Extreme™
Customer-Driven Networking

Extreme NetIron
Security Configuration Guide, 6.3.00a

Supporting NetIron OS 6.3.00a

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

Software Licensing

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

Preface

This section discusses the conventions used in this guide, ways to provide feedback, additional help, and other Extreme Networks® publications.

Conventions

This section discusses the conventions used in this guide.

Notes, cautions, and warnings

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

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

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

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

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

Text formatting conventions

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

Format	Description
bold text	Identifies command names.
	Identifies keywords and operands.
	Identifies the names of GUI elements.
	Identifies text to enter in the GUI.
italic text	Identifies emphasis.
	Identifies variables.
	Identifies document titles.

Format	Description
`Courier font`	Identifies CLI output.
	Identifies command syntax examples.

Command syntax conventions

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

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

Documentation and Training

To find Extreme Networks product guides, visit our documentation pages at:

Current Product Documentation	www.extremenetworks.com/documentation/
Archived Documentation (for earlier versions and legacy products)	www.extremenetworks.com/support/documentation-archives/
Release Notes	www.extremenetworks.com/support/release-notes
Hardware/Software Compatibility Matrices	https://www.extremenetworks.com/support/compatibility-matrices/
White papers, data sheets, case studies, and other product resources	https://www.extremenetworks.com/resources/

Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For more information, 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: 1-800-998-2408 (toll-free in U.S. and Canada) or +1 408-579-2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

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

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products

- A description of the failure

- A description of any action(s) already taken to resolve the problem

- A description of your network environment (such as layout, cable type, other relevant environmental information)

- Network load at the time of trouble (if known)

- The device history (for example, if you have returned the device before, or if this is a recurring problem)

- Any related RMA (Return Material Authorization) numbers

Subscribing 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 with your information (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. Click **Submit**.

Providing Feedback to Us

Quality is our first concern at Extreme Networks, and we have 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 especially want to know about:

- Content errors or confusing or conflicting information.

- Ideas for improvements to our documentation so you can find the information you need faster.

- Broken links or usability issues.

If you would like to provide feedback to the Extreme Networks Information Development team, you can do so in two ways:

- Use our short online feedback form at https://www.extremenetworks.com/documentation-feedback/.

- Email us at documentation@extremenetworks.com.

Please 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.

About This Document

What's new in this document

NOTE
The NetIron 6.3.00 release (the image files and the documentation) is no longer available from the Extreme Portal. New software features introduced in release 6.3.00 are included in release 6.3.00a.

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

For the complete list of supported features and the summary of enhancements and configuration notes for this release, refer to the *Extreme NetIron OS Release Notes*.

The following table describes new information added to this document for the NetIron OS 6.3.00 software release.

TABLE 1 Summary of enhancements in NetIron OS 6.3.00 release

Feature	Description	Described in
Regular expression support in RADIUS command authorization.	The Extreme-specific RADIUS attribute **foundry-command-string** now supports specifying a range of data for a CLI command.	Configuring Extreme-specific attributes on the RADIUS server on page 50

Supported hardware and software

End of Support for ExtremeSwitching CES 2000 Series devices

Beginning with NetIron OS 6.3.00a and later, the ExtremeSwitching CES 2000 Series devices are not supported. Refer to the End of Sale and End of Support page for additional information.

The hardware platforms in the following table are supported by this release of this guide.

TABLE 2 Supported devices

ExtremeRouting XMR Series	ExtremeRouting MLX Series	ExtremeRouting CER 2000 Series
XMR 4000	MLX-4	CER 2024C
XMR 8000	MLX-8	CER-RT 2024C
XMR 16000	MLX-16	CER 2024F
XMR 32000	MLX-32	CER-RT 2024F
	MLXe-4	CER 2048C
	MLXe-8	CER-RT 2048C
	MLXe-16	CER 2048CX
	MLXe-32	CER-RT 2048CX
		CER 2048F

TABLE 2 Supported devices (continued)

ExtremeRouting XMR Series	ExtremeRouting MLX Series	ExtremeRouting CER 2000 Series
		CER-RT 2048F
		CER 2048FX
		CER-RT 2048FX

Supported software

For the complete list of supported features and the summary of enhancements and configuration notes for this release, refer to the *Extreme NetIron Release Notes*.

Managing User Accounts

Securing access methods

The table below lists the management access methods available on the Extreme devices, how they are secured by default, and the ways in which they can be secured.

Access method	How the access method is secured by default	Ways to secure the access method
Serial access to the CLI	Not secured	Establish passwords for management privilege levels
		Establish username and password to log in to the console.
Access to the Privileged EXEC and CONFIG levels of the CLI	Not secured	Establish a password for Telnet access to the CLI
		Establish passwords for management privilege levels
		Set up local user accounts
		Configure TACACS or TACACS+ security
		Configure RADIUS security
Telnet access Telnet server is turned off by default.		Regulate Telnet access using ACLs
Allow Telnet access only from specific IP addresses		
Allow Telnet access only to clients connected to a specific VLAN		
Regulate telnet access using Management VRF.		
Disable Telnet access		
Establish a password for Telnet access		
Establish passwords for privilege levels of the CLI		
Set up local user accounts		
Configure TACACS or TACACS+ security		
Configure RADIUS security		
Secure Shell (SSH) access For more information on SSH, refer to *Extreme NetIron Layer 2 Switching Configuration Guide*.	Not configured	Configure DSA or RSA host keys

Access method	How the access method is secured by default	Ways to secure the access method
Disable SSH server.		
Password Authentication		
Public key authentication using client's public key (excludes use of username and password credentials)		
Regulate SSH access using ACLs		
Allow SSH access only from specific IP addresses		
Establish passwords for privilege levels of the CLI		
Set up local user accounts		
Configure TACACS or TACACS+ security		
Configure RADIUS security		
Web management access	SNMP read or read-write community strings Web server is turned off by default. **Note** : Web access is not allowed in CES 2000 Series and CER 2000 Series devices.	Regulate Web management access using ACLs
Allow Web management access only from specific IP addresses		
Allow Web management access only to clients connected to a specific VLAN		
Disable Web management access		
Configure SSL security for the Web Management Interface		
Set up local user accounts		
Establish SNMP read or read-write community strings for SNMP versions 1 and 2		
Configure AAA command for Web access		
Configure TACACS or TACACS+ security		
Configure RADIUS security		
SNMP (Brocade Network Advisor) access	SNMP read or read-write community strings and the password to the Super User privilege level NOTE SNMP read or read-write community strings are always required for SNMP access to the device. SNMP access is disabled by default.	Regulate SNMP access using ACLs
		Allow SNMP access only from specific IP addresses
		Disable SNMP access
		Allow SNMP access only to clients connected to a specific VLAN
		Establish passwords to management levels of the CLI
		Set up local user accounts
		Configure AAA command for SNMP access
		Establish SNMP read or read-write community strings
TFTP access	Not secured	Allow TFTP access only to clients connected to a specific VLAN
Secure Copy access	Secured access if SSH server is enabled	Configure DSA or RSA host keys

Restricting remote access to management functions

Access method	How the access method is secured by default	Ways to secure the access method
Disable SSH server.		
Password Authentication		
Public key authentication using client's public key		
(excludes use of username and password credentials)		
Regulate SSH access using ACLs		
Allow SSH access only from specific IP addresses		
Establish passwords for privilege levels of the CLI		
Set up local user accounts		
Configure TACACS or TACACS+ security		
Configure RADIUS security		

Restricting remote access to management functions

You can restrict access to management functions from remote sources, including Telnet, SSH, the Web Management Interface, and SNMP. The following methods for restricting remote access are supported:

- Using ACLs to restrict Telnet, SSH, Web Management Interface, or SNMP access.
- Allowing remote access only from specific IP addresses.
- Allowing remote access only to clients connected to a specific VLAN.
- Specifically disabling Telnet, SSH, Web Management Interface, or SNMP access to the device.
- Using Management VRF to restrict access from certain physical ports.

    NOTE
    The Web Management Interface is only supported on the XMR Series and MLX Series devices.

    NOTE
    If the display on the front panel of the Web Management Interface is distorted, manually click on the link to reset the display to normal.

Using ACLs to restrict remote access

You can use ACLs to control the following access methods to management functions on the Extreme device:

- Telnet access
- SSH access
- Web management access
- SNMP access

Follow the steps listed below to configure access control for these management access methods.

1. Configure an ACL with the IP addresses you want to allow to access the device. You can specify a numbered standard IPv4 ACL, or a named standard IPv4 ACL.

9035624-02   21

2.   Configure a Telnet access group, SSH access group, web access group, and SNMP community strings for SNMPv1, SNMPv2c or SNMPv3 user. Each of these configuration items accepts an ACL as a parameter. The ACL contains entries that identify the IP addresses that can use the access method.

The following sections present examples of how to secure management access using ACLs. Refer to "Access Control List" chapter "Configuring an IPv6 Access Control List" for more information on configuring ACLs.

> **NOTE**
> ACL filtering for remote management access is done in hardware.

Using an ACL to restrict Telnet access

To configure an ACL that restricts Telnet access to the device, enter commands such as the following:

```
device(config)# access-list 10 deny host 10.157.22.32
device(config)# access-list 10 deny 10.157.23.0 0.0.0.255
device(config)# access-list 10 deny 10.157.24.0 0.0.0.255
device(config)# access-list 10 deny 10.157.25.0/24
device(config)# access-list 10 permit any
device(config)# telnet access-group 10
device(config)# write memory
```

The commands configure ACL 10, then apply it as the access list for Telnet access. The device allows Telnet access to all IP addresses except those listed in ACL 10.

Syntax: [no] telnet access-group { *num* | *name* }

The *num* variable specifies the number of a standard IPv4 ACL, 1 - 99.

The *name* variable specifies the standard IPv4 access list name.

> **NOTE**
> ACLs for Telnet sessions will be applied only to inbound sessions.

To configure a more restrictive ACL, create permit entries and omit the **permit any** entry at the end of the ACL.

```
device(config)# access-list 10 permit host 10.157.22.32
device(config)# access-list 10 permit 10.157.23.0 0.0.0.255
device(config)# access-list 10 permit 10.157.24.0 0.0.0.255
device(config)# access-list 10 permit 10.157.25.0/24
device(config)# telnet access-group 10
device(config)# write memory
```

The ACL in the example permits Telnet access only from the IPv4 addresses in the **permit** entries and denies Telnet access from all other IP addresses.

Using an ACL to restrict SSH access

To configure an ACL that restricts SSH access to the device, enter commands such as the following:

```
device(config)# access-list 12 deny host 10.157.22.98
device(config)# access-list 12 deny 10.157.23.0 0.0.0.255
device(config)# access-list 12 deny 10.157.24.0/24
device(config)# access-list 12 permit any
device(config)# ssh access-group 12
device(config)# write memory
```

Syntax: [no] ssh access-group {*num* | *name* }

The *num* variable specifies the number of a standard IPv4 ACL, 1 - 99.

The *name* variable specifies the standard IPv4 access list name.

hese commands configure ACL 12, then apply the ACL as the access list for SSH access. The device denies SSH access from the IPv4 addresses listed in ACL 12 and permits SSH access from all other IP addresses. Without the last ACL entry for permitting all packets, this ACL would deny SSH access from all IP addresses.

> **NOTE**
> In this example, the **command ssh access-group**could have been used to apply the ACL configured in the example for Telnet access. You can use the same ACL multiple times.

Using an ACL to restrict Web management access

> **NOTE**
> The Web Management Interface is only supported on the XMR Series and MLX Series devices.

To configure an ACL that restricts Web management access to the device, enter commands such as the following:

```
device(config)# access-list 12 deny host 10.157.22.98
device(config)# access-list 12 deny 10.157.23.0 0.0.0.255
device(config)# access-list 12 deny 10.157.24.0/24
device(config)# access-list 12 permit any
device(config)# web access-group 12
device(config)# write memory
```

Syntax: [no] web access-group { *num* | *name* }

The *num* variable specifies the number of a standard IPv4 ACL, 1 - 99.

The *name* variable specifies the standard IPv4 access list name.

These commands configure ACL 12, then apply the ACL as the access list for Web management access. The device denies Web management access from the IP addresses listed in ACL 12 and permits Web management access from all other IP addresses. Without the last ACL entry for permitting all packets, this ACL would deny Web management access from all IP addresses.

Using ACLs to restrict SNMP access

To restrict SNMP access to the device using ACLs, enter commands such as the following.

> **NOTE**
> The syntax for using ACLs for SNMP access is different from the syntax for controlling Telnet, SSH, and Web management access using ACLs.

```
device(config)# access-list 25 deny host 10.157.22.98
device(config)# access-list 25 deny 10.157.23.0 0.0.0.255
device(config)# access-list 25 deny 10.157.24.0 0.0.0.255
device(config)# access-list 25 permit any
device(config)# access-list 30 deny 10.157.25.0 0.0.0.255
device(config)# access-list 30 deny 10.157.26.0/24
device(config)# access-list 30 permit any
device(config)# snmp-server community public ro 25
device(config)# snmp-server community private rw 30
device(config)# write memory
```

These commands configure ACLs 25 and 30, then apply the ACLs to community strings. ACL 25 is used to control read-only access using the "public" community string. ACL 30 is used to control read-write access using the "private" community string.

Syntax: **[no] snmp-server community string** { **ro** | **rw** } { *standard-acl-name* | *standard-acl-id* }

The *string* variable specifies the SNMP community string the user must enter to gain SNMP access.

The **ro** parameter indicates that the community string is for read-only ("get") access. The **rw** parameter indicates the community string is for read-write ("set") access.

The *standard-acl-name* or *standard-acl-id* variable specifies which ACL will be used to filter incoming SNMP packets.

The *standard-acl-id* variable specifies the number of a standard IPv4 ACL, 1 - 99.

The *standard-acl-name* variable specifies the standard IPv4 access list name.

> **NOTE**
> When **snmp-server community** is configured, all incoming SNMP packets are validated first by their community strings and then by their bound ACLs. Packets are permitted if no filters are configured for an ACL.

Defining the console idle time

By default, an Extreme device does not time out serial console sessions. A serial session remains open indefinitely until you close it. You can however define how many minutes a serial management session can remain idle before it is timed out.

To configure the idle time for a serial console session, use the following command.

```
device(config)# console timeout 120
```

Syntax: **[no] console timeout value**

Possible values: 0 - 240 minutes

Default value: 0 minutes (no timeout)

> **NOTE**
> The standard for the idle-timeout RADIUS attribute is for it to be implemented in seconds as opposed to the minutes that the Extreme device uses. If this attribute is used for setting idle time instead of this configuration, the value from the idle-timeout RADIUS attribute will be converted from seconds to minutes and truncated to the nearest minute.

Restricting remote access to the device to specific IP addresses

By default, an Extreme device does not control remote management access based on the IP address of the managing device. You can restrict remote management access to a single IP address for the following access methods:

- Telnet access
- Web management access
- SNMP access
- SSH access

In addition, if you want to restrict all three access methods to the same IP address, you can do so using a single command.

The following examples show the CLI commands for restricting remote access. You can specify only one IP address with each command. However, you can enter each command ten times to specify up to ten IP addresses.

> **NOTE**
> You cannot restrict remote management access using the Web Management Interface.

Restricting Telnet access to a specific IP address

To allow Telnet access to the Extreme device only to the host with IP address 10.157.22.39, enter the following command.

```
device(config)# telnet client 10.157.22.39
```

Syntax: [no] telnet client *ip-address*

Restricting SSH access to a specific IP address

To allow SSH access to the Extreme device only to the host with IP address 10.157.22.39, enter the following command.

```
device(config)# ip ssh client 10.157.22.39
```

Syntax: [no] ip ssh client *ip-address*

Restricting Web management access to a specific IP address

> NOTE
> The Web Management Interface is only supported on the XMR Series and MLX Series devices.

To allow Web management access to the Extreme device only to the host with IP address 10.157.22.26, enter the following command.

```
device(config)# web client 10.157.22.26
```

Syntax: [no] web client *ip-address*

Restricting SNMP access to a specific IP address

To allow SNMP access (which includes Brocade Network Advisor) to the Extreme device only to the host with IP address 10.157.22.14, enter the following command.

```
device(config)# snmp-client 10.157.22.14
```

Syntax: [no] snmp-client *ip-address*

Restricting all remote management access to a specific IP address

> NOTE
> The Web Management Interface is only supported on the XMR Series and MLX Series devices.

To allow Telnet, SSH, Web, and SNMP management access to the Extreme device only to the host with IP address 10.157.22.69, you can enter three separate commands (one for each access type) or you can enter the following command.

```
device(config)# all-client 10.157.22.69
```

Syntax: [no] all-client *ip-address*

Defining the Telnet idle time

You can define how many minutes a Telnet session can remain idle before it is timed out. An idle Telnet session is a session that is still sending TCP ACKs in response to keepalive messages from the device, but is not being used to send data.

To configure the idle time for a Telnet session, use the following command.

```
device(config)# telnet timeout 120
```

Syntax: [no] telnet timeout 0 - 240

Possible values: 0 - 240 minutes

Default value: 0 minutes (no timeout)

> **NOTE**
> The standard for the idle-timeout RADIUS attribute is for it to be implemented in seconds as opposed to the minutes that the Extreme device uses. If this attribute is used for setting idle time instead of this configuration, the value from the idle-timeout RADIUS attribute will be converted from seconds to minutes and truncated to the nearest minute.

