                                                                                                                                           |

*Table continues...*

| Name                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                              | <p>is enabled on the port, then the physical link process on the port will be disabled (if hardware supports this ability).</p> <p>Any change to this configuration restarts the auto-negotiation process, which has the same effect as physically unplugging and reattaching the cable attached to the port.</p> <p>If you select <b>default</b>, all capabilities supported by the hardware are advertised.</p>                                         |
| <b>AdminDuplex</b>           | Configures the administrative duplex setting for the port.                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>OperDuplex</b>            | Indicates the operational duplex setting for the port.                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>AdminSpeed</b>            | <p>Configures the administrative speed for the port.</p> <p> <b>Important:</b></p> <p>If Auto-Negotiation is disabled and you change the administrative speed on a port that results in a configuration mismatch in speed between two ports, VSP 4450 Series and VSP 4900 Series switches may show an incorrect operational status of "up" for the mismatched ports.</p> |
| <b>OperSpeed</b>             | Indicates the operational speed for the port.                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>QoSLevel</b>              | Selects the Quality of Service (QoS) level for this port. The default is level1.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>DiffServ</b>              | Enables the Differentiated Service feature for this port. The default is disabled.                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Layer3Trust</b>           | Configures if the system should trust Layer 3 packets coming from access links or core links only. The default is core.                                                                                                                                                                                                                                                                                                                                   |
| <b>Layer2Override8021p</b>   | Specifies whether Layer 2 802.1p override is enabled (selected) or disabled (cleared) on the port. The default is disabled (clear).                                                                                                                                                                                                                                                                                                                       |
| <b>MltId</b>                 | Shows the MLT ID associated with this port. The default is 0.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Locked</b>                | Shows if the port is locked. The default is disabled.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>UnknownMacDiscard</b>     | Discards packets that have an unknown source MAC address, and prevents other ports from sending packets with that same MAC address as the destination MAC address. The default is disabled.                                                                                                                                                                                                                                                               |
| <b>DirectBroadcastEnable</b> | Specifies if this interface forwards direct broadcast traffic.                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>OperRouting</b>           | Shows the routing status of the port.                                                                                                                                                                                                                                                                                                                                                                                                                     |

*Table continues...*

| Name                                                                                                                                                                                            | Description                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>HighSecureEnable</b>                                                                                                                                                                         | Enables or disables the high secure feature for this port.                                                                                                                                                                                                                                                                                |
| <b>RmonEnable</b>                                                                                                                                                                               | Enables or disables Remote Monitoring (RMON) on the interface. The default is disabled.                                                                                                                                                                                                                                                   |
| <b>FlexUniEnable</b>                                                                                                                                                                            | Enables Flex UNI on the port. The default is disabled.                                                                                                                                                                                                                                                                                    |
| <b>IngressRateLimit</b><br>* <b>Note:</b><br>Exception: only supported on , VSP 4900 Series, VSP 7200 Series, , VSP 8200 Series, VSP 8400 Series, and XA1400 Series.                            | Limits the traffic rate accepted by the specified ingress port.                                                                                                                                                                                                                                                                           |
| <b>IngressRatePeak</b>                                                                                                                                                                          | Configures the peak rate in Kbps. The default is 0.                                                                                                                                                                                                                                                                                       |
| <b>IngressRateSvc</b>                                                                                                                                                                           | Configures the service rate in Kbps. The default is 0.                                                                                                                                                                                                                                                                                    |
| <b>EgressRateLimitState</b>                                                                                                                                                                     | Enables or disables egress port-based shaping to bind the maximum rate at which traffic leaves the port. The default is disabled.                                                                                                                                                                                                         |
| <b>EgressRateLimit</b><br>* <b>Note:</b><br>Exception: only supported on VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, , VSP 8200 Series, and VSP 8400 Series.                             | Configures the egress rate limit in Kbps. Different hardware platforms support different egress rate limits, depending on the port with the highest speed available on the platform. You cannot configure the egress shaper rate to exceed the port capability.<br><br>If you configure this value to 0, shaping is disabled on the port. |
| <b>TxFlowControl</b><br>* <b>Note:</b><br>Exception: only supported on VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and XA1400 Series. | Configures if the port sends pause frames. By default, an interface does not send pause frames.<br><br>You must also enable the flow control feature globally before an interface can send pause frames.                                                                                                                                  |
| <b>TxFlowControlOperState</b>                                                                                                                                                                   | Shows the operational state of flow control.                                                                                                                                                                                                                                                                                              |
| <b>BpduGuardTimerCount</b>                                                                                                                                                                      | Shows the time, starting at 0, since the port became disabled. When the BpduGuardTimerCount reaches the BpduGuardTimeout value, the port is enabled. Displays in 1/100 seconds.                                                                                                                                                           |
| <b>BpduGuardTimeout</b>                                                                                                                                                                         | Specifies the value to use for port-state recovery. After a BPDU guard disables a port, the port remains in the disabled state until this timer expires.<br><br>You can configure a value of 0 or to 65535. The default is 120 seconds. If you configure the value to 0, the expiry is infinity.                                          |

*Table continues...*

| Name                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>BpduGuardAdminEnabled</b>               | Enables BPDU Guard on the port. The default is disabled.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>ForwardErrorCorrection</b>              | <p>Configures one of the following options for Forward Error Correction (FEC) on the port:</p> <ul style="list-style-type: none"> <li>• CL 91</li> <li>• CL 108</li> <li>• CL 74</li> <li>• disable</li> <li>• auto</li> </ul> <p>The disable option disables this configuration on the port.</p>                                                                                                                                 |
| <b>ForwardErrorCorrectionApplicability</b> | Displays whether FEC is applicable on the interface.                                                                                                                                                                                                                                                                                                                                                                              |
| <b>OperAutoNegotiate</b>                   | Shows the operational state of Auto-Negotiation.                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>OperForwardErrorCorrection</b>          | <p>Shows the negotiated operational FEC clause.</p> <p>If the value is off, the port supports FEC and is up but not configured for FEC. If the value is notApplicable, the port does not support FEC. If the value is unknown, the port supports FEC but is down.</p>                                                                                                                                                             |
| <b>IsPortShared</b>                        | <p>Indicates whether the port is combo or not.</p> <ul style="list-style-type: none"> <li>• portShared—Combo port.</li> <li>• portNotShared—Not a combo port.</li> </ul>                                                                                                                                                                                                                                                          |
| <b>PortActiveComponent</b>                 | <p>Specifies whether the copper port is active or fabric port is active if port is a combo port.</p> <ul style="list-style-type: none"> <li>• fixed port—Copper port is active.</li> <li>• gbic port—Fabric port is active.</li> </ul>                                                                                                                                                                                            |
| <b>Action</b>                              | <p>Performs one of the following actions on the port</p> <ul style="list-style-type: none"> <li>• none - none of the following actions</li> <li>• flushMacFdb - flush the MAC forwarding table</li> <li>• flushArp - flush the ARP table</li> <li>• flushIp - flush the IP route table</li> <li>• flushAll - flush all tables</li> <li>• triggerRipUpdate — manually triggers a RIP update</li> </ul> <p>The default is none.</p> |

*Table continues...*

| Name   | Description                                                      |
|--------|------------------------------------------------------------------|
| Result | Displays the result of the selected action. The default is none. |

## Configure Boot Flags

### About this task

Change the boot configuration to determine the services available after the system starts.

### Procedure

1. In the navigation pane, expand **Configuration > Edit > Chassis**.
2. Select the **Boot Config** tab.
3. Select the services you want to enable.
4. Select **Apply**.





## Boot Config Field Descriptions

Use the data in the following table to use the Boot Config tab.





| Name                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SwVersion                 | Specifies the software version that currently runs on the chassis.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| LastRuntimeConfigSource   | Specifies the last source for the run-time image.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| PrimaryConfigSource       | Specifies the primary configuration source.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| PrimaryBackupConfigSource | Specifies the backup configuration source to use if the primary does not exist.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| EnableFactoryDefaultsMode | <p>Specifies whether the switch uses the factory default settings at startup.</p> <ul style="list-style-type: none"> <li>• <b>false</b>: The node does not use factory default settings at startup.</li> <li>• <b>fabric</b>: The node uses the factory default fabric mode settings at startup. Zero Touch Fabric Configuration is enabled.</li> <li>• <b>noFabric</b>: The node uses the factory default mode settings at startup.</li> </ul> <p>The default value is false. This flag is automatically reset to the default setting after the switch restarts. If you change this parameter, you must restart the switch for the change to take effect.</p> |
| EnableDebugMode           | Enabling the debugmode will provide the opportunity to allow user to enable TRACE on any port by                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

*Table continues...*







| Name                                                                                                                                                                                                                                                                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                                                                                                                                                                                                                | <p>prompting the selection on the console during boot up. This allows the user start trace for debugging earlier on specified port. It only works on console connection. By default, it is disabled.</p> <p> <b>Important:</b><br/>Do not change this parameter.</p>                                                                      |
| <b>EnableRebootOnError</b>                                                                                                                                                                                                                                                     | <p>Activates or disables automatic reboot on a fatal error. The default value is activated.</p> <p> <b>Important:</b><br/>Do not change this parameter.</p>                                                                                                                                                                               |
| <b>EnableTelnetServer</b>                                                                                                                                                                                                                                                      | <p>Activates or disables the Telnet server service. The default is disabled.</p>                                                                                                                                                                                                                                                                                                                                           |
| <b>EnableRloginServer</b>                                                                                                                                                                                                                                                      | <p>Activates or disables the rlogin and rsh server. The default value is disabled.</p>                                                                                                                                                                                                                                                                                                                                     |
| <b>EnableFtpServer</b>                                                                                                                                                                                                                                                         | <p>Activates or disables the FTP server on the switch. The default value is disabled. To enable FTP, ensure that the TFTPD flag is disabled.</p>                                                                                                                                                                                                                                                                           |
| <b>EnableTftpServer</b>                                                                                                                                                                                                                                                        | <p>Activates or disables Trivial File Transfer Protocol server service. The default value is disabled.</p>                                                                                                                                                                                                                                                                                                                 |
| <b>EnableSshServer</b>                                                                                                                                                                                                                                                         | <p>Activates or disables the SSH server service. The default value is disabled.</p>                                                                                                                                                                                                                                                                                                                                        |
| <b>EnableSpbmConfigMode</b>                                                                                                                                                                                                                                                    | <p>Enables you to configure SPB and IS-IS, but you cannot configure PIM and IGMP either globally or on an interface.</p> <p>The boot flag is enabled by default.</p>                                                                                                                                                                                                                                                       |
| <p><b>EnableIpv6Mode</b></p> <p> <b>Note:</b><br/>Exception: only supported on VSP 4900 Series VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and VSP 8600 Series.</p> | <p>Enable this flag to support IPv6 routes with prefix-lengths greater than 64 bits. This flag is disabled by default.</p>                                                                                                                                                                                                                                                                                                 |
| <b>EnableEnhancedsecureMode</b>                                                                                                                                                                                                                                                | <p>Enables or disables the enhanced secure mode. Select either <b>jitc</b> or <b>non-jitc</b> to enable the enhanced secure mode in one of these sub-modes. The default is disabled.</p> <p> <b>Note:</b><br/>It is recommended that you enable the enhanced secure mode in the non-JITC sub-mode because the JITC sub-mode is more</p> |