Specifying the maximum login attempts for Telnet access

If you are connecting to the Extreme device using Telnet, the device prompts you for a username and password. By default, you have up to 4 chances to enter a correct username and password. If you do not enter a correct username or password after 4 attempts, the Extreme device disconnects the Telnet session.

You can specify the number of attempts a Telnet user has to enter a correct username and password before the Extreme device disconnects the Telnet session. For example, to allow a Telnet user up to 5 chances to enter a correct username and password, enter the following command.

```
device(config)# telnet login-retries 5
```

Syntax: [no] telnet login-retries number

You can specify from 0 - 5 attempts. The default is 4 attempts.

Restricting remote access to the device to specific VLAN IDs

You can restrict management access to an Extreme device to ports within a specific port-based VLAN. VLAN-based access control applies to the following access methods:

- Telnet access
- Web management access
- SNMP access
- TFTP access

By default, access is allowed for all the methods listed above on all ports. Once you configure security for a given access method based on VLAN ID, access to the device using that method is restricted to only the ports within the specified VLAN.

VLAN-based access control works in conjunction with other access control methods. For example, suppose you configure an ACL to permit Telnet access only to specific client IP addresses, and you also configure VLAN-based access control for Telnet access. In this case, the only Telnet clients that can access the device are clients that have one of the IP addresses permitted by the ACL and are connected to a port that is in a permitted VLAN. Clients who have a permitted IP address but are connected to a port in a VLAN that is not permitted still cannot access the device through Telnet.

Restricting Telnet access to a specific VLAN

To allow Telnet access only to clients in a specific VLAN, enter a command such as the following.

```
device(config)# telnet server enable vlan 10
```

The command configures the device to allow Telnet management access only to clients connected to ports within port-based VLAN 10. Clients connected to ports that are not in VLAN 10 are denied management access.

Syntax: [no] telnet server enable vlan vlan-id

Restricting Web management access to a specific VLAN

> NOTE
> The Web Management Interface is only supported on the XMR Series and MLX Series devices.

To allow Web management access only to clients in a specific VLAN, enter a command such as the following.

```
device(config)# web-management enable vlan 10
```

The command configures the device to allow Web management access only to clients connected to ports within port-based VLAN 10. Clients connected to ports that are not in VLAN 10 are denied management access.

Syntax: [no] web-management enable vlan vlan-id

Restricting SNMP access to a specific VLAN

To allow SNMP access only to clients in a specific VLAN, enter a command such as the following.

```
device(config)# snmp-server enable vlan 40
```

The command configures the device to allow SNMP access only to clients connected to ports within port-based VLAN 40. Clients connected to ports that are not in VLAN 40 are denied access.

Syntax: [no] snmp-server enable vlan vlan-id

Restricting TFTP access to a specific VLAN

To allow TFTP access only to clients in a specific VLAN, enter a command such as the following.

```
device(config)# tftp client enable vlan 40
```

The command in this example configures the device to allow TFTP access only to clients connected to ports within port-based VLAN 40. Clients connected to ports that are not in VLAN 40 are denied access.

Syntax: [no] tftp client enable vlan vlan-id

Enabling specific access methods

You can specifically enable the following access methods:

- Telnet access
- Web management access
- SNMP access

> NOTE
> If you do not enable Telnet access, you can access the CLI using a serial connection to the management module. If you do not enable SNMP access, you will not be able to use Brocade Network Advisor or third-party SNMP management applications.

Enabling Telnet access

Telnet access is disabled by default. You can use a Telnet client to access the CLI on the device over the network.

To enable Telnet operation, enter the following command.

```
device(config)# telnet server
```

If you do not plan to use the CLI over the network and want to disable Telnet access to prevent others from establishing CLI sessions with the device , enter the following command.

```
device(config)# no telnet server
```

Syntax: [no] telnet-server

Enabling Web management access for a device

Web Management is disabled by default. You can enable it through HTTP or HTTPS as described in the following sections.

> **NOTE**
> The Web Management Interface is only supported on the XMR Series and MLX Series devices.

Web management through HTTP

To allow web management through HTTP for an Extreme device, you enable web management as shown in the following command.

```
device(config)# web-management
```

Syntax: [no] web-management [**http** | **https**]

Using the web-management command without the http or https option makes web management available for both.

The **http** option specifies that web management is enabled for HTTP access.

The **https** option specifies that web management is enabled for HTTPS access.

Web management through HTTPS

The following encryption cipher algorithm are supported for HTTPS. They are listed in order of preference:
- **3des-cbc:** Triple-DES
- **rc4-des:** RC4 DES

To allow web management through HTTPS for an Extreme device you must enable web management as shown in Web management through HTTP on page 28. Additionally, you must generate a crypto SSL certificate or import digital certificates issued by a third-party Certificate Authority (CA).

To generate a crypto SSL certificate use the following command.

```
device(config)# crypto-ssl certificate generate
```

Syntax: [no] crypto-ssl certificate [**generate** | **zeroize**]

Using the web-management command without the http or https option makes web management available for both.

The **generate** parameter generates an ssl certificate.

The **zeroize** parameter deletes the currently operative ssl certificate.

To import a digital certificate issued by a third-party Certificate Authority (CA) and save it in the flash memory, use the following command.

```
device# copy tftp flash 10.10.10.1 cacert.pem server-certificate
```

Syntax: copy tftp flash ip-address file-name server-certificate

The *ip-address* variable is the IP address of the TFTP server where the digital certificate file is being downloaded from.

The *file-name* variable is the file name of the digital certificate that you are importing to the device.

Disabling Web management access by HP ProCurve Manager

By default, TCP ports 80 is enabled on the Extreme device. TCP port 80 (HTTP) allows access to the device's Web Management Interface.

By default, TCP port 280 for HP Top tools is disabled. This tool allows access to the device by HP ProCurve Manager.

The **no web-management** command disables both TCP ports. However, if you want to disable only port 280 and leave port 80 enabled, use the **hp-top-tools** option with the command.

```
device(config)# no web-management hp-top-tools
```

Syntax: [no] web-management hp-top-tools

The **hp-top-tools** parameter disables TCP port 280.

Enabling SNMP access

SNMP is disabled by default on the Extreme devices. SNMP is required if you want to manage an Extreme device using Brocade Network Advisor.

To enable SNMP management of the device.

```
device(config)#snmp-server
```

To later disable SNMP management of the device.

```
device(config)#no snmp-server
```

Syntax: [no] snmp-server

Setting passwords

Passwords can be used to secure the following access methods:

- Telnet access can be secured by setting a Telnet password. Refer to Setting a Telnet password on page 29.
- Access to the Privileged EXEC and CONFIG levels of the CLI can be secured by setting passwords for management privilege levels. Refer to Setting passwords for management privilege levels on page 30.

This section also provides procedures for enhancing management privilege levels, recovering from a lost password, and disabling password encryption.

> NOTE
> You can configure up to 32 user accounts consisting of a user name and password, and assign each user account a management privilege level. Refer to Setting up local user accounts on page 33.

Setting a Telnet password

By default, the device does not require a user name or password when you log in to the CLI using Telnet.

To set the password "letmein" for Telnet access to the CLI, enter the following command at the global CONFIG level.

```
device(config)# enable telnet password letmein
```

Syntax: **[no] enable telnet password** *string*

> NOTE
> If enable strict-password-enforcement is enabled, when a user is logged in and is attempting to change their own user password, the following prompt is displayed: Enter old password. After validating the old password, the following prompt is displayed: Enter new password.

Suppressing Telnet connection rejection messages

By default, if an Extreme device denies Telnet management access to the device, the software sends a message to the denied Telnet client. You can optionally suppress the rejection message. When you enable the option, a denied Telnet client does not receive a message from the Extreme device. Instead, the denied client simply does not gain access.

To suppress the connection rejection message sent by the device to a denied Telnet client, enter the following command at the global CONFIG level of the CLI.

```
device(config)# telnet server suppress-reject-message
```

Syntax: **[no] telnet server suppress-reject-message**

Setting passwords for management privilege levels

You can set one password for each of the following management privilege levels:

- **Super User level** - Allows complete read-and-write access to the system. This is generally for system administrators and is the only management privilege level that allows you to configure passwords.
- **Port Configuration level** - Allows read-and-write access for specific ports but not for global (system-wide) parameters.
- **Read Only level** - Allows access to the Privileged EXEC mode and CONFIG mode of the CLI but only with read access.

You can assign a password to each management privilege level. You also can configure up to 16 user accounts consisting of a user name and password, and assign each user account to one of the three privilege levels. Refer to Setting up local user accounts on page 33.

> NOTE
> You must use the CLI to assign a password for management privilege levels. You cannot assign a password using the Web Management Interface.

If you configure user accounts in addition to privilege level passwords, the device will validate a user's access attempt using one or both methods (local user account or privilege level password), depending on the order you specify in the authentication-method lists. Refer to Configuring authentication-method lists on page 40.

Follow the steps listed below to set passwords for management privilege levels.

1. At the opening CLI prompt, enter the following command to change to the Privileged level of the EXEC mode.

```
device> enable
device#
```

2. Access the CONFIG level of the CLI by entering the following command.

```
device# configure terminal
device(config)#
```

3. Enter the following command to set the Super User level password.

```
device(config)# enable super-user-password text
```

> **NOTE**
> You must set the Super User level password before you can set other types of passwords. The Super User level password can be an alphanumeric string, but cannot begin with a number.

4. Enter the following commands to set the Port Configuration level and Read Only level passwords.

```
device(config)# enable port-config-password text
device(config)# enable read-only-password text
```

> **NOTE**
> When **enable strict-password-enforcement** is in effect and you create a password using the following commands, the characters you type are masked. The examples in this guide do not mask the passwords for clarity.

Syntax: **enable super-user-password text**

Syntax: **enable port-config-password text**

Syntax: **enable read-only-password text**

> **NOTE**
> If you forget your Super User level password, refer to Recovering from a lost password on page 32.

> **NOTE**
> When **enable strict-password-enforcement** is enabled, the user uses the **enable super-user-password** to log in, and the **enable-super-user password** command is used, the following prompt is displayed: Enter old password. After validating the old password, the following prompt is displayed: Enter new password.

Augmenting management privilege levels

Each management privilege level provides access to specific areas of the CLI by default:

- Super User level provides access to all commands and displays.
- Port Configuration level gives access to the following:
  - The User EXEC and Privileged EXEC levels
  - The port-specific parts of the CONFIG level
  - All interface configuration levels
- Read Only level gives access to the following:
  - The User EXEC and Privileged EXEC levels

You can grant additional access to a privilege level on an individual command basis. To grant the additional access, you specify the privilege level you are enhancing, the CLI level that contains the command, and the individual command.

> **NOTE**
> This feature applies only to management privilege levels on the CLI. You cannot augment management access levels for the Web Management Interface.

To enhance the Port Configuration privilege level so users also can enter IP commands at the global CONFIG level.

```
device(config)# privilege configure level 4 ip
```

In this command, **configure** specifies that the enhanced access is for a command at the global CONFIG level of the CLI. The **level 4** parameter indicates that the enhanced access is for management privilege level 4 (Port Configuration). All users with Port Configuration privileges will have the enhanced access. The ip parameter indicates that the enhanced access is for the IP commands. Users who log in with valid Port Configuration level user names and passwords can enter commands that begin with "ip" at the global CONFIG level.

Syntax: [no] privilege cli-level level privilege-level command-string

The *cli-level* parameter specifies the CLI level and can be one of the following values:

- **exec** - EXEC level; for example, device`#`
- **configure** - CONFIG level; for example, device`(config)#`
- **interface** - Interface level; for example, device`(config-if-e10000-6)#`
- **virtual-interface** - Virtual-interface level; for example, device`(config-vif-6)#`
- **rip-router** - RIP router level; for example, device`(config-rip-router)#`
- **ospf-router** - OSPF router level; for example, device`(config-ospf-router)#`
- **bgp-router** - BGP4 router level; for example,device`(config-bgp-router)#`
- **port-vlan** - Port-based VLAN level; for example, device`(config-vlan)#`
- **protocol-vlan** - Protocol-based VLAN level
- **dot1x**
- **loopback-interface**
- **tunnel-interface**
- **vrrp-router**

The *privilege-level* indicates the number of the management privilege level you are augmenting. You can specify one of the following:

- **0** - Super User level (full read-write access)
- **4** - Port Configuration level
- **5** - Read Only level

The *command-string* parameter specifies the command you are allowing users with the specified privilege level to enter. To display a list of the commands at a CLI level, enter "?" at that level's command prompt.

Recovering from a lost password

Recovery from a lost password requires direct access to the serial port and a system reset.

> NOTE
> You can perform this procedure only from the CLI.

Follow the steps listed below to recover from a lost password.

1. Start a CLI session over the serial interface to the device.
2. Reboot the device.
3. At the initial boot prompt at system startup, enter **b** to enter the boot monitor mode.
4. Enter **no password** at the prompt. (You cannot abbreviate this command.) This command will cause the device to bypass the system password check.
5. Enter **boot system flash primary** at the prompt, and enter **y** when asked "Are you sure?".
6. After the console prompt reappears, assign a new password.

Displaying the SNMP community string

If you want to display the SNMP community string, enter the following commands.

```
device(config)# enable password-display
device(config)# show snmp server
```

The **enable password-display** command enables display of the community string, but only in the output of the **show snmp server** command. Display of the string is still encrypted in the startup configuration file and running configuration. Enter the command at the global CONFIG level of the CLI.

Disabling password encryption

When you configure a password, then save the configuration to the device's flash memory, the password is also saved to flash as part of the configuration file. By default, the passwords are encrypted so that the passwords cannot be observed by another user who displays the configuration file. Even if someone observes the file while it is being transmitted over TFTP, the password is encrypted.

If you want to remove the password encryption, you can disable encryption by entering the following command.

```
device(config)# no service password-encryption
```

Syntax: **[no] service password-encryption**

Specifying a minimum password length

By default, the device imposes no minimum length on the Line (Telnet), Enable, or Local passwords. You can configure the device to require that Line, Enable, and Local passwords be at least a specified length.

For example, to specify that the Line, Enable, and Local passwords be at least 8 characters, enter the following command.

```
device(config)# enable password-min-length 8
```

Syntax: **[no] enable password-min-length number-of-characters**

The *number-of-characters* can be from 1 - 48.

Setting up local user accounts

You can define up to 32 local user accounts on an Extreme device. User accounts regulate who can access the management functions in the CLI using the following methods:

- Telnet access
- SSH access
- Console access
- Web management access
- SNMP access

Local user accounts provide greater flexibility for controlling management access to the Extreme device than do management privilege level passwords and SNMP community strings of SNMP versions 1 and 2. You can continue to use the privilege level passwords and the SNMP community strings as additional means of access authentication. Alternatively, you can choose not to use local user accounts and instead continue to use only the privilege level passwords and SNMP community strings. Local user accounts are backward-compatible with configuration files that contain privilege level passwords. Refer to Setting passwords for management privilege levels on page 30.

If you configure local user accounts, you also need to configure an authentication-method list for Telnet access, Web management access, and SNMP access. Refer to Configuring authentication-method lists on page 40.

For each local user account, you specify a user name which can have up to 48 characters. You also can specify the following parameters:

- A password
- A management privilege level, which can be one of the following:
    - **Super User level** - Allows complete read-and-write access to the system. This is generally for system administrators and is the only privilege level that allows you to configure passwords. This is the default.
    - **Port Configuration level** - Allows read-and-write access for specific ports but not for global (system-wide) parameters.
    - **Read Only level** - Allows access to the Privileged EXEC mode and CONFIG mode but only with read access.

Configuring a local user account

To configure a local user account, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)# username wonka password willy
```

This command adds a local user account with the user name "wonka" and the password "willy". This account has the Super User privilege level; this user has full access to all configuration and display features.

> **NOTE**
> If you configure local user accounts, you must grant Super User level access to at least one account before you add accounts with other privilege levels. You need the Super User account to make further administrative changes.

```
device(config)# username waldo privilege 5 password whereis
```

This command adds a user account for user name "waldo", password "whereis", with the Read Only privilege level. Waldo can look for information but cannot make configuration changes.

To display user account information, enter the following command.

```
device(config)# show users
**USER SESSIONS**
Username    Role    Host IP       Device    Time Logged In
jsmith      user    192.0.2.0     Cli       2016-04-30 01:59:35
jdoe        admin   192.0.2.1     Cli       2016-05-30 01:57:41
waldo       user    192.0.2.0     Cli       2016-06-30 03:39:33

**LOCKED USERS**
testUser
```

Note about changing local user passwords

The Extreme device stores not only the current password configured for a local user, but the previous two passwords configured for the user as well. The local user's password cannot be changed to one of the stored passwords.

Consequently, if you change the password for a local user, you must select a password that is different from the current password, as well as different from the previous two passwords that had been configured for that user.

For example, say local user waldo originally had a password of "whereis", and the password was subsequently changed to "whois", then later changed to "whyis". If you change waldo's password again, you cannot change it to "whereis", "whois", or "whyis".

The current and previous passwords are stored in the device's running configuration file in encrypted form.

```
device# show run
...
username waldo password 8 $1$Ro2..Ox0$udBu7pQT5XyuaXMUiUHy9. history $1$eq...T62$IfpxIcxnDWX7CSVQKIodu.
```

```
$1$QD3..2Q0$DYxgxCI64ZOSsYmSSaA28/
...
```

In the running configuration file, the user's previous two passwords are displayed in encrypted form following the **history** parameter.

Enabling strict password enforcement

Additional security to the local username and password by configuring the **enable strict-password-enforcement** CLI command. Note the rules for passwords if the strict password is disabled and when it is enabled.

Configuring the strict password rules

Use the **enablestrict-password-enforcement** command to enable the strict password enforcement feature. Enter a command such as the following.