*Table continues...*

| Name                                                                                                                                                                                                                                                                                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                                                                                                                                                                                                                         | restrictive and prevents the use of some troubleshooting utilities.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>EnableUrpMode</b>                                                                                                                                                                                                                                                                    | Enables Unicast Reverse Path Forwarding (uRPF) globally. You must enable uRPF globally before you configure it on a port or VLAN. The default is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>EnableVxlanGwFullInterworkingMode</b><br> <b>Note:</b><br>Exception: only supported on VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, and VSP 8400 Series.                                      | Enables VXLAN Gateway in Full Interworking Mode, which supports SPB, SMLT, and vIST.<br><br>By default, the Base Interworking Mode is enabled and Full Interworking Mode is disabled. You change modes by enabling this boot configuration flag.<br><br>In Base Interworking Mode, VXLAN Gateway supports Layer 2 gateway communication between VXLAN and traditional VLAN environments.                                                                                                                                                                                                                                                                    |
| <b>EnableFlowControlMode</b><br> <b>Note:</b><br>Exception: only supported on VSP 4000 Series, VSP 4900 Series, VSP 7200 Series, VSP 7400 Series, VSP 8200 Series, VSP 8400 Series, and XA1400 Series. | Enables or disables flow control globally. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames.<br><br>The default is disabled.                                                                                                                                                                                                                                                                                       |
| <b>AdvancedFeatureBwReservation</b><br> <b>Note:</b><br>Exception: only supported on VSP 7400 Series and XA1480.                                                                                     | Enables the switch to support advanced features. The default is enabled with low level configuration.<br><br>The high level means that the switch reserves the maximum bandwidth for the advanced features. The low level means that the switch reserves less bandwidth to support minimum functionality for advanced features.<br><br>If you change this parameter, you must restart the switch.<br><br>You must ensure your configuration does not include reserved ports before you enable this feature. If the configuration includes reserved ports after you enable this feature and restart the switch, the switch aborts loading the configuration. |
| <b>InsightPortConnectType</b><br> <b>Note:</b><br>Exception: only supported on VSP 7400-48Y.                                                                                                         | Determines the connection type the Insight port can use with virtual machine (VM) virtual ports. The default is vtd.<br><br>The VT-d connection type supports only one VM virtual port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

*Table continues...*

| Name                                                                                                                                                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|                                                                                                                                                                               | If you change this parameter, the switch automatically saves the configuration and restarts.                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>EnableDvrLeafMode</b>                                                                                                                                                      | Enables the switch to be configured as a DvR Leaf.<br><br>When enabled, you cannot configure the switch to operate as a DvR Controller.                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>EnablevrfScaling</b>                                                                                                                                                       | Changes the maximum number of VRFs and Layer 3 VSNs that the switch supports. If you select this check box, the maximum number increases. The default is disabled.<br><br> <b>Important:</b><br><br>If you select both this check box and the <b>EnableSpbmConfigMode</b> check box, the switch reduces the number of configurable VLANs. For more information about maximum scaling numbers, see <a href="#">Release Notes for VOSS</a> . |
| <b>EnableSyslogRfc5424Format</b>                                                                                                                                              | Enables or disables the RFC 5424 syslog format.<br><br>The default is enabled. If the pre-existing configuration file is for a release prior to this enhancement, then the flag is disabled automatically.                                                                                                                                                                                                                                                                                                                  |
| <b>NniMstp</b>                                                                                                                                                                | Enables MSTP, and allows non SPBM B-VLAN configuration on SPBM NNI ports. The default is disabled.<br><br> <b>Note:</b><br><br>Spanning Tree is disabled on all SPBM NNIs. You cannot add an SPBM NNI port or MLT port to any non SPBM B-VLAN.                                                                                                                                                                                           |
| <b>EnableIpv6EgressFilterMode</b>                                                                                                                                             | Enables IPv6 egress filters. The default is disabled.<br><br>If you change this parameter, you must restart the switch.                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>MasterCPUSlot</b><br><br> <b>Note:</b><br>Exception: only supported on VSP 8600 Series. | Specifies the slot number, either 1 or 2, for the master CPU. The default value is 1.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>EnableHaCpu</b><br><br> <b>Note:</b><br>Exception: only supported on VSP 8600 Series.   | Enables or disables the CPU High Availability feature.<br><br>If you enable or disable HA mode, the secondary CPU automatically resets to load settings from the previously-saved configuration file. The default is enabled.                                                                                                                                                                                                                                                                                               |

*Table continues...*

| Name                                                                                          | Description                                                                                                    |
|-----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------|
| <b>EnableSavetoStandby</b><br>* <b>Note:</b><br>Exception: only supported on VSP 8600 Series. | Enables or disables automatic save of the configuration file to the standby CPU. The default value is enabled. |
| <b>Slot</b>                                                                                   | Specifies the slot number.                                                                                     |
| <b>TftpHash</b>                                                                               | Enables TFTP hashing.                                                                                          |
| <b>TftpRetransmit</b>                                                                         | Set TFTP retransmit timeout counter.                                                                           |
| <b>TftpTimeout</b>                                                                            | Set TFTP timeout counter.                                                                                      |
| <b>User</b>                                                                                   | Configure host user.                                                                                           |
| <b>Password</b>                                                                               | Configure host password.                                                                                       |

## Configure IEEE 802.3X Pause frame transmit

Configure IEEE 802.3X Pause frame transmit to eliminate or minimize packet loss.