```
device(config)# enable strict-password-enforcement
```

Syntax: **[no] enable strict-password-enforcement**

This feature is disabled by default.

When enabled, the system verifies uniqueness against the history of passwords of the user whose password is being set. Passwords must not share four or more concurrent characters with any other password configured for that user on the device. If the user tries to create a password which shares four or more concurrent characters for that user, the following error message is returned:

```
Error - The substring <str> within the password has been used earlier, please choose a different password.
```

Also, if the user tries to configure a password that was previously configured, the local user account configuration is not allowed and the following message is displayed.

```
Error - This password was used earlier, please choose a different password.
```

When you create a password, the characters you type are masked.

: To assign a password for a user account.

```
device(config)# username sandy password [Enter]
Enter new password: ********
```

Syntax: **[no] username name password**

Enter a password such as TesT12$! that contains the required character combination.

> **NOTE**
> If enable strict-password-enforcement is enabled, when a user is logged in and is attempting to change their own user password, the following prompt is displayed: Enter old password. After validating the old password, the following prompt is displayed: Enter new password.

Password history

If the **enable strict-password-enforcement** command is enabled, the CLI keeps the last 15 passwords used by the user. A user is prevented from changing the password to one that has already been used.

Setting passwords to expire

If the **enable strict-password-enforcement** command is enabled, passwords can be set to expire, early warning periods can be configured, and grace login reset attempts can be configured.

To configure a user password to expire, enter the following.

```
device(config)# enable strict-password-enforcement
device(config)# username sandy expires 20
```

Syntax: **[no] username** *name* **expires** *days*

The *name* variable specifies the user that the expiration time is applied to.

The *days* variable specifies the number of day before the password will expire. The following values can be used 1 - 365 days. The default is 90 days.

> NOTE
> The **enable strict-password-enforcement** command must be enabled before this command is configured. Otherwise, the following message will be displayed: "`Password expire time is enabled only if strict-password-enforcement is set.`"

If the **enable strict-password-enforcement** command is enabled, the administrator can configure an early warning period to warn users for a particular number of days prior to their password expiring.

To configure the early warning period for password expiration, enter the following:

```
device(config)# enable strict-password-enforcement expiration early-warning-period 5
```

Syntax: **[no] enable strict-password-enforcement expiration early-warning-period** *days*

The *days* variable specifies the number of days prior to password expiration of a user that a notification of password expiration is printed at user login. The default is 10 days, the minimum is 1 day, and maximum is 365 days.

Once the early warning period is set, when the user successfully logs in within the early warning period time frame, the following message is displayed: "password will expire in *x* day(s)", where *x* is the number of days remaining before the password expires.

Once the password is expired, the user is permitted a configurable amount of subsequent login attempts. There is no limit on the time-period before requiring a new password. The only limit is the number of subsequent login attempts allowed for that user.

To configure the maximum grace login attempts allowed for a user once the user's password has expired, enter the following:

```
device(config)# enable strict-password-enforcement expiration grace-login-attempts 2
```

Syntax: **[no] enable strict-password-enforcement expiration grace-login-attempts** *times*

The *times* variable specifies the maximum number of times a user can log in after password expiration. The default is 3 times, the maximum is 3 times, and the minimum is 0 times.

The **show user** command can be used to display the expiration date or remaining grace logins as shown in **bold** in the following:

```
device(config)#show users
Username          Password                        Encrypt   Priv Status   Expire Time/Grace Logins
=================================================================================
user1             $1$E81..sj4$Kv25UrYDLYHaSv.SQY8fB.  enabled   0    enabled  90 days
user2          $1$Tm3..r91$O7l5L98/V7ivvRxgJKPNU0  enabled   0    enabled  90 days
user3          $1$Qn/..9n2$3lAwyrYolr2Pe.5x5wdYw.  enabled   0    expired  1 grace
```

Login lockout

If the **enable strict-password-enforcement** command is enabled, users have up to three login attempts. If a user fails to login after third attempts, that user is locked out (disabled).

To re-enable a user that has been locked out, perform one of the following tasks:

- Reboot the device to re-enable all disabled users.
- Enable the user by entering the following command.

```
device(config)# username sandy enable
```

Syntax: [no] username name enable

The *name* variable specifies the username to be enabled.

Requirement to accept the message of the day

If a message of the day (MOTD) is configured and the **enable strict-password-enforcement** command is enabled, user is required to press the Enter key before he or she can login. MOTD is configured using the **banner motd** command.

```
device(config)# banner motd require-enter-key
```

Syntax: [no] banner motd require-enter-key

Regular password rules

The following rules apply to passwords unless the **enable strict-password-enforcement** command is executed:

- A minimum of one character is required to create a password.
- The last 3 passwords are stored in the CLI.
- No password expiration.
- Users are not locked out (disabled) after failed login attempts.

Strict password rules

> **NOTE**
> If enable strict-password-enforcement is enabled, when a user is logged in and is attempting to change their own user password, the following prompt is displayed: Enter old password. After validating the old password, the following prompt is displayed: Enter new password.

Rules for passwords are different if the strict password enforcement is used. By default, the following rules apply when the **enable strict-password-enforcement** command is executed:

- Users are required to accept the message of the day (enabled).

In addition to the rule above, the following rules can be enabled:

- The device can store the last 15 passwords in the CLI.
- Password can be set to expire.
- Password grace login attempts can be configured by administrator.
- Password expiration early warning period can be configured by administrator.
- Passwords are masked during password creation.

- Passwords may not share four or more concurrent characters with any other password configured on the device.
- Passwords that were previously configured for a user can be rejected.

When you create an enable and a user password, you must enter a minimum of eight characters containing the following combinations:

- At least two upper case characters
- At least two lower case characters
- At least two numeric characters
- At least two special character

> NOTE
> Password minimum and combination requirements are strictly enforced.

Web interface login lockout

The Web interface provides up to three login attempts. If a user fails to login after three attempts, that user is locked out (disabled).

To re-enable a user that has been locked out, reboot the device to re-enable all disabled users.

> NOTE
> The Web Management Interface is only supported on the XMR Series and MLX Series devices.

Creating an encrypted all-numeric password

To create a password that is made up of all numeric values, use the command "**username** *user-string* **privilege** *privilege-level* **password** *password-string* ." To allow backward compatibility with the **username** command, the new keyword **create-password** has been created and it is used as shown in the following.

```
device# username customer1 create-password 9999
```

Syntax: [no] username user-string create-password password-string

The **create-password** option allows you to create a password with a numeric value in the *password-string* variable. The generated password will be encrypted. The **show running-config** command will display the password as shown.

```
username <user-string> 8 <encrypted-password>
```

> NOTE
> The **create-password** option is not supported when the **strict-password-enforcement** is in effect.

Granting access by time of day

To configure an Extreme device to restrict access to a specified user to a specified time of day, use the following command.

```
device(config)# username admin1 access-time 10:00:00 to 13:00:00
```

Syntax: [no] username user-string access-time hh:mm:ss to hh:mm:ss

The *user-string* variable specifies the user that you want to limit access time for.

The first instance of the *hh:mm:ss* variable specifies the start of the access time and the second instance of the *hh:mm:ss* variable specifies the end of the access time.

AAA

Configuring AAA on the console

Only enable-level authentication is available on the console by default. Command authorization and accounting and exec accounting must be explicitly configured.

1. Enter the **enable aaa console** to access the console.

2. At the console, type "**end**" to go to the Privileged EXEC level.

3. Type "**exit**" to go to the User EXEC level.

   Once the AAA support is enabled on the console, a new command, exit is available at the User EXEC level.

4. Enter "**exit**" to display the following login prompt on the console window.

   ```
   "Press Enter key to login".
   ```

5. Press the **Enter**, key to begin the login process.

   The next prompt to appear is determined by the first method configured in the login authentication configuration. If it is not TACACS+, the default prompts are used.

   > NOTE
   > If you use the use the **aaa console** command to enable AAA, you must make sure that the method lists are configured to allow access. Otherwise, you will be locked out of the console.

Configuring AAA authentication-method lists for login

With AAA is enabled on the console, you must configure an authentication-method list to set the conditions for granting access to the console. The authentication methods supported on the Extreme devices include the following:

* enable
* line
* local
* radius
* tacacs
* tacacs+
* local-auth-fallback
* none

When a list is configured, the first method listed is attempted to provide authentication at login. If that method is not available, (for example, a TACACs server can not be reached) the next method is tried until a method in the list is available or all methods have been tried. You can place the method none at the end of a list to ensure that access will always be available if all active methods fail.

NOTE

If you configure the **aaa authentication login default radius local-auth-fallback** command, then if RADIUS rejects it and authentication is performed using local, any configured command and EXEC authorization using RADIUS will not work, even if it is configured.

To configure a AAA authentication-method list for login, use the following command.

```
device(config)# aaa authentication login default tacacs+ local none
```

In this configuration, tacacs+ would be tried first. If a tacacs+ server cannot be reached, the local system password would be used. If this method fails, authentication would default to none.

The **enable** option uses the enable password configured on the device to grant access to the console.

The **line** option uses the line password configured on the device to grant access to the console.

The **local** option uses the local password configured on the device to grant access to the console.

The **radius** option uses authentication provided by a radius server to grant access to the console.

The **tacacs** option uses authentication provided by a tacacs server to grant access to the console.

The **tacacs+** option uses authentication provided by a tacacs+ server to grant access to the console.

The **local-auth-fallback** option indicates that authentication should fall-back to local authentication in the event that the authentication at any server in an earlier authentication method in the authentication method list is denied. It can only be used as the last authentication method when the other authentication methods do not include local and none. In other respects, it is similar to local authentication.

The **none** option eliminates the requirement for any authentication method to grant access to the console.

Configuring authentication-method lists

To implement one or more authentication methods for securing access to the device, you configure authentication-method lists that set the order in which the authentication methods are consulted.

In an authentication-method list, you specify the access method (Telnet, Web, SNMP, and so on) and the order in which the device tries one or more of the following authentication methods:

- Local Telnet login password
- Local password for the Super User privilege level
- Local user accounts configured on the device
- Database on a TACACS or TACACS+ server
- Database on a RADIUS server
- No authentication

    NOTE

    The TACACS or TACACS+, RADIUS, and Telnet login password authentication methods are not supported for SNMP access.

    NOTE

    To authenticate Telnet access to the CLI, you also must enable the authentication by entering the **enable telnet authentication** command at the global CONFIG level of the CLI. You cannot enable Telnet authentication using the Web Management Interface.

> **NOTE**
> You do not need an authentication-method list to secure access based on ACLs or a list of IP addresses. Refer to Using ACLs to restrict remote access on page 21 or Restricting remote access to the device to specific IP addresses on page 24.

In an authentication-method list for a particular access method, you can specify up to seven authentication methods. If the first authentication method is successful, the software grants access and stops the authentication process. If the access is rejected by the first authentication method, the software denies access and stops checking.

However, if an error occurs with an authentication method, the software tries the next method on the list, and so on. For example, if the first authentication method is the RADIUS server, but the link to the server is down, the software will try the next authentication method in the list.

> **NOTE**
> If an authentication method is working properly and the password (and user name, if applicable) is not known to that method, this is not an error. The authentication attempt stops, and the user is denied access.

The software will continue this process until either the authentication method is passed or the software reaches the end of the method list. If the Super User level password is not rejected after all the access methods in the list have been tried, access is granted.

> **NOTE**
> If a user cannot be authenticated using local authentication, then the next method on the authentication methods list is used to try to authenticate the user. If there is no method following local authentication, then the user is denied access to the device.

> **NOTE**
> Telnet login does not work with the local option. When the authentication list for login is configured with " tacacs+ line enable local", and all previous methods timeout except local, the telnet login does not authenticate.

The following login combinations do not function with the local option:

- aaa authentication login default tacacs+ line local
- aaa authentication login default tacacs+ enable local
- aaa authentication login default tacacs+ line radius local
- aaa authentication login default tacacs+ enable radius local
- aaa authentication login default tacacs+ enable line radius local

Configuration considerations for authentication-method lists

The configuration considerations for authentication-method lists are as follows:

- For CLI access, you must configure authentication-method lists if you want the device to authenticate access using local user accounts or a RADIUS server. Otherwise, the device will authenticate using only the locally based password for the Super User privilege level.
- When no authentication-method list is configured specifically for Web management access, the device performs authentication using the SNMP community strings:
  - For read-only access, you can use the user name "get" and the password "public". The default read-only community string is "public".
  - There is no default read-write community string. Thus, by default, you cannot open a read-write management session using the Web Management Interface. You first must configure a read-write community string using the CLI. Then you can log on using "set" as the user name and the read-write community string you configure as the password.
- If you configure an authentication-method list for Web management access and specify "local" as the primary authentication method, users who attempt to access the device using the Web Management Interface must supply a user name and password

configured in one of the local user accounts on the device. The user cannot access the device by entering "set" or "get" and the corresponding SNMP community string.

- For devices that can be managed using Brocade Network Advisor, the default authentication method (if no authentication-method list is configured for SNMP) is the CLI Super User level password. If no Super User level password is configured, then access through Brocade Network Advisor is not authenticated. To use local user accounts to authenticate access through Brocade Network Advisor, configure an authentication-method list for SNMP access and specify "local" as the primary authentication method.

Examples of authentication-method lists

The following example shows how to configure authentication-method lists for the Web Management Interface, and the Privileged EXEC and CONFIG levels of the CLI. In this example, the primary authentication method for each is "local". The device will authenticate access attempts using the locally configured user names and passwords first.

These examples also apply to the Brocade Network Advisor.

To configure an authentication-method list for the Web Management Interface, enter a command such as the following.

```
device(config)# aaa authentication web-server default local
```

This command configures the device to use the local user accounts to authenticate access to the device through the Web Management Interface. If the device does not have a user account that matches the user name and password entered by the user, the user is not granted access.

To configure an authentication-method list for Brocade Network Advisor, enter a command such as the following.

```
device(config)# aaa authentication snmp-server default local
```

This command configures the device to use the local user accounts to authenticate access attempts through any network management software.

To configure an authentication-method list for the Privileged EXEC and CONFIG levels of the CLI, enter the following command.

```
device(config)# aaa authentication enable default local
```

This command configures the device to use the local user accounts to authenticate attempts to access the Privileged EXEC and CONFIG levels of the CLI.

To configure the device to consult a RADIUS server first to authenticate attempts to access the Privileged EXEC and CONFIG levels of the CLI, then consult the local user accounts if the RADIUS server is unavailable, enter the following command.

```
device(config)# aaa authentication enable default radius local
```

The **snmp-server | web-server | enable | login | dot1x** parameter specifies the type of access this authentication-method list controls. You can configure one authentication-method list for each type of access.

> NOTE
> If you configure authentication for Web management access, authentication is performed each time a page is requested from the server. When frames are enabled on the Web Management Interface, the browser sends an HTTP request for each frame. The Extreme device authenticates each HTTP request from the browser. To limit authentications to one per page, disable frames on the Web Management Interface.

> NOTE
> TACACS or TACACS+ and RADIUS are not supported with the **snmp-server** parameter.

The *method1* parameter specifies the primary authentication method. The remaining optional *method* parameters specify additional methods to try if an error occurs with the primary method. A method can be one of the values listed in the Method Parameter column in the table below.

Method parameter	Description
line	Authenticate using the password you configured for Telnet access. The Telnet password is configured using the **enable telnet password...** command.
	Refer to Setting a Telnet password on page 29.
enable	Authenticate using the password you configured for the Super User privilege level. This password is configured using the **enable super-user-password...** command.
	Refer to Setting passwords for management privilege levels on page 30.
local	Authenticate using a local user name and password you configured on the device. Local user names and passwords are configured using the **username...** command.
	Refer to Configuring a local user account on page 34.
tacacs	Authenticate using the database on a TACACS server. You also must identify the server to the device using the **tacacs-server** command.
tacacs+	Authenticate using the database on a TACACS+ server. You also must identify the server to the device using the **tacacs-server** command.
radius	Authenticate using the database on a RADIUS server. You also must identify the server to the device using the **radius-server** command.
none	Do not use any authentication method. The device automatically permits access.

RADIUS Authentication

Configuring RADIUS security

You can use a Remote Authentication Dial In User Service (RADIUS) server to secure the following types of access to the Extreme devices:

* Telnet access
* SSH access
* Web management access
* Access to the Privileged EXEC level and CONFIG levels of the CLI

> **NOTE**
> The Extreme devices do not support RADIUS security for SNMP (Brocade Network Advisor) access.

RADIUS authentication, authorization, and accounting

When RADIUS authentication is implemented, the Extreme device consults a RADIUS server to verify user names and passwords. Optionally, you can configure RADIUS authorization, in which the Extreme device consults a list of commands supplied by the RADIUS server to determine whether a user can execute a command that has been entered. You can configure RADIUS accounting, which causes the Extreme device to log information on a RADIUS accounting server when specified events occur on the device.

> **NOTE**
> By default, a user logging into the device through Telnet or SSH first enters the User EXEC level. The user can then enter the **enable** command to get to the Privileged EXEC level.

> **NOTE**
> A user that is successfully authenticated can be automatically placed at the Privileged EXEC level after login.

RADIUS authentication

The following events occur when RADIUS authentication takes place.

1. A user triggers RADIUS authentication by doing one of the following:
   * Logging in to the Extreme device using Telnet, SSH, or the Web Management Interface
   * Entering the Privileged EXEC level or CONFIG level of the CLI
2. The user is prompted for a username and password.

   > **NOTE**
   > The SSH client prompts the user for the password before the SSH server is contacted. If the SSH server has a non-standard multi-factor authentication prompt, it may not match the prompt displayed.

3. The user enters a username and password.
4. The Extreme device sends a RADIUS Access-Request packet containing the username and password to the RADIUS server.
5. The RADIUS server validates the Extreme device using a shared secret (the RADIUS key).

6. The RADIUS server looks up the username in its database.

7. If the username is found in the database, the RADIUS server validates the password.

8. If the password is valid, then:

    a) If the RADIUS server is configured to use multi-factor authentication, it may send an Access-Challenge packet to the Extreme device. If so, the user may be asked for additional input (for example, an RSA SecurID PIN or RSA SecurID next tokencode) which the Extreme device will forward to the RADIUS server. If the additional input is valid, then the process moves to the next step.

    b) If the RADIUS server is configured to use single-factor authentication, then the process moves immediately to the next step.

9. The RADIUS server sends an Access-Accept packet to the Extreme device, authenticating the user. Within the Access-Accept packet are three Extreme vendor-specific attributes that indicate:

    - – The privilege level of the user
      – A list of commands
      – Whether the user is allowed or denied usage of the commands in the list

    The last two attributes are used with RADIUS authorization, if configured.

10. The user is authenticated, and the information supplied in the Access-Accept packet for the user is stored on the Extreme device. The user is granted the specified privilege level. If you configure RADIUS authorization, the user is allowed or denied usage of the commands in the list.

Multi-factor RADIUS authentication

The Extreme device supports multi-factor authentication (for example, RSA SecurID) through a RADIUS server. For access by Telnet, no further configuration is needed on the Extreme device to enable multi-factor RADIUS authentication.

The default is **yes** (interactive authentication is supported by default); therefore, all you must do is configure SSH to use multi-factor authentication.

```
device(config)# ip ssh interactive-authentication
```

Syntax: ip ssh interactive-authentication [**no** | **yes**]

Refer to "Configuring Secure Shell and Secure Copy" for SSH configuration details.

> **NOTE**
> The SSH client prompts the user for the password before the SSH server is contacted. If the SSH server has a non-standard multi-factor authentication prompt, it may not match the prompt displayed.

A sample interactive authentication session (with RSA SecurID) is shown below.

```
Telnet_DMT_MLX_16k - 08-25-2010  -- 11:20:18  Session Log Start -- 10.20.179.55|Telnet - 08-25-2010  --
11:20:18
Telnet - 08-25-2010  -- 11:20:18  This is the message of the day
Telnet - 08-25-2010  -- 11:20:18
Telnet - 08-25-2010  -- 11:20:18  User Access Verification
Telnet - 08-25-2010  -- 11:20:18
Telnet - 08-25-2010  -- 11:20:38  Please Enter Login Name: pbikram3
Telnet - 08-25-2010  -- 11:20:58  Please Enter Password: <enter-token-code-for-user-here>
Telnet - 08-25-2010  -- 11:21:01
Telnet - 08-25-2010  -- 11:21:06  Enter a new PIN having from 4 to 8 alphanumeric characters:<new-pin>
Telnet - 08-25-2010  -- 11:21:07
Telnet - 08-25-2010  -- 11:21:10  Please re-enter new PIN:<new-pin>
Telnet - 08-25-2010  -- 11:21:12
Telnet - 08-25-2010  -- 11:21:12  PIN Accepted.
Telnet - 08-25-2010  -- 11:21:12  Wait for the token code to change,
Telnet - 08-25-2010  -- 11:21:36  then enter the new passcode:<new-pin>+<enter-token-code-for-user-here>
```

```
Telnet - 08-25-2010  -- 11:21:38
Telnet - 08-25-2010  -- 11:21:38  User login successful.
Telnet - 08-25-2010  -- 11:21:38
Telnet - 08-25-2010  -- 11:21:55  Session Log End --10.20.179.55|Telnet - 08-25-2010  -- 11:21:55
```

RADIUS authorization

The following events occur when RADIUS authorization takes place.

1. A user previously authenticated by a RADIUS server enters a command on the Extreme device.

2. The Extreme device looks at its configuration to see if the command is at a privilege level that requires RADIUS command authorization.

3. If the command belongs to a privilege level that requires authorization, the Extreme device looks at the list of commands delivered to it in the RADIUS Access-Accept packet when the user was authenticated. (Along with the command list, an attribute was sent that specifies whether the user is permitted or denied usage of the commands in the list.)

   > **NOTE**
   > After RADIUS authentication takes place, the command list resides on the Extreme device. The RADIUS server is not consulted again once the user has been authenticated. This means that any changes made to the user's command list on the RADIUS server are not reflected until the next time the user is authenticated by the RADIUS server, and the new command list is sent to the Extreme device.

4. If the command list indicates that the user is authorized to use the command, the command is executed.

RADIUS accounting

The following steps explain the working of RADIUS accounting.

1. One of the following events occur on an Extreme device:
   - – A user logs into the management interface using Telnet or SSH
   - – A user enters a command for which accounting has been configured
   - – A system event occurs, such as a reboot or reloading of the configuration file

2. The Extreme device checks its configuration to see if the event is one for which RADIUS accounting is required.

3. If the event requires RADIUS accounting, the Extreme device sends a RADIUS Accounting Start packet to the RADIUS accounting server, containing information about the event.

4. The RADIUS accounting server acknowledges the Accounting Start packet.

5. The RADIUS accounting server records information about the event.

6. When the event is concluded, the Extreme device sends an Accounting Stop packet to the RADIUS accounting server.

7. The RADIUS accounting server acknowledges the Accounting Stop packet.

AAA operations for RADIUS

The following table lists the sequence of authentication, authorization, and accounting operations that take place when a user gains access to an Extreme device that has RADIUS security configured.

User action	Applicable AAA operations
User attempts to gain access to the Privileged EXEC and CONFIG levels of the CLI	Enable authentication: aaa authentication enable default *method-list*
	System accounting start:

User action	Applicable AAA operations
	aaa accounting system default start-stop *method-list*
User logs in using Telnet or SSH	Login authentication: aaa authentication login default *method-list*
	EXEC accounting Start: aaa accounting exec default start-stop *method-list* System accounting Start: aaa accounting system default start-stop *method-list*
User logs into the Web Management Interface	Web authentication: aaa authentication web-server default *method-list*
User logs out of Telnet or SSH session	Command authorization for **logout** command: aaa authorization commands *privilege-level* default *method-list*
	Command accounting: aaa accounting commands *privilege-level* default start-stop *method-list* EXEC accounting stop: aaa accounting exec default start-stop *method-list*
User enters system commands (for example, **reload** , **boot system**)	Command authorization: aaa authorization commands *privilege-level* default *method-list*
	Command accounting: aaa accounting commands *privilege-level* default start-stop *method-list* System accounting stop: aaa accounting system default start-stop *method-list*
User enters the command: [no] aaa accounting system defaultstart-stop *method-list*	Command authorization: aaa authorization commands *privilege-level* default *method-list*
	Command accounting: aaa accounting commands *privilege-level* default start-stop *method-list* System accounting start: aaa accounting system default start-stop *method-list*
User enters other commands	Command authorization: aaa authorization commands *privilege-level* default *method-list*
	Command accounting: aaa accounting commands *privilege-level* default start-stop *method-list*

AAA security for commands pasted into the running configuration

If AAA security is enabled on the device, commands pasted into the running configuration are subject to the same AAA operations as if they were entered manually.

When you paste commands into the running configuration, and AAA command authorization or accounting is configured on the device, AAA operations are performed on the pasted commands. The AAA operations are performed before the commands are actually added to the running configuration. The server performing the AAA operations should be reachable when you paste the commands into the running configuration file. If the device determines that a pasted command is invalid, AAA operations are halted on the remaining commands. The remaining commands may not be executed if command authorization is configured.

NOTE
Since RADIUS command authorization relies on a list of commands received from the RADIUS server when authentication is performed, it is important that you use RADIUS authentication when you also use RADIUS command authorization.

RADIUS configuration considerations and restrictions

Consider the following for configuring the RADIUS server:

- You must deploy at least one RADIUS server in your network.
- The Extreme device supports authentication using up to eight RADIUS servers. The device tries to use the servers in the order you add them to the device's configuration. If one RADIUS server is not responding, the device tries the next one in the list.
- You can select only one primary authentication method for each type of access to a device (CLI through Telnet, CLI Privileged EXEC and CONFIG levels). For example, you can select RADIUS as the primary authentication method for Telnet CLI access, but you cannot also select TACACS+ authentication as the primary method for the same type of access. However, you can configure backup authentication methods for each access type.
- When a radius-server host is configured, a status-server request is sent automatically to determine the current status of the server. You must configure the radius-server key before entering the radius host command. The radius-server key may also be configured along with radius-server host command.

  Example 1:

  ```
  device(config)# radius-server key key
  device(config)# radius-server host a.b.c.d
  ```

  Example 2:

  ```
  device(config)# radius-server host a.b.c.d auth-port n acct-port n default key key
  ```

  - There will not be any retransmission of Status-Server packets, therefore, the timeout counter will not increase in case of a non-response.
  - Not all radius servers support status-server requests. For those servers to perform successful authentication, the one of following commands must be included in the radius-server configuration.

    › To disable radius health check at the global level:

    ```
    device(config)# no radius-server enable-health-check
    ```

  - › To disable radius health check for a specific server:

    ```
    device(config)# radius-server host a.b.c.d health-check disable
    ```

RADIUS configuration procedure

Use the following procedure to configure an Extreme device for RADIUS.

1. Configure Extreme vendor-specific attributes on the RADIUS server. Refer to Configuring Extreme-specific attributes on the RADIUS server on page 50.
2. Enabling Radius. Refer to Enabling SNMP traps for RADIUS on page 52.
3. Identify the RADIUS server to the Extreme device. Refer to Identifying the RADIUS server to the device on page 53.
4. Set RADIUS parameters. Refer to Setting RADIUS parameters on page 55.
5. Configure authentication-method lists. Refer to Configuring authentication-method lists for RADIUS on page 56.
6. Optionally configure RADIUS authorization. Refer to Configuring RADIUS authorization on page 57.

7. Optionally configure RADIUS accounting. Configuring RADIUS accounting on page 58.

Configuring Extreme-specific attributes on the RADIUS server

During the RADIUS authentication process, if a user supplies a valid username and password, the RADIUS server sends an Access-Accept packet to the Extreme device, authenticating the user. Within the Access-Accept packet, the RADIUS server could send attribute "Vendor-Specific" whose value could inform the Extreme device on the runtime environment for this session. This section will detail all the vendor-specific attributes defined by Extreme.

Attribute name	Attribute ID	Data type	Description
foundry-privilege-level	1	integer	Specifies the privilege level for the user. This attribute can be set to one of the following: **0** -**Super User level** - Allows complete read-and-write access to the system. This is generally for system administrators and is the only management privilege level that allows you to configure passwords. **4** - **Port Configuration level** - Allows read-and-write access for specific ports but not for global (system-wide) parameters. **5** - **Read Only level** - Allows access to the Privileged EXEC mode and CONFIG mode of the CLI but only with read access.
foundry-command-string	2	string (maximum of 247 characters allowed)	Specifies a list of CLI commands that are permitted or denied to the user when RADIUS authorization is configured. The commands are delimited by semi-colons (;) and without a space between the command and the semi-colon. You can specify an asterisk (*) as a wildcard at the end of a command string. For example, the following command list specifies all **show**, **debug ip**, and **write terminal** commands: show *;debug ip *;write terminal * On XMR Series, MLX Series, and CER 2000 Series devices only, you can optionally specify a range of data for a CLI command. For example, the following command list specifies **show vlan** commands in the range 1 through 10, as well as all **debug ip** and **write terminal** commands: show vlan [1-10];debug ip *;write terminal *

Attribute name	Attribute ID	Data type	Description
foundry-command-exception-flag	3	integer	Specifies whether the commands indicated by the foundry-command-string attribute are permitted or denied to the user. This attribute can be set to one of the following: **0** - Permit execution of the commands indicated by foundry-command-string, deny all other commands. **1** - Deny execution of the commands indicated by foundry-command-string, permit all other commands.
foundry-INM-privilege	4	integer	Specifies the Brocade Network Advisor user privilege level. This attribute can take a value range from 0 to 15. In Brocade Network Advisor, this attribute value will be mapped to the preconfigured roles "AAA privilege level 0" through "AAA privilege level 15". The admin user has to configure these roles with the appropriate sets of privileges in order for the AAA user to get the correct set of feature access.
foundry-access-list	5	string	Specifies the access control list to be used for RADIUS authorization. Enter the access control list in the following format. type=string, value="ipacl.[e|s].[in|out] = [*acl-name* |*acl-number*] *separator* macfilter.in = [*acl-name*	*acl-number*] Where: • separator can be a space, new line, semicolon, comma, or null character • ipacl.e is extended ACL; ipacl.s is standard ACL. **NOTE** Outbound MAC filters are not supported, but outbound ACLs with 802.1X authentication is supported.
foundry-MAC-authent-needs-802x	6	integer	Specifies whether or not 802.1x authentication is required and enabled. **0** - Disabled

Attribute name	Attribute ID	Data type	Description
			1 - Enabled
foundry-802.1x-valid-lookup	7	integer	Specifies if 802.1x lookup is enabled: **0** - Disabled **1** - Enabled
foundry-MAC-based-VLAN-QOS	8	integer	Specifies the priority for MAC-based VLAN QOS: **0** - qos_priority_0 **1** - qos_priority_1 **2** - qos_priority_2 **3** - qos_priority_3 **4** - qos_priority_4 **5** - qos_priority_5 **6** - qos_priority_6 **7** - qos_priority_7
foundry-INM-Role-AOR-List	9	string	Specifies the list of Roles and Area of Responsibility (AOR) that are allowed for an Brocade Network Advisor user. These values are mapped to Brocade Network Advisor Roles and AORs when the user logs in. For example, to configure an Brocade Network Advisor user to have "Administrator" and "Report User" roles and "New York Region" and "Santa Clara Region" AORs, specify "NmRoles=Administrator, Report User; NmAORs=New York Region, Santa Clara Region". The keys "NmRoles" and "NmAORs" are delimited by semi colon (;) and the values for the keys are delimited by a comma (,).

Enabling SNMP traps for RADIUS

To enable SNMP traps for RADIUS on an Extreme device, you must execute the **enable snmp config-radius** command as shown in the following.

```
device(config)# enable snmp config-radius
```

The **enable snmp config-radius** command supports two keywords; **config-radius** and **config-tacacs**. The **config-radius** parameter specifies that traps will be enabled for RADIUS. Generation of Radius traps is disabled by default. The **config-tacacs** parameter specifies that traps will be enabled for TACACS. Generation of TACACS traps is disabled by default.

Identifying the RADIUS server to the device

To use a RADIUS server to authenticate access to an Extreme device, you must identify the server to the Extreme device.

```
device(config)# radius-server host 10.157.22.99
```

The **host** *ip-addr* | *server-name* parameter is either an IP address or an ASCII text string.

The *auth-port* parameter is the Authentication port number; it is an optional parameter. The default is 1812.

The *acct-port* parameter is the Accounting port number; it is an optional parameter. The default is 1813.

> **NOTE**
> The RADIUS server is supported only through the management VRF, the default-vrf, or through the out-of-band management port.

Specifying different servers for individual AAA functions

In a RADIUS configuration, you can designate a server to handle a specific AAA task. For example, you can designate one RADIUS server to handle authorization and another RADIUS server to handle accounting. You can specify individual servers for authentication and accounting, but not for authorization. You can set the RADIUS key for each server.

To specify different RADIUS servers for authentication and accounting, enter a command such as the following.

```
device(config)#
 radius-server host 10.2.3.4 auth-port 1812 acct-port 1813 authentication-only key abc
device(config)#
 radius-server host 10.2.3.6 auth-port 1812 acct-port 1813 accounting-only key ghi

device(config)#
 radius-server host 10.2.3.4 authentication-only key abc
```

The software adds a prefix to the authentication key in the configuration. For example,

```
radius-server host 10.2.3.6 auth-port 1812 acct-port 1813 default key 2 $D?@d=8
```

Configuring RadSec

The RADIUS protocoll, described in RFC 2865, RFC 2868 and RFC 2869, is a widely deployed authentication and authorization protocol. The supplementary RADIUS Accounting specification, described in RFC 2866 and RFC 2869, provides accounting mechanisms, thus delivering a full Authentication, Authorization, and Accounting (AAA) solution. The main focus of RadSec is to provide a means to secure the communication between RADIUS/TCP peers using TLS 1.2. To reinforce RADIUS, RadSec uses the secure transport layer protocol (TLS/SSL) in place of UDP for sending and receiving packets. RadSec wraps the entire RADIUS packet payload into a TLS stream and thus mitigates the risk of attacks. Because RadSec utilizes TCP, re-transmission of requests is not supported.

TLS/SSL protects confidential information using cryptography. Sensitive data is encrypted across public networks to achieve a high level of confidentiality. Primarily, PKI utilizes asymmetric cryptography that is considered more secure than symmetric cryptography.

Your environment may require that no default keys are permitted on a device, and that all keys/passwords must be configurable. In other cases, some deployments use different servers for administrative access and 802.1x access, and therefore utilize different keys. To this end, the RADIUS server keys can be unique and configured for each RADIUS server. Extreme NetIron supports up to eight RADIUS servers.

> **NOTE**
> Encryption of the RADIUS keys is completed by default.

The **ssl-auth-port** specifies that the server is a RadSec TLS-encrypted TCP session. Only one auth-port or ssl-auth-port can be specified. If neither is specified, it defaults to existing default behavior, which is to use the default auth-port of 1812 and 1813 for accounting with UDP transport.

To specify different RadSec servers for authentication and accounting, enter commands such as the following.

1. Download the local client certificate.