### About this task

By default, flow control mode is disabled. When disabled, the system does not generate nor configure the transmission of flow control messages. The system always honors received flow control messages regardless of the flow control mode status. You must enable this mode before you configure an interface to send pause frames.

By default, an interface does not send pause frames.

### Procedure

1. In the navigation pane, expand the **Configuration > Edit** folders.
2. Click **Chassis**.
3. Click the **Boot Config** tab.
4. For EnableFlowControlMode, select **enable**.
5. Click **Apply**.
6. Save the switch configuration.
7. Reboot the chassis, and log in again.
8. In the Device Physical View, select a port or ports.
9. In the navigation pane, expand the **Configuration > Edit > Port** folders.
10. Click **General**.
11. Click the **Interface** tab.
12. For TxFlowControl, select **enable** to enable the interface to generate pause frames.
13. Click **Apply**.

## Filenames for this Release

**! Important:**

Do not use Google Chrome or Safari to download software files. Google Chrome can change the file sizes. Safari changes the .tgz extension to .tar.

After you download the software, calculate and verify the md5 checksum. For more information, see [Administering VOSS](#).

The following table provides the filenames and sizes for this release.

**Table 12: Software Filenames and Sizes**

| Description                            | XA1400 Series                     | File size       |
|----------------------------------------|-----------------------------------|-----------------|
| SHA512 Checksum files                  | VOSS1400.8.1.50.0.sha512          | 1107 bytes      |
| MD5 Checksum files                     | VOSS1400.8.1.50.0.md5             | 435 bytes       |
| MIB - supported object names           | VOSS1400.8.1.50.0_mib_sup.txt     | 1067381 bytes   |
| MIB - zip file of all MIBs             | VOSS1400.8.1.50.0_mib.zip         | 1155302 bytes   |
| MIB - objects in the OID compile order | VOSS1400.8.1.50.0_mib.txt         | 7667943 bytes   |
| Open source software notice            | VOSS1400.8.1.50.0_oss-notice.html | 2766416 bytes   |
| EDM Help files                         | VOSS1400v8150_HELP_EDM_gzip.zip   | 4334109 bytes   |
| Logs reference                         | VOSS1400.8.1.50.0_edoc.tar        | 65945600 bytes  |
| Software image                         | VOSS1400.8.1.50.0.tgz             | 347110868 bytes |

The Open Source license text for the switch is included on the product. You can access it by entering the following command in the CLI:

```
more release/w.x.y.z.GA /release/oss-notice.txt
```

where *w.x.y.z* represents a specific release number.

## Documentation Changes

The Features by Release table has been removed from this document. Product support information for features is now described in product support tables at the beginning of each feature description throughout the documentation suite, and in the [VOSS Feature Support Matrix](#).

# Chapter 3: Upgrade and Downgrade Considerations

## Important:

*Notices for XA1400 Series.*

If your XA1400 Series runs VOSS 8.1.1, you must first downgrade to VOSS 8.1 before you can upgrade to VOSS 8.1.50. Ensure that you save and back up your existing configuration before and after the intermediate release 8.1 installation.

An upgrade from VOSS 8.1.1 to 8.1.50 is not supported.

The same restriction applies to downgrades from VOSS 8.1.1 and later to releases to VOSS 8.1.50 for XA1400 Series. You must first downgrade to VOSS 8.1 before upgrading to 8.1.50.

## Important:

If you downgrade from 8.1.50 to a prior release, you must ensure the config file does not contain command syntax for the FE tunnel MTU, or IPsec NAT-T, or other features only available in 8.1.50 and later releases.

See the [Administering VOSS](#) document for detailed image management procedures that includes information about these specific upgrade considerations:

- Fabric:
  - Pre-upgrade instructions for IS-IS metric type
  - Upgrade considerations for IS-IS enabled links with HMAC-MD5 authentication
  - The following releases included modified Zero Touch Fabric Configuration support that impacts upgrades from earlier releases: VOSS 8.1 and later.
- Upgrade considerations regarding MACsec replay-protect configuration
- Upgrade support for the nni-mstp boot configuration flag
- TACACS+ upgrade consideration
- Considerations for VLANs or MLTs where the VLAN or MLT name uses all numbers.
- Considerations for digital certificates configured prior to VOSS 8.1.

If your configuration includes one of the preceding scenarios or features, read the upgrade information in [Administering VOSS](#) before you begin an image upgrade.

# Chapter 4: XA1400 Series Hardware

| Part number | Model number                             | Initial release | Supported new feature release |     |       |        |
|-------------|------------------------------------------|-----------------|-------------------------------|-----|-------|--------|
|             |                                          |                 | 8.0.50                        | 8.1 | 8.1.1 | 8.1.50 |
| XA1440      | ExtremeAccess Platform 1440<br>(XA1440)  | 8.0.50          | Y                             | Y   | Y     | Y      |
| XA1480      | ExtremeAccess Platform 1480<br>(XA1480 ) | 8.0.50          | Y                             | Y   | Y     | Y      |

# Chapter 5: Software Scaling

This section lists software scaling capabilities for the XA1400 Series.

---

## Layer 2

**Table 13: Layer 2 Maximums**

|                                     |                                      |
|-------------------------------------|--------------------------------------|
| LACP aggregators                    | 8                                    |
| Layer 2 VSNs                        | 124                                  |
| MAC table size                      | 2,000 for XA1440<br>4,000 for XA1480 |
| Microsoft NLB cluster IP interfaces | N/A                                  |
| MLT groups                          | 8                                    |
| MSTP instances                      | 12                                   |
| Port-based VLANs                    | 500                                  |
| Ports per LACP aggregator           | 8                                    |
| Ports per MLT group                 | 8                                    |
| RSTP instances                      | 1                                    |
| SLPP VLANs                          | 128                                  |
| VLACP interfaces                    | 8                                    |

---

## IP Unicast

**Table 14: IP Unicast Maximums**

|                                      |       |
|--------------------------------------|-------|
| BGP+ peers                           | N/A   |
| DHCP Relay forwarding entries (IPv4) | 128   |
| ECMP groups/paths per group          | 500/8 |

*Table continues...*



|                                                    |                                    |
|----------------------------------------------------|------------------------------------|
| IP interfaces (IPv4)                               | 500                                |
| IPv4 ARP table                                     | 2000 for XA1440<br>4000 for XA1480 |
| IPv4 BGP peers                                     | 12                                 |
| IPv4 CLIP interfaces                               | 64                                 |
| IPv4 RIP interfaces                                | 200                                |
| IPv4 route policies (per VRF/per switch)           | 500/5,000                          |
| IPv4 static ARP entries (per VRF/per switch)       | 200/1,000                          |
| IPv4 static routes (per VRF/per switch)            | 1,000/5,000                        |
| IPv4 UDP forwarding entries                        | 128                                |
| IPv4 VRF instances                                 | 24 including GRT                   |
| IPv6 CLIP interfaces                               | N/A                                |
| IPv6 Ingress ACEs (Security and QoS)               | N/A                                |
| IPv6 Neighbor table                                | N/A                                |
| IPv6 OSPFv3 routes - GRT only                      | N/A                                |
| IPv6 RIPng peers                                   | N/A                                |
| IPv6 RIPng routes                                  | N/A                                |
| IPv6 Route Table size                              | N/A                                |
| IPv6 static neighbor records                       | N/A                                |
| IPv6 static routes                                 | N/A                                |
| Layer 3 VSNs                                       | 23                                 |
| OSPFv2 interfaces                                  | 48                                 |
| OSPF v2 neighbors (adjacencies)                    | 24                                 |
| OSPF v2 areas (per VRF/per switch)                 | 12/64                              |
| OSPF v3 areas                                      | N/A                                |
| Routed Split Multi-LinkTrunking (RSMLT) interfaces | N/A                                |
| VRRP interfaces (IPv4)                             | 64                                 |
| VRRP interfaces with fast timers (200ms)           | 24                                 |
| VRRP VRIDs                                         | 8                                  |
| Manually configured 6-in-4 tunnels                 | N/A                                |

## Layer 3 Route Table Size

**Table 15: Layer 3 Route Table Size Maximums**

|                                      |        |
|--------------------------------------|--------|
| IPv4 BGP routes (control plane only) | 15,488 |
| IPv4 OSPF routes                     | 15,488 |
| IPv4 RIP routes                      | 15,488 |
| IPv4 routes                          | 15,488 |
| IPv4 SPB Shortcut routes             | 15,488 |

## IP Multicast

**Table 16: IP Multicast Maximums**

|                                                        |     |
|--------------------------------------------------------|-----|
| IGMP interfaces                                        | N/A |
| PIM interfaces (Active/Passive )                       | N/A |
| Multicast receivers/IGMP receiver entries (per switch) | N/A |
| Multicast senders/IGMP sender entries (per switch)     | N/A |
| PIM-SSM static channels                                | N/A |
| Total multicast routes (S,G,V) (per switch)            | N/A |

**\* Note:**

IPv4 Routes, IPv4 SGV sender records, IPv6 Routes and IPv6 neighbor records reside in the same shared hardware table. If records of all 4 types are present together in this shared table, then the actual numbers that can be supported might be less than the scaling numbers indicated in the above tables.

## Filters, QoS, and Security

**Table 17: Filters, QoS, and Security Maximums**

|                     |     |
|---------------------|-----|
| Total ACE - Ingress | 500 |
| Total ACE - Egress  | 500 |
| Total ACL - Ingress | 500 |
| Total ACL - Egress  | 500 |

## Fabric Scaling

**Table 18: Fabric Scaling Maximums**

|                                                                                                                                             |       |
|---------------------------------------------------------------------------------------------------------------------------------------------|-------|
| Number of SPB regions                                                                                                                       | 1     |
| Number of B-VIDs                                                                                                                            | 2     |
| Number of SPB adjacencies                                                                                                                   | 64    |
| SPBM enabled nodes per region (BEB + BCB)                                                                                                   | 550*  |
| * <b>NOTE</b> : If there are VSP 4000 switches in the network, then the total number of SPBM enabled switches per region is reduced to 550. |       |
| Maximum number of IP multicast S,Gs when operating as a BCB                                                                                 | 2000  |
| Number of BEBs this node can share services with (Layer 2 VSNs, Layer 3 VSNs)                                                               | N/A** |
| ** <b>NOTE</b> : vIST clusters are counted as 3 nodes.                                                                                      |       |
| Maximum number of SPB Layer 2 multicast UNI I-SIDs                                                                                          | N/A   |
| Maximum number of SPB Layer 3 multicast UNI I-SIDs                                                                                          | N/A   |
| Maximum number of IP multicast S,Gs when operating as a BCB                                                                                 | 2000  |