```
device(config)# scp rsacert2048_days1095_sha256_SAN.pem <user>@<ip>:sslclientcert
```

2. Download the local client key.

```
device(config)# scp rsakey2048.pem <user>@<ip>:sslclientprivkey
```

3. Configure the RADIUS server, including the accounting definition.

```
device(config)# radius-server host 10.25.105.44 ssl-auth-port 2083 default key abc
```

The **radius-server host** command supports IPv4, IPv6, and host names. When authenticating any device that is configured with an IPv4/IPv6 address, RadSec accesses the RADIUS servers in the order they were configured, regardless of the IP address family of the RadSec server.

When you display the configuration of the Extreme device, the RADIUS key is encrypted, as shown in the following example.

```
device(config)# radius-server key 1 abc
device(config)# write memory
...
radius-server host 10.2.3.5
radius key 1 $!2d
```

4. Configure the AAA method.

```
device(config)# aaa authentication login default radius
```

Radius health check

Radius health check pro actively polls the radius server and checks for the radius-server availability. If the checks fail, radius health check marks the status of the radius-server as not available. This feature is disabled by default.

To create an authentication-method list that specifies RADIUS as the primary authentication method for securing Telnet or ssh access to the CLI, enter the following command.

```
device(config)# radius-server host 10.2.3.4 auth-port 1812 acct-port 1813
device(config)# enable telnet authentication
device(config)# aaa authentication login default radius local
```

Global radius configuration

The following global configurations are for all radius servers, and can be used to configure defaults. If the individual radius servers are configured, the instance value takes precedence.

Health-check is disabled by default. Use the following command to globally enable health-check.

```
device(config)# radius-server enable-health-check
```

Use the **no radius-server enable-health-check** command to globally disable health-check.

Setting the poll time and dead time intervals

The poll interval sets how often the status-server packets are sent. The status-server packets are sent every 15minutes. The minimum that can be configured is 1 and maximum is 96. This translates to one status check for every 15m to one per day (4*24)

The dead time interval is the period for which we wait after declared dead or not reachable and before sending the status-server packet again. By default, it waits for 45m when server is declared dead, before sending again health checks. For example, you can configure one status check for every 15minutes to one per week(4*24*7). Use a command such as the following to enable the radius health check pool time interval and dead time inter

```
device(config)# radius-health-check  poll-time 5 dead-time 4
```

Setting RADIUS parameters

You can set the following parameters in a RADIUS configuration:

- **RADIUS key** - This parameter specifies the value that the Extreme device sends to the RADIUS server when trying to authenticate user access.
- **Retransmit interval** - This parameter specifies how many times the Extreme device will resend an authentication request when the RADIUS server does not respond. The retransmit value can be from 1 - 5 times. The default is 3 times.
- **Timeout** - This parameter specifies how many seconds the Extreme device waits for a response from a RADIUS server before either retrying the authentication request, or determining that the RADIUS servers are unavailable and moving on to the next authentication method in the authentication-method list. The timeout can be from 1 - 15 seconds. The default is 3 seconds.

Setting the RADIUS key

The **key** parameter in the **radius-server** command is used to encrypt RADIUS packets before they are sent over the network. The value for the **key** parameter on the Extreme device should match the one configured on the RADIUS server. The key length can be from 1 - 64 characters and cannot include any space characters.

Your environment may require that no default keys are permitted on a device, and that all keys/passwords must be configurable. In other cases, some deployments use different servers for administrative access and 802.1x access, and therefore utilize different keys. To this end, the RADIUS server keys can be unique and configured for each RADIUS server. Extreme NetIron supports up to eight RADIUS servers.

To specify a RADIUS server key, enter a command such as the following.

```
device(config)# radius-server key abc
```

When you display the configuration of the Extreme device, the RADIUS key is encrypted.

```
device(config)# radius-server key 1 abc
device(config)# write memory
...
radius-server host 10.2.3.5
radius key 1 $!2d
```

> **NOTE**
> Encryption of the RADIUS keys is done by default. The **0** parameter disables encryption. The **1** parameter is not required; it is provided for backwards compatibility.

Setting the retransmission limit

The **retransmit** parameter specifies the maximum number of retransmission attempts. When an authentication request times out, the software will retransmit the request up to the maximum number of retransmissions configured. The default retransmit value is 3 retries. The range of retransmit values is from 1 - 5.

To set the RADIUS retransmit limit, enter a command such as the following.

```
device(config)# radius-server retransmit 5
```

Setting the timeout parameter

The **timeout** parameter specifies how many seconds the Extreme device waits for a response from the RADIUS server before either retrying the authentication request, or determining that the RADIUS server is unavailable and moving on to the next authentication method in the authentication-method list. The timeout can be from 1 - 15 seconds. The default is 3 seconds.

```
device(config)# radius-server timeout 5
```

Configuring authentication-method lists for RADIUS

You can use RADIUS to authenticate Telnet or SSH access and access to Privileged EXEC level and CONFIG levels of the CLI. When configuring RADIUS authentication, you create authentication-method lists specifically for these access methods, specifying RADIUS as the primary authentication method.

Within the authentication-method list, RADIUS is specified as the primary authentication method and up to six backup authentication methods are specified as alternates. If RADIUS authentication fails due to an error, the device tries the backup authentication methods in the order they appear in the list.

When you configure authentication-method lists for RADIUS, you must create a separate authentication-method list for Telnet or SSH CLI access and for CLI access to the Privileged EXEC level and CONFIG levels of the CLI.

To create an authentication-method list that specifies RADIUS as the primary authentication method for securing Telnet access to the CLI, enter the following command.

```
device(config)#
 enable telnet authentication
device(config)# aaa authentication login default radius local
```

The commands above cause RADIUS to be the primary authentication method for securing Telnet access to the CLI. If RADIUS authentication fails due to an error with the server, local authentication is used instead.

To create an authentication-method list that specifies RADIUS as the primary authentication method for securing access to Privileged EXEC level and CONFIG levels of the CLI, enter the following command.

```
device(config)# aaa authentication enable default radius local none
```

The command above causes RADIUS to be the primary authentication method for securing access to Privileged EXEC level and CONFIG levels of the CLI. If RADIUS authentication fails due to an error with the server, local authentication is used instead. If local authentication fails, no authentication is used; the device automatically permits access.

Entering privileged EXEC mode after a Telnet or SSH login

By default, a user enters User EXEC mode after a successful login through Telnet or SSH. You can configure the device so that a user enters Privileged EXEC mode after a Telnet or SSH login. To do this, use the following command.

```
device(config)#
aaa authentication login privilege-mode
```

Syntax: [no] aaa authentication login privilege-mode

The user's privilege level is based on the privilege level granted during login.

Configuring enable authentication to prompt for password only

If Enable authentication is configured on the device, by default, a user is prompted for a username and password. when the user attempts to gain Super User access to the Privileged EXEC and CONFIG levels of the CLI. You can configure the Extreme device to prompt only for a password. The device uses the username (up to 48 characters) entered at login, if one is available. If no username was entered at login, the device prompts for both username and password.

To configure the Extreme device to prompt only for a password when a user attempts to gain Super User access to the Privileged EXEC and CONFIG levels of the CLI, enter the following command.

```
device(config)# aaa authentication enable implicit-user
```

Syntax: [no] aaa authentication enable implicit-user

Configuring RADIUS authorization

The Extreme device supports RADIUS authorization for controlling access to management functions in the CLI. Two kinds of RADIUS authorization are supported:

- Exec authorization determines a user's privilege level when they are authenticated
- Command authorization consults a RADIUS server to get authorization for commands entered by the user

Configuring Exec authorization

> NOTE
> Before you configure RADIUS exec authorization on an Extreme device, make sure that the **aaa authentication enable default radius** command exists in the configuration.

When RADIUS exec authorization is performed, the Extreme device consults a RADIUS server to determine the privilege level of the authenticated user.

To configure RADIUS exec authorization on an Extreme device, enter the following command.

```
device(config)# aaa authentication login default radius
device(config)# aaa authorization exec default radius
```

Syntax: [no] aaa authorization exec default radius | none

If you specify **none**, or omit the **aaa authorization exec** command from the device's configuration, no exec authorization is performed.

> NOTE
> If the **aaa authorization exec default radius** command exists in the configuration, following successful authentication the device assigns the user the privilege level specified by the brocade-privilege-level attribute received from the RADIUS server. If the **aaa authorization exec default radius** command does not exist in the configuration, then the value in the brocade-privilege-level attribute is ignored, and the user is granted Super User access.For the **aaa authorization exec default radius** command to work, either the **aaa authentication login default radius** command, or the **aaa authentication enable default radius** command must also exist in the configuration.

Configuring command authorization

When RADIUS command authorization is enabled, the Extreme device consults the list of commands supplied by the RADIUS server during authentication to determine whether a user can execute a command he or she has entered.

You enable RADIUS command authorization by specifying a privilege level whose commands require authorization. For example, to configure the Extreme device to perform authorization for the commands available at the Super User privilege level (that is; all commands on the device), enter the following command.

```
device(config)# aaa authorization commands 0 default radius
```

Syntax: [no] aaa authorization commands *privilege-level* default radius | tacacs+ | none

The **privilege-level** parameter can be one of the following:

- **0** - Authorization is performed (that is, the Extreme device looks at the command list) for commands available at the Super User level (all commands)
- **4** - Authorization is performed for commands available at the Port Configuration level (port-config and read-only commands)
- **5** - Authorization is performed for commands available at the Read Only level (read-only commands)

> **NOTE**
> RADIUS command authorization can be performed only for commands entered from Telnet or SSH sessions, or from the console. No authorization is performed for commands entered at the Web Management Interface or Brocade Network Advisor.

> **NOTE**
> Since RADIUS command authorization relies on the command list supplied by the RADIUS server during authentication, you cannot perform RADIUS authorization without RADIUS authentication.

Command authorization and accounting for console commands

The Extreme devices support command authorization and command accounting for CLI commands entered at the console. To configure the device to perform command authorization and command accounting for console commands, enter the following.

```
device(config)# enable aaa console
```

Syntax: [no] enable aaa console

> ⚠️ **CAUTION**
> **If you have previously configured the device to perform command authorization using a RADIUS server, entering the enable aaa console command may prevent the execution of any subsequent commands entered on the console.**

> **NOTE**
> This happens because RADIUS command authorization requires a list of allowable commands from the RADIUS server. This list is obtained during RADIUS authentication. For console sessions, RADIUS authentication is performed only if you have configured Enable authentication and specified RADIUS as the authentication method (for example, with the **aaa authentication enable default radius** command). If RADIUS authentication is never performed, the list of allowable commands is never obtained from the RADIUS server. Consequently, there would be no allowable commands on the console.

Configuring RADIUS accounting

The Extreme devices support RADIUS accounting for recording information about user activity and system events. When you configure RADIUS accounting on an Extreme device, information is sent to a RADIUS accounting server when specified events occur, such as when a user logs into the device or the system is rebooted.

Configuring RADIUS accounting for Telnet or SSH (shell) access

To send an Accounting Start packet to the RADIUS accounting server when an authenticated user establishes a Telnet or SSH session on the Extreme device, and an Accounting Stop packet when the user logs out, enter the following command.

```
device(config)# aaa accounting exec default start-stop radius
```

Syntax: [no] aaa accounting exec default start-stop radius | tacacs+ | none

Configuring RADIUS accounting for CLI commands

You can configure RADIUS accounting for CLI commands by specifying a privilege level whose commands require accounting. For example, to configure an Extreme device to perform RADIUS accounting for the commands available at the Super User privilege level (that is; all commands on the device), enter the following command.

```
device(config)# aaa accounting commands 0 default start-stop radius
```

An Accounting Start packet is sent to the RADIUS accounting server when a user enters a command, and an Accounting Stop packet is sent when the service provided by the command is completed.

> **NOTE**
> If authorization is enabled, and the command requires authorization, then authorization is performed before accounting takes place. If authorization fails for the command, no accounting takes place.

Syntax: [no] aaa accounting commands privilege-level default start-stop radius | tacacs | none

The **privilege-level** parameter can be one of the following:

- **0** - Records commands available at the Super User level (all commands)
- **4** - Records commands available at the Port Configuration level (port-config and read-only commands)
- **5** - Records commands available at the Read Only level (read-only commands)

Configuring RADIUS accounting for system events

You can configure RADIUS accounting to record when system events occur on an Extreme device. System events include rebooting and when changes to the active configuration are made.

The following command causes an Accounting Start packet to be sent to the RADIUS accounting server when a system event occurs, and a Accounting Stop packet to be sent when the system event is completed.

```
device(config)# aaa accounting system default start-stop radius
```

Syntax: [no] aaa accounting system default start-stop radius | tacacs+ | none

Configuring an interface as the source for all RADIUS packets

You can designate the lowest-numbered IP address configured an Ethernet port, loopback interface, or virtual interface as the source IP address for all RADIUS packets from the Extreme device. Identifying a single source IP address for RADIUS packets provides the following benefits:

- If your RADIUS server is configured to accept packets only from specific links or IP addresses, you can use this feature to simplify configuration of the RADIUS server by configuring the Extreme device to always send the RADIUS packets from the same link or source address.
- If you specify a loopback interface as the single source for RADIUS packets, RADIUS servers can receive the packets regardless of the states of individual links. Thus, if a link to the RADIUS server becomes unavailable but the client or server can

be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

The software contains separate CLI commands for specifying the source interface for Telnet, TACACS or TACACS+, and RADIUS packets. You can configure a source interface for one or more of these types of packets.

To specify an Ethernet or a loopback or virtual interface as the source for all RADIUS packets from the device, use the following CLI method. The software uses the lowest-numbered IP address configured on the port or interface as the source IP address for RADIUS packets originated by the device.

To specify the lowest-numbered IP address configured on a virtual interface as the device's source for all RADIUS packets, enter commands such as the following.

```
device(config)# int ve 1
device(config-vif-1)# ip address 10.0.0.3/24
device(config-vif-1)# exit
device(config)# ip radius source-interface ve 1
```

The commands in this example configure virtual interface 1, assign IP address 10.0.0.3/24 to the interface, then designate the interface as the source for all RADIUS packets from the Extreme device.

Syntax: [no] ip radius source-interface ethernet portnum | loopback num | ve num

The *num* parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the *portnum* is the port's number (including the slot number, if you are configuring a device).

> NOTE
> The NAS-IP-ADDR attribute is added into the RADIUS Access-Request when **ip radius source-interface** command is configured or when the Access-Request is for IPv4 RADIUS server.

Configuring an IPv6 interface as the source for all RADIUS packets

Use the **ipv6 radius source-interface** command to specify the IPv6 address of the interface that is chosen for the NAS-IPv6-Attribute. This feature is applicable only if an IPv6 interface is configured and authentication happens through RADIUS.

```
device(config)# int ve 1
device(config-vif-1)# ipv6 address
2001:DB8::2004
device(config-vif-1)# exit
device(config)# ipv6 radius source-interface ve 1
```

Syntax: [no] ipv6 radius source-interface ethernet *port-num* | loopback *num* | ve *num*

The *num* parameter is a loopback interface or virtual interface number. If you specify an Ethernet port, the *portnum* is the port's number (including the slot number, if you are configuring a device).

The [**no**] option removes the configuration.

This command configures the designate interface *port-num* or num as the source for all RADIUS packets from the device.

> NOTE
> The NAS-IPv6-ADDR attribute is added into the RADIUS Access-Request when **ipv6 radius source-interface** command is configured or when the Access-Request is for IPv6 RADIUS server.

Displaying RADIUS configuration information

The **show aaa** command displays information about all TACACS or TACACS+ and RADIUS servers identified on the device.

```
***** TACACS server not configured
Radius default key: ...
Radius retries: 3
Radius timeout: 3 seconds
IPv4 Radius source-interface: loopback 1
IPv6 Radius source-interface: loopback 1
Radius Server:  IP=10.25.105.201 Auth Port=1812 Acct Port=1813 Usage=any
                Key=...
                opens=0 closes=0 timeouts=0 errors=0
                packets in=0 packets out=6
                Health-check=disabled dead-time-interval=45  auto-authenticate-time-interval=30 available
                IPv4 Radius Source address:    IP=172.26.65.207              IPv6 Radius Source
Address:     IP=2001:DB8::18
no connection
Radius Server:  IP=fe80::7ae7:d1ff:fe8d:1b82 Auth Port=1812 Acct Port=1813 Usage=any
                Key=...
                opens=0 closes=0 timeouts=0 errors=0
                packets in=0 packets out=0
                Health-check=disabled dead-time-interval=45  auto-authenticate-time-interval=30 available
                IPv4 Radius Source address:    IP=172.26.65.207              IPv6 Radius Source
Address:     IP=2001:DB8::18
no connection
```

Syntax: show aaa

The following table describes the RADIUS information displayed by the **show aaa** command.

Field	Description
Radius default key	The setting configured with the **radius-server key** command. At the Super User privilege level, the actual text of the key is displayed. At the other privilege levels, a string of periods (....) is displayed instead of the text.
Radius retries	The setting configured with the **radius-server retransmit** command.
Radius timeout	The setting configured with the **radius-server timeout** command.
IPv4 Radius source-interface	The setting configured with the **ip radius source-interface** command.
IPv6 Radius source-interface	The setting configured with the **ipv6 radius source-interface** command.
Radius Server	For each RADIUS server, the IP address, and the following statistics are displayed: **Auth Port** - RADIUS authentication port number (default 1645) **Acct Port** - RADIUS accounting port number (default 1646) **opens** - Number of times the port was opened for communication with the server **closes** - Number of times the port was closed normally **timeouts** - Number of times port was closed due to a timeout **errors** - Number of times an error occurred while opening the port **packets in** - Number of packets received from the server **packets out** - Number of packets sent to the server
connection	The current connection status. This can be "no connection" or "connection active".

The **show web** command displays the privilege level of Web Management Interface users.

```
device(config)# show web
User                           Privilege       IP address
set                                0            192.168.1.234
```

Syntax: show web

TACACS+ Server Authentication

Configuring TACACS or TACACS+ security

You can use the security protocol Terminal Access Controller Access Control System (TACACS) or TACACS+ to authenticate the following kinds of access to the Extreme devices:

* Telnet access
* SSH access
* Console access
* Web management access
* Access to the Privileged EXEC level and CONFIG levels of the CLI

   NOTE
   You cannot authenticate Brocade Network Advisor (SNMP) access to an Extreme device using TACACS or TACACS+.

The TACACS and TACACS+ protocols define how authentication, authorization, and accounting information is sent between an Extreme device and an authentication database on a TACACS or TACACS+ server. TACACS or TACACS+ services are maintained in a database, typically on a UNIX workstation or PC with a TACACS or TACACS+ server running.

How TACACS+ differs from TACACS

TACACS is a simple UDP-based access control protocol originally developed by BBN for MILNET. TACACS+ is an enhancement to TACACS and uses TCP to ensure reliable delivery.

TACACS+ is an enhancement to the TACACS security protocol. TACACS+ improves on TACACS by separating the functions of authentication, authorization, and accounting (AAA) and by encrypting all traffic between the Extreme device and the TACACS+ server. TACACS+ allows for arbitrary length and content authentication exchanges, which allow any authentication mechanism to be utilized with the Extreme device. TACACS+ is extensible to provide for site customization and future development features. The protocol allows the Extreme device to request very precise access control and allows the TACACS+ server to respond to each component of that request.

   NOTE
   TACACS+ provides for authentication, authorization, and accounting, but an implementation or configuration is not required to employ all three.

TACACS or TACACS+ authentication, authorization, and accounting

When you configure an Extreme device to use a TACACS or TACACS+ server for authentication, the device prompts users who are trying to access the CLI for a user name and password, then verifies the password with the TACACS or TACACS+ server.

If you are using TACACS+, it is recommended that you also configure *authorization*, in which the Extreme device consults a TACACS+ server to determine which management privilege level (and which associated set of commands) an authenticated user is allowed to use. You can also optionally configure *accounting*, which causes the Extreme device to log information on the TACACS+ server when specified events occur on the device.

NOTE

By default, a user logging into the device through Telnet or SSH would first enter the User EXEC level. The user can enter the **enable** command to get to the Privileged EXEC level.

NOTE

A user that is successfully authenticated can be automatically placed at the Privileged EXEC level after login. Refer to Entering privileged EXEC mode after a console Telnet or SSH login on page 71.

TACACS authentication

NOTE

Also, multiple challenges are supported for TACACS+ login authentication.

The following events occur when TACACS authentication takes place.

1. A user attempts to gain access to the Extreme device by doing one of the following:

   - Logging into the device using console, Telnet, SSH, or the Web Management Interface.
   - Entering the Privileged execution level or configuration level of the CLI.

2. The user is prompted for a username and password.

3. The user enters a username and password.

4. The Extreme device sends a request containing the username and password to the TACACS server.

5. The username and password are validated in the TACACS server's database.

6. If the password is valid, the user is authenticated.

TACACS+ authentication

The following events occur when TACACS+ authentication takes place.

1. A user attempts to gain access to the Extreme device by doing one of the following:

   - – Logging into the device using console, telnet, SSH, or the Web Management Interface
     – Entering the Privileged EXEC level or CONFIG level of the CLI

2. The user is prompted for a username.

3. The user enters a username.

4. The Extreme device obtains a password prompt from a TACACS+ server.

5. The user is prompted for a password.

6. The user enters a password.

7. The Extreme device sends the password to the TACACS+ server.

8. The password is validated in the TACACS+ server's database.

9. If the password is valid, the user is authenticated.

TACACS+ authorization

The Extreme devices support two kinds of TACACS+ authorization:

- Exec authorization determines a user's privilege level when they are authenticated.
- Command authorization consults a TACACS+ server to get authorization for commands entered by the user.

The following events occur when TACACS+ exec authorization takes place.

1. A user logs into the Extreme device using console, Telnet, SSH, or the Web Management Interface
2. The user is authenticated.
3. The Extreme device consults the TACACS+ server to determine the privilege level of the user.
4. The TACACS+ server sends back a response containing an A-V (Attribute-Value) pair with the privilege level of the user.
5. The user is granted the specified privilege level.

   The following events occur when TACACS+ command authorization takes place.

   1. A Telnet, SSH, or console interface user previously authenticated by a TACACS+ server enters a command on the Extreme device.
   2. The Extreme device looks at its configuration to see if the command is at a privilege level that requires TACACS+ command authorization.
   3. If the command belongs to a privilege level that requires authorization, the Extreme device consults the TACACS+ server to see if the user is authorized to use the command.
   4. If the user is authorized to use the command, the command is executed.

TACACS+ accounting

The following steps explain the working of TACACS+ accounting.

1. One of the following events occur on the Extreme device:
   - – A user logs into the management interface using console, Telnet or SSH
     – A user enters a command for which accounting has been configured
     – A system event occurs, such as a reboot or reloading of the configuration file
2. The Extreme device checks its configuration to see if the event is one for which TACACS+ accounting is required.
3. If the event requires TACACS+ accounting, the Extreme device sends a TACACS+ Accounting Start packet to the TACACS+ accounting server, containing information about the event.
4. The TACACS+ accounting server acknowledges the Accounting Start packet.
5. The TACACS+ accounting server records information about the event.
6. When the event is concluded, the Extreme device sends an Accounting Stop packet to the TACACS+ accounting server.
7. The TACACS+ accounting server acknowledges the Accounting Stop packet.

AAA operations for TACACS or TACACS+

The following table lists the sequence of authentication, authorization, and accounting operations that take place when a user gains access to an Extreme device that has TACACS or TACACS+ security configured.

User action	Applicable AAA operations
User attempts to gain access to the Privileged EXEC and CONFIG levels of the CLI	Enable authentication: aaa authentication enable default *method-list*
	Exec authorization (TACACS+): aaa authorization exec default tacacs+
	System accounting start (TACACS+): aaa accounting system default start-stop *method-list*
User logs in using console, Telnet, or SSH	Login authentication:

User action	Applicable AAA operations
	aaa authentication login default *method-list*
	Exec authorization (TACACS+):
	aaa authorization exec default tacacs+
	Exec accounting start (TACACS+):
	aaa accounting exec default *method-list*
	System accounting start (TACACS+):
	aaa accounting system default start-stop *method-list*
User logs into the Web Management Interface	Web authentication:
	aaa authentication web-server default *method-list*
	Exec authorization (TACACS+):
	aaa authorization exec default tacacs+
User logs out of console, Telnet, or SSH session	Command accounting (TACACS+):
	aaa accounting commands *privilege-level* default start-stop *method-list*
	EXEC accounting stop (TACACS+):
	aaa accounting exec default start-stop *method-list*
User enters system commands (for example, **reload** , **boot system**)	Command authorization (TACACS+):
	aaa authorization commands *privilege-level* default *method-list*
	Command accounting (TACACS+):
	aaa accounting commands *privilege-level* default start-stop *method-list*
	System accounting stop (TACACS+):
	aaa accounting system default start-stop *method-list*
User enters the command: [no] aaa accounting system defaultstart-stop *method-list*	Command authorization (TACACS+):
	aaa authorization commands *privilege-level* default *method-list*
	Command accounting (TACACS+):
	aaa accounting commands *privilege-level* default start-stop *method-list*
	System accounting start (TACACS+):
	aaa accounting system default start-stop *method-list*
User enters other commands	Command authorization (TACACS+):
	aaa authorization commands *privilege-level* default *method-list*
	Command accounting (TACACS+):
	aaa accounting commands *privilege-level* default start-stop *method-list*

AAA Security for commands pasted Into the running configuration

If AAA security is enabled on an Extreme device, commands pasted into the running configuration are subject to the same AAA operations as if they were entered manually.

When you paste commands into the running configuration, and AAA command authorization or accounting is configured on the device, AAA operations are performed on the pasted commands. The AAA operations are performed before the commands are actually added to the running configuration. The server performing the AAA operations should be reachable when you paste the commands into the running configuration file. If the device determines that a pasted command is invalid, AAA operations are halted on the remaining commands. The remaining commands may not be executed if command authorization is configured.

TACACS or TACACS+ configuration considerations

Consider the following for configuring TACACS or TACACS+ servers:

- You must deploy at least one TACACS or TACACS+ server in your network.
- The Extreme device supports authentication using up to eight TACACS or TACACS+ servers. The device tries to use the servers in the order you add them to the device's configuration.
- You can select only one primary authentication method for each type of access to a device (CLI through Telnet, CLI Privileged EXEC and CONFIG levels). For example, you can select TACACS+ as the primary authentication method for Telnet CLI access, but you cannot also select RADIUS authentication as a primary method for the same type of access. However, you can configure backup authentication methods for each access type.
- You can configure the Extreme device to authenticate using a TACACS or TACACS+ server, not both.

TACACS configuration procedure

Use the following procedure for TACACS configurations.

1. Enable TACACS. Enabling SNMP traps for TACACS on page 67.
2. Identify TACACS servers. Refer to Identifying the TACACS or TACACS+ servers on page 68.
3. Set optional parameters. Refer to XMR Series and MLX Series: Setting optional TACACS or TACACS+ parameters on page 69.
4. Configure authentication-method lists. Refer to Configuring authentication-method lists for TACACS or TACACS+ on page 70.

TACACS+ configuration procedure

Use the following procedure for TACACS+ configurations.

1. Enable TACACS. Enabling SNMP traps for TACACS on page 67
2. Identify TACACS+ servers. Refer to Identifying the TACACS or TACACS+ servers on page 68.
3. Set optional parameters. Refer to XMR Series and MLX Series: Setting optional TACACS or TACACS+ parameters on page 69.
4. Configure authentication-method lists. Refer to Configuring authentication-method lists for TACACS or TACACS+ on page 70.
5. Optionally configure TACACS+ authorization. Refer to Configuring TACACS+ authorization on page 73.
6. Optionally configure TACACS+ accounting. Refer to Configuring TACACS+ accounting on page 75.

Enabling SNMP traps for TACACS

To enable SNMP access to the TACACS MIB objects on an Extreme device, you must execute the enable **snmp config-tacacs** command as shown in the following.

```
device(config)# enable snmp config-tacacs
```

Syntax: [no] **enable snmp** [**config-radius** | **config-tacacs**]

The **config-radius** parameter specifies the MIBs accessible for RADIUS. Generation of Radius traps is disabled by default.

The **config-tacacs** parameter specifies the MIBs accessible for TACACS. Generation of TACACS traps is disabled by default.

Identifying the TACACS or TACACS+ servers

To use TACACS or TACACS+ servers to authenticate access to an Extreme device, you must identify the servers to the Extreme device.

For example, to identify three TACACS or TACACS+ servers, enter commands such as the following.

```
device(config)# tacacs-server host 10.94.6.161
device(config)# tacacs-server host 10.94.6.191
device(config)# tacacs-server host 10.94.6.122
```

Syntax: [no] tacacs-server host ip-addr | hostname [auth-port number]

The *ip-addr |hostname* parameter specifies the IP address or host name of the server. You can enter up to eight **tacacs-server host** commands to specify up to eight different servers.

> **NOTE**
> To specify the server's host name instead of its IP address, you must first identify a DNS server using the **ip dns server-address** *ip-addr* command at the global CONFIG level.

If you add multiple TACACS or TACACS+ authentication servers to the Extreme device, the device tries to reach them in the order you add them. For example, if you add three servers in the following order, the software tries the servers in the same order.

1. 10.94.6.161

2. 10.94.6.191

3. 10.94.6.122

   You can remove a TACACS or TACACS+ server by entering **no** followed by the **tacacs-server** command. For example, to remove 10.94.6.161, enter the following command.

   ```
   device(config)# no tacacs-server host 10.94.6.161
   ```

   > **NOTE**
   > If you erase a **tacacs-server** command (by entering "**no**" followed by the command), make sure you also erase the **aaa** commands that specify TACACS or TACACS+ as an authentication method. (Refer to Configuring authentication-method lists for TACACS or TACACS+ on page 70.) Otherwise, when you exit from the CONFIG mode or from a Telnet session, the system continues to believe it is TACACS or TACACS+ enabled and you will not be able to access the system.

   The **auth-port** parameter specifies the UDP (for TACACS) or TCP (for TACACS+) port number of the authentication port on the server. The default port number is 49.

Specifying different servers for individual AAA TACACS functions

In a TACACS+ configuration, you can designate a server to handle a specific AAA task. For example, you can designate one TACACS+ server to handle authorization and another TACACS+ server to handle accounting. You can set the TACACS+ key for each server.

To specify different TACACS+ servers for authentication, authorization, and accounting, enter a command such as the following.

```
device(config)#
 tacacs-server host 1.2.3.4 auth-port 49 authentication-only key abc
device(config)#
 tacacs-server host 1.2.3.5 auth-port 49 authorization-only key define
device(config)#
 tacacs-server host 1.2.3.6 auth-port 49 accounting-only key ghi
```

Syntax: [no] tacacs-server host ip-addr | server-name [auth-port number [authentication-only | authorization-only | accounting-only | default] [key string]]

The **host***ip-addr* | *server-name* parameter is either an IP address or an ASCII text string.

The **auth-port** *number* parameter is the Authentication port number; it is an optional parameter.

Enter **accounting-only** if the server is used only for TACACS accounting. Enter **authentication-only** if the server is used only for TACACS authentication. Enter **authorization-only** if the server is used only for TACAC authorization. Entering the **default** parameter causes the server to be used for all AAA TACACS functions.

After authentication takes place, the server that performed the authentication is used for authorization, accounting or both. If the authenticating server cannot perform the requested function, then the next server in the configured list of servers is tried; this process repeats until either a server that can perform the requested function is found, or every server in the configured list has been tried.

Enter **key** and configure a key for the server if an authentication key is to be used. By default, **key** is encrypted. If you want key to be in clear text, insert a **0** between **key** and *string*.

```
device(config)#
  tacacs-server host 10.2.3.5 auth-port 49 authorization-only key 0  report
```

The software adds a prefix to the authentication key string in the configuration. For example,

```
tacacs-server host 10.2.3.6 auth-port 49 authorization-only key  $D?@d=8
```

The prefix can be one of the following:

- 0 = the key string is not encrypted and is in clear text
- 1 = the key string uses proprietary simple crytographic 2-way algorithm

XMR Series and MLX Series: Setting optional TACACS or TACACS+ parameters

You can set the following optional parameters in a TACACS or TACACS+ configuration:

- **TACACS+ key** - This parameter specifies the value that the Extreme device sends to the TACACS+ server when trying to authenticate user access.
- **Retransmit interval** - This parameter specifies how many times the Extreme device will resend an authentication request when the TACACS or TACACS+ server does not respond. The retransmit value can be from 1 - 5 times. The default is 3 times.
- **Dead time** - This parameter specifies how long the Extreme device waits for the primary authentication server to reply before deciding the server is dead and trying to authenticate using the next server. The dead-time value can be from 1 - 5 seconds. The default is 3 seconds.
- **Timeout** - This parameter specifies how many seconds the Extreme device waits for a response from a TACACS or TACACS+ server before either retrying the authentication request, or determining that the TACACS or TACACS+ servers are unavailable and moving on to the next authentication method in the authentication-method list. The timeout can be from 1 - 15 seconds. The default is 3 seconds.

Setting the TACACS+ key

The **key** parameter in the **tacacs-server** command is used to encrypt TACACS+ packets before they are sent over the network. The value for the **key** parameter on the Extreme device should match the one configured on the TACACS+ server. The key length can be from 1 - 64 characters and cannot include any space characters.

> NOTE
> The **tacacs-server key** command applies only to TACACS+ servers, not to TACACS servers. If you are configuring TACACS, do not configure a key on the TACACS server and do not enter a key on the Extreme device.

To specify a TACACS+ server key, enter the following command.

```
device(config)# tacacs-server key rkwong
```

Syntax: [no] tacacs-server key [0 | 1] string

When you display the configuration of the Extreme device, the TACACS+ keys are encrypted.