## OAM and Diagnostics

**Table 19: OAM and Diagnostics Maximums**

|                            |     |
|----------------------------|-----|
| EDM sessions               | 5   |
| FTP sessions               | 4   |
| Mirrored destination ports | 4   |
| Mirrored source ports      | 7   |
| Rlogin sessions            | 8   |
| sFlow sampling rate        | N/A |
| SSH sessions               | 8   |
| Telnet sessions            | 8   |

# Chapter 6: Important Notices

Unless specifically stated otherwise, the notices in this section apply to all XA1400 Series platforms.

---

## Subscription Licensing for XA1400 Series

Each XA1400 Series device requires a subscription license.

Licenses are tied to the switch Base MAC address and switch model type. After you generate the license through Extreme Networks Support Portal at <https://extremeportal.force.com/ExtrLicenseLanding>, you can install the license on the switch.

### **Note:**

VOSS Release 8.0.50 or later is required to support subscription licenses generated through the Extreme Networks Support Portal.

The following sections detail the different categories of licenses supported on the XA1400 Series switch.

### **Factory Default Trial License**

A new switch includes a 60-day Factory Default Trial License starting from the time the switch is first booted. You can configure all features (except MACsec), without restrictions and save the configuration. No license file is required.

The system generates warning messages to inform you about the time remaining in the license period. The alerts appear once every 5 days for the first 55 days, and then once daily for the last 5 days. If you reboot the switch after the 60-day period, and a valid software license is not present, the licensed features in the configuration are not loaded. You must install a valid license to enable the licensed features.

### **Subscription License**

All subscription licenses support all VOSS features on the switch, plus software upgrades and technical support services entitlement during the license term. A one, three, or five year subscription license is required for each XA1400 Series device. Three services entitlement tiers of license are available: ExtremeWorks, PartnerWorks, and ExtremeWorks Premier.

A Subscription License is available in two bandwidth tiers of licenses: Small License and Medium License. A Small License enables up to 100 Mbps aggregate throughput Fabric Extend WAN tunneling connectivity, and a Medium License enables up to 500 Mbps aggregate throughput Fabric Extend WAN tunneling connectivity.

License expiry notifications are sent to the console and management station every 30 days until the last 30 days of the subscription. Then every 5 days until the last 9 days of the subscription, and then daily until the Subscription License expires.

Once the Subscription License expires, you get a limited 30 day grace period. When a subscription expires, notification messages are shown as the grace period counts down. Messages are shown on the console and in the alarms database indicating that the license is expired. If the system reboots after a license expiration, the grace period immediately ends and the system does not load or support any saved configurations or software services. License expired messages continue to show on the console and in the alarms database until a valid subscription license is installed.

**! Important:**

The 30 day grace period is lost if the system reboots after a license expires. You must renew your Subscription License to allow the software features to continue to function.

## XA1400 Series License Types and Part Numbers

The following table provides the part numbers for the various licenses the XA1400 Series supports.

**Table 20: Supported licenses**

| Small Subscription Licenses (up to 100 Mbps) | Part number/ Order code |
|----------------------------------------------|-------------------------|
| 1 year, ExtremeWorks                         | FCVPN-100-EW-1Y         |
| 1 year, PartnerWorks                         | FCVPN-100-PW-1Y         |
| 1 year, ExtremeWorks Premier                 | FCVPN-100-EWP-1Y        |
| 3 years, ExtremeWorks                        | FCVPN-100-EW-3Y         |
| 3 years, PartnerWorks                        | FCVPN-100-PW-3Y         |
| 3 years, ExtremeWorks Premier                | FCVPN-100-EWP-3Y        |
| 5 years, ExtremeWorks                        | FCVPN-100-EW-5Y         |
| 5 years, PartnerWorks                        | FCVPN-100-PW-5Y         |
| 5 years, ExtremeWorks Premier                | FCVPN-100-EWP-5Y        |

| Medium Subscription Licenses (up to 500 Mbps) | Part number/ Order code |
|-----------------------------------------------|-------------------------|
| 1 year, ExtremeWorks                          | FCVPN-500-EW-1Y         |
| 1 year, PartnerWorks                          | FCVPN-500-PW-1Y         |
| 1 year, ExtremeWorks Premier                  | FCVPN-500-EWP-1Y        |
| 3 years, ExtremeWorks                         | FCVPN-500-EW-3Y         |
| 3 years, PartnerWorks                         | FCVPN-500-PW-3Y         |
| 3 years, ExtremeWorks Premier                 | FCVPN-500-EWP-3Y        |
| 5 years, ExtremeWorks                         | FCVPN-500-EW-5Y         |
| 5 years, PartnerWorks                         | FCVPN-500-PW-5Y         |
| 5 years, ExtremeWorks Premier                 | FCVPN-500-EWP-5Y        |

 **Note:**

500 Mbps Subscription Licenses are only supported on XA1480 devices.

---

## Limitations on license filename size

When you dynamically load a named license file, ensure that the file name has a maximum of 42 characters *including* the .xml extension. In other words, the length of the file name must be less than or equal to 42 characters, including the extension.

Otherwise, the license file does not load successfully on system reboot.

---

## Supported Browsers

Use the following browser versions to access Enterprise Device Manager (EDM):

- Microsoft Edge 41+
- Microsoft Internet Explorer 11.0+
- Mozilla Firefox 58.0+
- Google Chrome 64+

# Chapter 7: Known Issues and Restrictions

This section details the known issues and restrictions found in this release. Where appropriate, use the workarounds provided.

---

## Known Issues and Restrictions

This chapter details the known issues and restrictions found in this release. Where appropriate, use the workarounds provided.

### General Issues and Restrictions

| Issue number | Description                                                                                                                                                                                                                                                                                                                                                                  | Workaround                                                                                                        |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| VOSS-13463   | Out port statistics for MLT port interfaces are not accurate.                                                                                                                                                                                                                                                                                                                | Use the command <code>show io nic-counters</code> to display detailed port stats and error info on XA1400 Series. |
| VOSS-13680   | Interface error statistics display is inaccurate in certain scenarios.                                                                                                                                                                                                                                                                                                       | Use the command <code>show io nic-counters</code> to display detailed port stats and error info on XA1400 Series. |
| VOSS-13681   | QoS: <code>show qos cosq-stats cpu-port</code> command output is not supported.                                                                                                                                                                                                                                                                                              | Use the command <code>show io cpu-cosq-counters</code> to display detailed cosq-stats on XA1400 Series.           |
| VOSS-14150   | CLI remote console might stop wrapping text after some usage.                                                                                                                                                                                                                                                                                                                | Reset the CLI window or open a new remote console window.                                                         |
| VOSS-14494   | Layer 2 VSN and Layer 3 VSN UNI to NNI traffic between two Backbone Edge Bridges does not hash to different ports of a MLT network-to-network interface. MLT hashing for XA1400 devices occurs after the mac-in-mac encapsulation is done. The hash keys used are the Backbone destination and Backbone source MAC addresses (BMAC DA and BMAC SA) in the Mac-in-Mac header. | None.                                                                                                             |

*Table continues...*

| Issue number             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Workaround                                                                                                 |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
|                          | Even for the Transit BCB case on XA 1400 devices for NNI to NNI traffic, the MLT hash keys used are the Backbone destination and Backbone source MAC addresses (BMAC DA and BMAC SA) in the Mac-in-Mac header.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |                                                                                                            |
| VOSS-14515               | <p>Console output errors and warnings are shown during an XA1400 Series reboot, such as:</p> <ul style="list-style-type: none"> <li>• error: no such device: ((hd0,gpt1)/EFI/BOOT)/EFI/BOOT/grub.cfg.</li> <li>error: file `/EFI/BOOT/grubenv' not found</li> <li>• error: no suitable video mode found.</li> <li>• vfio-pci 0000:05:00.0: Invalid PCI ROM header signature: expecting 0xaa55, got 0xbeef</li> <li>• [0.727012] ACPI: No IRQ available for PCI Interrupt Link [LNKS]. Try pci=noacpi or acpi=off</li> <li>• exportfs: can't open /etc/exports for reading</li> <li>• KCORE: WARNING can't find /boot/b/ulmage-gemini.bin. No kexec kernel will be configured.</li> </ul> | None. The errors or warnings are host OS or guest OS related with no functional impact and can be ignored. |
| VOSS-14590               | ISIS logical-interface displays the same egress port for different tunnels when the underlay reachability is from different port interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | None.                                                                                                      |
| VOSS-14597               | Ping (originated from local CP) fails for jumbo frames on Layer 3 VSN interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | None.                                                                                                      |
| VOSS-14616               | <p>Seeing Queue buffer usage logs when changing the logical interface source IP with 64 tunnels.</p> <p>When changing the source IP with 64 tunnels, seeing "GlobalRouter CPU INFO CPP: 60 percent of fbufs are in use: 0 in Tx queue,1843 in RxQueue0 0 in RxQueue1 0 in RxQueue2 0 in RxQueue3 0 in RxQueue4 0 in RxQueue5 0 in RxQueue6 0 in RxQueue7 ".</p>                                                                                                                                                                                                                                                                                                                          | None.                                                                                                      |
| VOSS-14656               | Console output "ErrLog: Error Level=2 [(null)] seen during OpenVas testing. No functional impact.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | None.                                                                                                      |
| VOSS-14805<br>VOSS-15305 | <p>The following transceivers are not supported on XA1400 Series switches:</p> <ul style="list-style-type: none"> <li>• 10 Gb Bidirectional 40 km SFP+ Module (10GB-BX40-D and 10GBBX40-U)</li> <li>• 1000BASE-BX10 Bidirectional 10 km DDI SFP Modules (AA1419069-E6 and AA1419070-E6)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                       | Use only supported transceivers.                                                                           |
| VOSS-15463               | XA1440 and XA1480 switches may experience intermittent Link Up and Link Down transitions on                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | No workaround, but there is no functional impact.                                                          |

Table continues...