```
device(config)#
 tacacs-server key 1 abc
device(config)# write terminal
...
tacacs-server host 10.2.3.5 auth-port 49
tacacs key 1 $!2d
```

> **NOTE**
> Encryption of the TACACS+ keys is done by default. The **0** parameter disables encryption. The **1** parameter is not required; it is provided for backwards compatibility.

Setting the retransmission limit

The **retransmit** parameter specifies how many times the Extreme device will resend an authentication request when the TACACS or TACACS+ server does not respond. The retransmit limit can be from 1 - 5 times. The default is 3 times.

To set the TACACS or TACACS+ retransmit limit, enter the following command.

```
device(config)# tacacs-server retransmit 5
```

Syntax: [no] tacacs-server retransmit number

Setting the timeout parameter

The **timeout** parameter specifies how many seconds the Extreme device waits for a response from the TACACS or TACACS+ server before either retrying the authentication request, or determining that the TACACS or TACACS+ server is unavailable and moving on to the next authentication method in the authentication-method list. The timeout can be from 1 - 15 seconds. The default is 3 seconds.

```
device(config)# tacacs-server timeout 5
```

Syntax: [no] tacacs-server timeout number

Configuring authentication-method lists for TACACS or TACACS+

You can use TACACS or TACACS+ to authenticate console, Telnet, or SSH access and access to Privileged EXEC level and CONFIG levels of the CLI. When configuring TACACS or TACACS+ authentication, you create authentication-method lists specifically for these access methods, specifying TACACS or TACACS+ as the primary authentication method.

Within the authentication-method list, TACACS or TACACS+ is specified as the primary authentication method and up to six backup authentication methods are specified as alternates. If TACACS or TACACS+ authentication fails due to an error, the device tries the backup authentication methods in the order they appear in the list. If a TACACS or TACACS+ server responds with a reject for a user, the system does not try the backup authentication methods.

When you configure authentication-method lists for TACACS or TACACS+ authentication, you must create a separate authentication-method list for Telnet or SSH CLI access, and for access to the Privileged EXEC level and CONFIG levels of the CLI.

To create an authentication-method list that specifies TACACS or TACACS+ as the primary authentication method for securing Telnet or SSH access to the CLI.

```
device(config)# enable telnet authentication
device(config)# aaa authentication login default tacacs+ local
```

> **NOTE**
> To enable AAA support for commands entered at the console you must follow the procedure described in the Configuring AAA on the console task.

The commands above cause TACACS or TACACS+ to be the primary authentication method for securing Telnet or SSH access to the CLI. If TACACS or TACACS+ authentication fails due to an error with the server, authentication is performed using local user accounts instead.

To create an authentication-method list that specifies TACACS or TACACS+ as the primary authentication method for securing access to Privileged EXEC level and CONFIG levels of the CLI.

```
device(config)# aaa authentication enable default tacacs+ local none
```

The command above causes TACACS or TACACS+ to be the primary authentication method for securing access to Privileged EXEC level and CONFIG levels of the CLI. If TACACS or TACACS+ authentication fails due to an error with the server, local authentication is used instead. If local authentication fails, no authentication is used; the device automatically permits access.

Entering privileged EXEC mode after a console Telnet or SSH login

By default, a user enters User EXEC mode after a successful login using a non-AAA method through console, Telnet or SSH. Optionally, you can configure the device so that a user enters Privileged EXEC mode after a console, Telnet or SSH login. To do this, use the following command.

```
device(config)#
aaa authentication login privilege-mode
```

Syntax: [no] aaa authentication login privilege-mode

The user's privilege level is based on the privilege level granted during login.

Limitations when automatically entering privilege EXEC mode for SSH session with public-key authentication

- Features that require user identity will continue to behave as if no user identity was provided.
- The authentication, authorization and accounting will not be performed through AAA.

Enabling automatically entering Privilege EXEC mode access for SSH session with public-key authentication

1:

```
device (config) # aaa authentication login default local
device (config) # aaa authentication login privilege-mode
```

> **NOTE**
> After successful key-authentication, the SSH session will be placed into the Privileged EXEC mode.

2:

```
device (config) # aaa authentication enable default local
device (config) # aaa authentication login privilege-mode
device (config) # ip ssh password-authentication no
device (config) # ip ssh interactive-authentication no
```

NOTE
After successful key-authentication, the SSH session will be placed into the privileged EXEC mode.

3:

```
device (config) # aaa authentication login privilege-mode
device (config) # ip ssh permit-empty-passwd yes
```

NOTE
After successful key-authentication, the SSH session will be placed into the privileged EXEC mode.

4:

```
device (config) # aaa authentication login privilege-mode
device (config) # ip ssh key-authentication no
device (config) # ip ssh password-authentication yes
device (config) # ip ssh interactive-authentication yes
```

NOTE
An authenticated SSH session using either password or interactive authentication will be placed into the privileged EXEC mode.

Disabling automatically entering Privilege EXEC mode access for SSH session with public-key authentication

1:

```
device (config) # aaa authentication login default local
device (config) # no aaa authentication login privilege-mode
```

NOTE
After successful key-authentication, the SSH session will be placed into the User EXEC mode.

Syntax: :[no] aaa authentication login privilege-mode

Configuring enable authentication to use enable password on TACACS+

TACACS+ server allows a common enable password to be configured on the TACACS+ server. To allow a user to authenticate against that enable password, instead of the login password, use this command.

```
device(config)# aaa authentication enable implicit-user
```

Syntax: [no] aaa authentication enable implicit-user

Telnet or SSH prompts when the TACACS+ server is unavailable

When TACACS+ is the first method in the authentication method list, the device displays the login prompt received from the TACACS+ server. If a user attempts to login through Telnet or SSH, but none of the configured TACACS+ servers are available, the following takes place:

- If the next method in the authentication method list is "enable", the login prompt is skipped, and the user is prompted for the Enable password (that is, the password configured with the **enable super-user-password** command).
- If the next method in the authentication method list is "line", the login prompt is skipped, and the user is prompted for the Line password (that is, the password configured with the **enable telnet password** command).

Configuring TACACS+ authorization

The Extreme device supports TACACS+ authorization for controlling access to management functions in the CLI. Two kinds of TACACS+ authorization are supported:

- Exec authorization determines a user's privilege level when they are authenticated
- Command authorization consults a TACACS+ server to get authorization for commands entered by the user

Configuring exec authorization

When TACACS+ exec authorization is performed, the Extreme device consults a TACACS+ server to determine the privilege level of the authenticated user.

To configure TACACS+ exec authorization on an Extreme device, enter the following command.

```
device(config)# aaa authorization exec default tacacs+
```

Syntax: aaa authorization exec default tacacs+ | radius | none

If you specify **none**, or omit the **aaa authorization** exec command from the device's configuration, no exec authorization is performed.

A user's privilege level is obtained from the TACACS+ server in the "foundry-privlvl" A-V pair. If the **aaa authorization exec default tacacs** command exists in the configuration, the device assigns the user the privilege level specified by this A-V pair. If the command does not exist in the configuration, then the value in the "foundry-privlvl" A-V pair is ignored, and the user is granted Super User access.

> **NOTE**
> If the **aaa authorization exec default tacacs+** command exists in the configuration, following successful authentication the device assigns the user the privilege level specified by the "foundry-privlvl" A-V pair received from the TACACS+ server. If the **aaa authorization exec default tacacs+** command does not exist in the configuration, then the value in the "foundry-privlvl" A-V pair is ignored, and the user is granted Super User access. Also note that in order for the **aaa authorization exec default tacacs+** command to work, either the **aaa authentication enable default tacacs+** command, or the **aaa authentication login default tacacs+** command must also exist in the configuration.

Configuring an attribute-value pair on the TACACS+ server

During TACACS+ exec authorization, the Extreme device expects the TACACS+ server to send a response containing an A-V (Attribute-Value) pair that specifies the privilege level of the user. When the Extreme device receives the response, it extracts an A-V pair configured for the Exec service and uses it to determine the user's privilege level.

To set a user's privilege level, you can configure the "foundry-privlvl" A-V pair for the Exec service on the TACACS+ server.

```
user=bob {
   default service = permit
   member admin
   # Global password
   global = cleartext "cat"
   service = exec {
     foundry-privlvl = 0
        }
}
```

In this example, the A-V pair `foundry-privlvl = 0` grants the user full read-write access. The value in the foundry-privlvl A-V pair is an integer that indicates the privilege level of the user. Possible values are 0 for super-user level, 4 for port-config level, or 5 for read-only level. If a value other than 0, 4, or 5 is specified in the foundry-privlvl A-V pair, the default privilege level of 5 (read-only) is used. The foundry-privlvl A-V pair can also be embedded in the group configuration for the user. Refer to your TACACS+ documentation for the configuration syntax relevant to your server.

If the foundry-privlvl A-V pair is not present, the Extreme device extracts the last A-V pair configured for the Exec service that has a numeric value. The Extreme device uses this A-V pair to determine the user's privilege level.

```
user=bob {
    default service = permit
    member admin
    # Global password
    global = cleartext "cat"
    service = exec {
      priv-lvl = 15
        }
}
```

The attribute name in the A-V pair is not significant; the Extreme device uses the last one that has a numeric value. However, the Extreme device interprets the value for a non-"foundry-privlvl" A-V pair differently than it does for a "foundry-privlvl" A-V pair. The following table lists how the Extreme device associates a value from a non-"foundry-privlvl" A-V pair with an Extreme privilege level.

Value for non-"foundry-privlvl" A-V pair	Privilege level
15	0 (super-user)
From 14 - 1	4 (port-config)
Any other number or 0	5 (read-only)

In the example above, the A-V pair configured for the Exec service is `priv-lvl = 15`. The Extreme device uses the value in this A-V pair to set the user's privilege level to 0 (super-user), granting the user full read-write access.

In a configuration that has both a "foundry-privlvl" A-V pair and a non-"foundry-privlvl" A-V pair for the Exec service, the non-"foundry-privlvl" A-V pair is ignored.

```
user=bob {
    default service = permit
    member admin
    # Global password
    global = cleartext "cat"
    service = exec {
      foundry-privlvl = 4
      priv-lvl = 15
        }
}
```

In this example, the user would be granted a privilege level of 4 (port-config level). The `privlvl = 15` A-V pair is ignored by the Extreme device.

If the TACACS+ server has no A-V pair configured for the Exec service, the default privilege level of 5 (read-only) is grated to the user.

Configuring command authorization

When TACACS+ command authorization is enabled, the device consults a TACACS+ server to get authorization for commands entered by the user.

You enable TACACS+ command authorization by specifying a privilege level whose commands require authorization. For example, to configure the Extreme device to perform authorization for the commands available at the Super User privilege level (that is, all commands on the device), enter the following command.

```
device(config)# aaa authorization commands 0 default tacacs+
```

Syntax: [no] aaa authorization commands privilege-level default tacacs+ | radius | none

The *privilege-level* parameter can be one of the following:

- **0** - Authorization is performed for commands available at the Super User level (all commands)

- 4 - Authorization is performed for commands available at the Port Configuration level (port-config and read-only commands)
- 5 - Authorization is performed for commands available at the Read Only level (read-only commands)

NOTE
TACACS+ command authorization can be performed only for commands entered from Telnet or SSH sessions, or from the console. No authorization is performed for commands entered at the Web Management Interface or Brocade Network Advisor.

TACACS+ command authorization is not performed for the following commands:

- **At all levels: exit** , **logout**, **end**, **quit** and **access-list** (Any command with "acc" prefix).
- **At the Privileged EXEC level: enable** or **enable**text, where *text* is the password configured for the Super User privilege level.

If configured, command accounting is performed for these commands.

NOTE
To enable AAA support for commands entered at the console you must follow the procedure described in the Configuring AAA on the console task.

Configuring TACACS+ accounting

The Extreme device supports TACACS+ accounting for recording information about user activity and system events. When you configure TACACS+ accounting on an Extreme device, information is sent to a TACACS+ accounting server when specified events occur, such as when a user logs into the device or the system is rebooted.

Configuring TACACS+ accounting for Telnet or SSH (shell) access

To send an Accounting Start packet to the TACACS+ accounting server when an authenticated user establishes a Telnet or SSH session on the Extreme device, and an Accounting Stop packet when the user logs out.

```
device(config)# aaa accounting exec default start-stop tacacs+
```

Syntax: [no] aaa accounting exec default start-stop radius | tacacs+ | none

Configuring TACACS+ accounting for CLI commands

You can configure TACACS+ accounting for CLI commands by specifying a privilege level whose commands require accounting. For example, to configure the Extreme device to perform TACACS+ accounting for the commands available at the Super User privilege level (that is; all commands on the device), enter the following command.

```
device(config)# aaa accounting commands 0 default start-stop tacacs+
```

An Accounting Start packet is sent to the TACACS+ accounting server when a user enters a command, and an Accounting Stop packet is sent when the service provided by the command is completed.

NOTE
If authorization is enabled, and the command requires authorization, then authorization is performed before accounting takes place. If authorization fails for the command, no accounting takes place.

Syntax: [no] aaa accounting commands privilege-level default start-stop radius | tacacs+ | none

The *privilege-level* parameter can be one of the following:

- 0 - Records commands available at the Super User level (all commands)
- 4 - Records commands available at the Port Configuration level (port-config and read-only commands)

- **5** - Records commands available at the Read Only level (read-only commands)

Configuring TACACS+ accounting for system events

You can configure TACACS+ accounting to record when system events occur on the Extreme device. System events include rebooting and when changes to the active configuration are made.

The following command causes an Accounting Start packet to be sent to the TACACS+ accounting server when a system event occurs, and a Accounting Stop packet to be sent when the system event is completed.

```
device(config)# aaa accounting system default start-stop tacacs+
```

Syntax: [no] aaa accounting system default start-stop radius | tacacs+ | none

Configuring an interface as the source for all TACACS or TACACS+ packets

You can designate the lowest-numbered IP address configured an Ethernet port, loopback interface, or virtual interface as the source IP address for all TACACS or TACACS+ packets from the Extreme device. Identifying a single source IP address for TACACS or TACACS+ packets provides the following benefits:

- If your TACACS or TACACS+ server is configured to accept packets only from specific links or IP addresses, you can use this feature to simplify configuration of the TACACS or TACACS+ server by configuring the Extreme device to always send the TACACS or TACACS+ packets from the same link or source address.
- If you specify a loopback interface as the single source for TACACS or TACACS+ packets, TACACS or TACACS+ servers can receive the packets regardless of the states of individual links. Thus, if a link to the TACACS or TACACS+ server becomes unavailable but the client or server can be reached through another link, the client or server still receives the packets, and the packets still have the source IP address of the loopback interface.

The software contains separate CLI commands for specifying the source interface for Telnet, TACACS or TACACS+, and RADIUS packets. You can configure a source interface for one or more of these types of packets.

To specify an Ethernet, loopback, or virtual interface as the source for all TACACS or TACACS+ packets from the device, use the following CLI method. The software uses the lowest-numbered IP address configured on the port or interface as the source IP address for TACACS or TACACS+ packets originated by the device.

To specify the lowest-numbered IP address configured on a virtual interface as the device source for all TACACS or TACACS+ packets, enter commands such as the following.

```
device(config)# int ve 1
device(config-vif-1)# ip address 10.0.0.3/24
device(config-vif-1)# exit
device(config)# ip tacacs source-interface ve 1
```

The commands in this example configure virtual interface 1, assign IP address 10.0.0.3/24 to the interface, then designate the interface as the source for all TACACS or TACACS+ packets from the Extreme device.

Syntax: [no] ip tacacs source-interface ethernet portnum | loopback num | ve num

The *num* parameter is a loopback interface or virtual interface number. If you specify an Ethernet, the *portnum* is the port's number (including the slot number, if you are configuring a device).

Displaying TACACS or TACACS+ statistics and configuration information

The **show aaa** command displays information about all TACACS+ and RADIUS servers identified on the device.

```
device# show aaa
TACACS default key: ...
TACACS retries: 3
TACACS timeout: 3 seconds
TACACS+ Server: IP=10.20.80.20 Port=49 Usage=any Key=...
               opens=0 closes=0 timeouts=0 errors=0
               packets in=0 packets out=0
Radius default key: ...
Radius retries: 3
Radius timeout: 3 seconds
Radius Server:  IP=10.20.99.134 Auth Port=1812 Acct Port=1813 Usage=any
               Key=...
               opens=7 closes=7 timeouts=24 errors=0
               packets in=7 packets out=79
               Health-check=disabled dead-time-interval=45  auto-authenticate-time-interval=30 available
Radius Server:  IP=10.20.99.135 Auth Port=1812 Acct Port=1813 Usage=any
               Key=...
               opens=72 closes=72 timeouts=0 errors=0
               packets in=72 packets out=72
               Health-check=disabled dead-time-interval=45  auto-authenticate-time-interval=30 available
device#
```

Syntax: **show aaa**

The following table describes the TACACS or TACACS+ information displayed by the **show aaa** command.

Field	Description
Tacacs+ key	The setting configured with the **tacacs-server key** command. At the Super User privilege level, the actual text of the key is displayed. At the other privilege levels, a string of periods (....) is displayed instead of the text.
Tacacs+ retries	The setting configured with the **tacacs-server retransmit** command.
Tacacs+ timeout	The setting configured with the **tacacs-server timeout** command.
Tacacs+ dead-time	The setting configured with the **tacacs-server dead-time** command.
Tacacs+ Server	For each TACACS or TACACS+ server, the IP address, port, and the following statistics are displayed: **opens** - Number of times the port was opened for communication with the server **closes** - Number of times the port was closed normally **timeouts** - Number of times port was closed due to a timeout **errors** - Number of times an error occurred while opening the port **packets in** - Number of packets received from the server **packets out** - Number of packets sent to the server
connection	The current connection status. This can be "no connection" or "connection active".

The **show web** command displays the privilege level of Web Management Interface users.