| Issue number | Description                                                                                                                                                                                                                                                                                                                                                                                                                 | Workaround                                                                                                                                                                                              |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|              | the 10/100/1000BASE-T Ethernet ports upon booting.                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                         |
| VOSS-16221   | Layer 2 ping is not working for packets larger than 1300 on an XA1400 Series.                                                                                                                                                                                                                                                                                                                                               | Use Layer 2 ping with packets smaller than 1300 bytes.                                                                                                                                                  |
| VOSS-16365   | Running the command <code>show pluggable-optical-module detail</code> on an XA1400 Series device is highly CPU intensive to read and reply with the EEPROM details. Due to a delay in ethtool response, a watchdog miss event can occur and the event is recorded in the <code>/intflash/wd_stats/1/wd_stats.ssio.1.log</code> file. This scenario occurs more often if 10Gb SFP+ optics with DDM capability are installed. | None. The high CPU usage and response delay for this command is expected and cannot be resolved. No console log is generated. When the scenario occurs, the Watchdog outage is approximately 5 seconds. |
| VOSS-16436   | Using the console connection on an XA1400 Series device while running a show command with large data output can result in drops of processing control packets.                                                                                                                                                                                                                                                              | Use Telnet or SSH connectivity instead of console connection.                                                                                                                                           |

---

## Filter Restrictions and Expected Behaviors

This section lists known restrictions and expected behaviors that may first appear to be issues.

The following list describes the expected behavior with filters:

- ACL: InVlan ACLs can match tagged or untagged traffic, with the port-default VLAN considered if the incoming packet is untagged. However, if an ACE of an InVlan ACL contains the qualifier `vlan-tag-prio`, it can be used to filter only tagged traffic and not the untagged traffic.
- ACL: The outPort ACLs cannot match on the fields that are changed in the packet during forwarding decisions. Hence, the fields (Destination MAC, Source MAC, VLAN ID, etc.), which get modified during Layer 3 routing, cannot be used to match on the new contents of these fields in the outgoing packet.
- ACL: The outPort ACLs cannot match on a destination port that is a member of an MLT. So if port 1/5 is a member of an MLT (static or via LACP), an ACE of an outPort filter with member 1/5 will not be hit.
- ACL: The outPort ACLs do not apply to mirrored ports.
- There can be a single ACE hit for a packet. Port-based ACLs have precedence over VLAN based ACLs. However, the default ACEs have a lower priority than the user ACEs.
  1. User ACE of InPort ACL
  2. User ACE of InVlan ACL
  3. Default ACE of InPort ACL

#### 4. Default ACE of InVlan ACL

**\* Note:**

If a packet matches a user ACE in both an inPort and inVLAN ACL, the inVLAN ACL is ignored.

If a packet matches a user ACE in VLAN-based ACL and the default ACE of an inPort ACL, the user ACE in the inVLAN ACL is hit and the inPort ACL is ignored.

- ACL: The monitor actions (monitor-dst-port or monitor-dst-mlt) are not supported for outPort ACLs. They are only applicable to Ingress ACLs (InPort or InVlan). For flow-based mirroring, you can configure these monitor actions at the ACE level. ACL global mirroring action is not supported.
- ACE: When an ACE with action count is disabled, the statistics associated with the ACE are reset.
- For ACEs of port-based ACLs, you can configure VLAN qualifiers. Configuring Port qualifiers are not permitted.

For ACEs of VLAN-based ACLs, you can configure port qualifiers. Configuring VLAN qualifiers are not permitted.

### Filters and QoS

Note the following filters:

- XA1400 Series does not support the following qualifiers in the egress direction (outPort). However, ingress support (inVlan/InPort) for these qualifiers are available.
  - `arprequest` and `arpresponse`
  - `ip-frag-flag`
  - `tcp-flags`
- The `ip-options` qualifier is not supported.

For more information, see [Configuring QoS and ACL-Based Traffic Filtering for VOSS](#).

# Chapter 8: Resolved Issues

This section details the issues that are resolved in this release.

## Fixes from Previous Releases

VOSS 8.1.50 incorporates all fixes from prior releases, up to and including VOSS 8.1.

## Resolved Issues in VOSS 8.1.50

| Issue number | Description                                                                                                                                                                                         |
|--------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VOSS-14592   | Operation down log messages display on console for an already shutdown port. Log messages are printed whenever a shut CLI command executes even if the port was previously shutdown.                |
| VOSS-14639   | Packets ingressing on the front panel ports with source mac as switch's local VLAN MAC are forwarded instead of being dropped.                                                                      |
| VOSS-15525   | If you enable the <code>filter-untagged-frame</code> option before you enable the <code>untag-port-default-vlan</code> option, the default VLAN untagging for the port does not function correctly. |

# Chapter 9: Related Information

The following section contains information related to the current release.

---

## MIB Changes

---

### Modified MIBs

| Object Name            | Object OID                     | Modification in Release 8.1.50   |
|------------------------|--------------------------------|----------------------------------|
| XA1400 Series:         |                                |                                  |
| rcPortIngressRateLimit | 1.3.6.1.4.1.2272.1.4.10.1.1.85 | Supports 10000000 maximum value. |

---

### New MIBs

| Object Name                                  | Object OID                      |
|----------------------------------------------|---------------------------------|
| rcIscisLogicalInterfaceIpssecResponderOnly   | 1.3.6.1.4.1.2272.1.63.26.21     |
| rcIscisLogicalInterfaceIpssecRemoteNatIPAddr | 1.3.6.1.4.1.2272.1.63.26.22     |
| rcIscisLogicalInterfaceMtu                   | 1.3.6.1.4.1.2272.1.63.26.1.17   |
| rcPrQosCosQueTunnelStatsTable                | 1.3.6.1.4.1.2272.1.202.1.1.1.21 |