```
device#show web
User                        Privilege      IP address
set                            0           192.168.1.234
```

Syntax: **show web**

Validating TACACS+ reply packets

The TACACS+ reply packets are validated for individual fields in the packet header and encrypted or unencrypted packet body to avoid any system failure due to processing invalid or corrupt reply packets. Since the packet body formats are different for authentication, authorization and accounting replies, packet body validation is done separately for each of these replies.

Validating TACACS+ packet header

The TACACS+ packet header validates:

- Minimum length of data (fixed size is 12 bytes) for a valid TACACS+ packet header before reading through individual fields in the header.
- Field type in the received packet header against type of TACACS+ reply from the server.
- Comparison between received packet length and full packet length (header-size + length field in the packet header).

Following table lists all possible error conditions and corresponding messages for the reply packet header validation.

Error warning message	Error condition
Warning: Received invalid TACACS+ packet header	The received packet size is less than minimum length for TACACS+ reply header
Warning: Received invalid TACACS+ packet type	Received packet having invalid or null packet type
Warning: Received invalid TACACS+ packet data	The received packet size is not matching data length specified in the packet header

```
Validating TACACS+ authentication reply
```

The TACACS+ authentication reply packet validates:

- Minimum length of data (fixed size is 6 bytes) for a valid TACACS+ authentication reply before reading through individual fields in the reply body.
- Reply packet is decrypted correctly, validate the status field received in the reply packet to be one of the legal values for TACACS+ authentication status.
- If **server-msg length** field is present in the reply packet, ensure server message is within the received packet and has non-null string message.
- If **data length** field is present in the reply packet, ensure data is within the received packet.
- Full packet length (header size + length field received in packet header) against number of bytes parsed successfully from the received reply packet.

Following table lists all possible error conditions and corresponding messages for the authentication reply validation.

Error warning message	Error condition
Warning: Invalid TACACS+ authentication reply packet	Received packet body size is less than minimum length for TACACS+ authentication reply body
Warning: Invalid TACACS+ authentication reply packet body	Received packet having invalid or null packet body
Warning: Invalid TACACS+ authentication reply packet body.. check key value	Invalid status field in the packet body. possibly key mismatch
Warning: Invalid server msg length in TACACS+ authentication reply	The server message length specified is not within packet boundary
Warning: Invalid server msg in TACACS+ authentication reply	Invalid or null data found in server message
Warning: Invalid data length in TACACS+ authentication reply	The data length specified is not within packet boundary

Error warning message	Error condition
Warning: Invalid TACACS+ authentication reply. packet total length mismatch	The total number of bytes parsed successfully from the received packet is not matching with data length specified in the packet

```
Validating TACACS+ authorization reply
```

The TACACS+ authorization reply packet validates:

- Minimum length of data (fixed size 6 bytes) for a valid TACACS+ authorization reply before reading through individual fields in the reply body.
- The reply packet is decrypted correctly, validate the status field received in the reply packet to be one of the legal values for TACACS+ authorization status.
- If **arg-count** field is present in the reply packet, ensure this is within the received packet and has non-null data.
- If **server-msg length** field is present in the reply packet, ensure server message is within the received packet and has non-null string message.
- If **data length** field is present in the reply packet, ensure data is within the received packet.
- If **arg-count** field is present in the reply packet, ensure this is within the received packet and has non-null data.
- Full packet length (header size + length field received in packet header) against number of bytes parsed successfully from the received reply packet.

Following table lists all possible error conditions and corresponding messages for the authorization reply validation.

Error warning message	Error condition
Warning: Invalid TACACS+ authorization reply packet	Received packet body size is less than minimum length for TACACS+ authorization reply body
Warning: Invalid TACACS+ authorization reply packet body	Received packet having invalid or null packet body
Warning: Invalid TACACS+ authorization reply packet body. check key value	Invalid status field in the packet body. possibly key mismatch
Warning: Invalid arg_cnt in TACACS+ authorization reply	The server argument count specified is not within packet boundary
Warning: Invalid arg_len in TACACS+ authorization reply	Invalid or null data found in argument length field
Warning: Invalid server msg length in TACACS+ authorization reply	The server message length specified is not within packet boundary
Warning: Invalid server msg in TACACS+ authorization reply	Invalid or null data found in server message
Warning: Invalid data length in TACACS+ authorization reply	The data length specified is not within packet boundary
Warning: Invalid arg length in TACACS+ authorization reply	The argument length specified is not within packet boundary
Warning: Invalid arg in TACACS+ authorization reply	Invalid or null data found in argument field
Warning: Invalid TACACS+ authorization reply. packet total length mismatch	The total number of bytes parsed successfully from the received packet is not matching with data length specified in the packet

```
Validating TACACS+ accounting reply
```

The TACACS+ accounting reply packet validates:

- Minimum length of data (fixed size 5 bytes) for a valid TACACS+ accounting reply before reading through individual fields in the reply body.
- Reply packet is decrypted correctly, validate the status field received in the reply packet to be one of the legal value for TACACS + accounting status.
- If **server-msg length** field is present in the reply packet, ensure server message is within the received packet and has non-null string message.
- If **data length** field is present in the reply packet, ensure data is within the received packet.

- Full packet length (header size + length field received in packet header) against number of bytes parsed successfully from the received reply packet.

Following table lists all possible error conditions and corresponding messages for the accounting reply validation.

Error warning message	Error condition
Warning: Invalid TACACS+ accounting reply packet	Received packet body size is less than minimum length for TACACS+ accounting reply body
Warning: Invalid TACACS+ accounting reply packet body	Received packet having invalid or null packet body
Warning: Invalid TACACS+ accounting reply packet body. check key value	Invalid status field in the packet body. possibly key mismatch
Warning: Invalid server msg length in TACACS+ accounting reply	The server message length specified is not within packet boundary
Warning: Invalid server msg in TACACS+ accounting reply	Invalid or null data found in server message
Warning: Invalid data length in TACACS+ accounting reply	The data length specified is not within packet boundary
Warning: Invalid TACACS+ accounting reply. packet total length mismatch	The total number of bytes parsed successfully from the received packet is not matching with data length specified in the packet

ACLs

ACL overview

An access control list (ACL) is a container for rules that permit or deny network traffic based on criteria that you specify.

When a frame or packet is received or sent, the device compares its header fields against the rules in applied ACLs. This comparison is done according to a rule sequence, which you can specify. Based on the comparison, the device either forwards or drops the frame or packet.

> **NOTE**
> In some ACL topics, the terms *entry*, *clause*, *statement*, or *filter* are used interchangeably with *rule*.

The benefits of ACLs include the following:

- Provide security and traffic management.
- Monitor network and user traffic.
- Save network resources by classifying traffic.
- Protect against denial of service (DOS) attacks.
- Reduce debug output.

There are five ACL types, supporting a range of filtering options:

- **MAC ACLs**—Filter traffic by Layer 2 frame contents: source and destination MAC addresses, priority, VLAN and etype.
- **IPv4 standard ACLs**—Filter IPv4 traffic by source address.
- **IPv4 extended ACLs**—Filter IPv4 traffic by source and destination address, IP protocol, QoS and other parameters. For example:
  - TCP or UDP source and destination port names and numbers
  - IP protocol names and numbers
  - TCP flags, ICMP messages and QoS attributes
- **IPv6 ACLs**—Filter routed IPv6 traffic by source and destination address, QoS, extension, and other options.
- **USER-DEFINED ACLs (UDAs)**—Use offsets to filter up to 128 bytes of specified traffic.

The following table indicates which ACL types support only named ACLs, or also ACLs identified by an integer only:

TABLE 3 Numbered and named ACLs

ACL type	Numbered / Named
MAC	Numbered or named
IPv4 standard	Numbered or named
IPv4 extended	Numbered or named
IPv6	Named only

TABLE 3 Numbered and named ACLs (continued)

ACL type	Numbered / Named
User-defined (UDA)	Numbered or named

The following guidelines apply to all named ACLs:

- ACL names can be 1 through 255 characters long, and must begin with a–z, A–Z or 0–9.
- You can use underscore (_) or hyphen (-) in ACL names, but not as the first character.
- ACL names must contain at least one alphabetic character.
- Although you can use the same name for different ACL types, Extreme recommends that you specify unique names across all ACL types.

ACL and rule limits

There are limits to the number of ACLs and rules supported. Some of these limits are configurable.

For each protocol or ACL type, the following table lists the maximum numbers of ACLs that you can define. For numbered ACLs, the table also lists the numbers you can assign to each ACL type.

TABLE 4 ACL maximums

Protocol or ACL type	Standard numbered	Extended numbered	Standard named	Extended named	Maximum total ACLs
MAC	1000 (400–1399)	NA	500 (CES/CER) 4000 (MLX/XMR)	NA	1500 (CES/CER) 5000 (MLX/XMR)
IPv4	99 (1–99)	100 (100–199)	100	500 (CES/CER) 4000 (MLX/XMR)	799 (CES/CER) 4299 (MLX/XMR)
IPv6	NA	NA			1000 (CES/CER) 4000 (MLX/XMR)
User-defined (UDAs) (MLX/XMR)	1000 (2000–2999)	NA	500 (4000–4499)	NA	1500

The following table lists the maximum total number of IPv4 and IPv6 ACL (Layer 3) rules supported:

TABLE 5 Maximum Layer 3 ACL rules

ACL type	Maximum ACL rules (in all ACLs)	How configured
IPv4 and IPv6	Lowest supported value: 1024 Default value: 4096 Maximum value: 102,400	**system-max ip-filter-sys** Refer to Modifying the Layer 3 ACL rule maximum on page 120.

For Layer 2 ACLs and for User-defined ACLs (UDAs), there are configurable parameters that specify the maximum number of rules that you can include in each ACL. There are no global parameters for these ACL types:

TABLE 6 Maximum rules per ACL (Layer 2 and UDA)

ACL type	Standard numbered	Extended numbered	Standard named	Extended named
Layer 2	64 (lowest and default max) 256 (max max)	NA	64 (lowest and default max) 256 (max max)	NA
User-defined ACLs (UDAs)	64 (lowest and default max) 256 (max max)	NA	64 (lowest and default max) 256 (max max)	NA

For details of how to modify the rule maximums for Layer 2 ACLs and for UDAs, refer to the following topics:

Layer 2 ACLs

Layer-2 Access Control Lists (ACLs) filter incoming traffic based on Layer-2 MAC header fields in the Ethernet IEEE 802.3 frame. Specifically, Layer-2 ACLs filter incoming traffic based on any of the following Layer-2 fields in the MAC header:

- Source CER 2000 Series MAC address and source MAC mask
- Destination MAC address and destination MAC mask
- VLAN ID
- Ethernet type
- 802.1p

Layer-2 ACLs filter traffic at line-rate speed.

Layer 2 ACL configuration guidelines

General Layer 2 ACL configuration guidelines

The following guidelines apply to all Layer 2 ACLs:

- On XMR Series and MLX Series devices, you cannot bind Layer 2 ACLs and IPv4 or IPv6 ACLs to the same port in the identical ingress/egress direction. However, you can perform the following configuration:
  - On one port, bind a Layer 2 ACL.
  - On a second port, bind one or both of the following:
    › An IPv4 ACL (standard or extended)
    › An IPv6 ACL
- CES 2000 Series and CER 2000 Series devices enable you to bind a Layer 2 ACL, an IPv4 ACL, and an IPv6 ACL on the same port.
- You cannot bind a Layer 2 ACL to a virtual interface.
- The Layer 2 ACL feature cannot perform SNAP and LLC encapsulation type comparisons.
- Extreme devices process ACLs in hardware.

- For all NetIron devices, if a port has an IPv4 or IPv6 ACL applied, you must remove the ACL bindings before adding that port to a VLAN that has a VE interface.
- You cannot edit or modify an existing Layer 2 ACL clause. If you want to change the clause, you must delete it first, then re-enter the new clause.
- You cannot add remarks to a Layer 2 ACL clause.
- When you bind a Layer 2 ACL that is not defined, it implicitly denies all traffic.
- The behavior of Layer 2 ACLs for dynamic LAG creation and deletion is that before a LAG is formed all ports which will be parts of the LAG must have the same configuration. For example, all of the ports can have no ACL, or have ACL 401 on inbound and outbound ports. After the LAG is removed, all ACL bindings (if there are any) are propagated to all of the secondary ports.
- Layer 2 inbound ACLs and Layer 2 inbound ACL-based rate limiting are not supported on Layer 3 VPNs.
- You can bind multiple rate limiting policies to a single port. However, once a matching ACL clause is found for a packet, the device does not evaluate subsequent clauses in that rate limiting ACL and subsequent rate limiting ACLs.
- If you need to downgrade to a version earlier than 5.6, refer to Upgrade and downgrade considerations (5.6.00) on page 197.

Configuration considerations for dual inbound ACLs on CES 2000 Series and CER 2000 Series devices

CES 2000 Series and CER 2000 Series devices enable you to bind a Layer 2 ACL, an IPv4 ACL, and an IPv6 ACL to the same port, as follows:

- A Layer 2 ACL is bound to the port.
- An IPv4 and an IPv6 ACL, or one of the two, are bound to the same port.

The filtering algorithm is as follows:

1. An incoming packet is first examined by the IPv4 ACL or the IPv6 ACL, depending on the protocol.
2. If the packet is denied by the IPv4/IPv6 ACL, the packet is dropped, without being examined by the L2 ACL.
3. If the packet is permitted by the IPv4/IPv6 ACL, it is then examined by the L2 ACL.

However, there is an implicit "deny" at the end of any ACL. To enable the above algorithm, IPv4/IPv6 ACLs intended for this scenario must include a "permit any" filter as the last rule. Such a rule ensures that even packets not explicitly permitted by the IPv4/IPv6 ACL are passed to the L2 ACL.

Dual inbound ACLs can also affect the behavior of ACL accounting. For details, refer to ACL accounting on CES 2000 Series and CER 2000 Series devices on page 94.

Configuration considerations for VPLS, VLL, and VLL-Local endpoints

L2 ACLs are supported on VPLS, VLL, and VLL-local endpoints with the following configuration considerations:

- First configure the port as a VPLS, VLL, or VLL-local endpoint and then bind the Layer-2 ACL on it.
- First remove the Layer-2 ACL from a VPLS, VLL, or VLL-local endpoint before removing the port from the VPLS, VLL, or VLL-local instance or corresponding VLAN.
- First remove the Layer-2 ACL from a VPLS, VLL, or VLL-local endpoint(s) before deleting the VPLS, VLL, or VLL-local instance or corresponding VLAN.
- If the VPLS, VLL, or VLL-local endpoint is a LAG port, you must first remove the Layer-2 ACL from the primary LAG port before deleting the LAG. This restriction is applicable even if you are deleting the LAG using the **force** keyword.

- If a VLL or VLL-local endpoint is a LAG port with Layer-2 ACL, you have to first remove the Layer-2 ACL from the primary LAG port before dynamically removing a port from the LAG.
- Ensure that no VPLS, VLL, or VLL-local endpoint exists with an Layer-2 ACL before entering the command: **no router mpls** .

Types of Layer-2 ACLs

Layer-2 ACLs can be numbered or named. Numbered Layer-2 ACL table IDs range from 400 through 1399 for a maximum of 1000 configurable numbered Layer-2 ACL tables.

Within each Layer-2 ACL table, you can configure, by default, up to 64 rules. (To change the default, refer to Increasing the maximum number of clauses per Layer-2 ACL table on page 89.) Each rule can define a set of Layer-2 parameters for filtering. Once you completely define a Layer-2 ACL table, you must bind it to the interface for filtering to take effect.

The maximum number of Layer 2 ACLs varies with the platform. For details, refer to ACL and rule limits on page 82. The maximum length of a named Layer-2 ACL is 255 characters. The Layer-2 ACL name cannot begin with digits 0 through 9, which prevents confusion with the numbered L2 ACLs.

The device evaluates traffic coming into the port against each ACL clause. Once a matching entry is found, the device either forwards or drops the traffic, depending upon the action specified for the clause. Once a matching entry is found, the device does not evaluate the traffic against subsequent clauses.

By default, if the traffic does not match any of the clauses in the ACL table, the device drops the traffic. To override this behavior, specify a "permit any ..." clause at the end of the table to match and forward all traffic not matched by the previous clauses.

> **NOTE**
> Use precaution when placing entries within ACL tables. You can also check conflict and duplication among ACL entries. Refer to:
> - Enabling ACL duplication check on page 125
> - Enabling ACL conflict check on page 125

Creating a numbered Layer-2 ACL table

You create a numbered Layer-2 ACL table by defining a Layer-2 ACL clause.

To create a numbered Layer-2 ACL table, enter commands (clauses) such as the following at the Global CONFIG level of the CLI. Note that you can add additional clauses to the ACL table at any time by entering the command with the same table ID and different MAC parameters.

```
device(config)# access-list 400 deny any any any etype arp
device(config)# access-list 400 deny any any any etype ipv6
device(config)# access-list 400 deny any any any etype 8848
device(config)# access-list 400 permit any any 100
```

The above configuration creates a Layer-2 ACL with an ID of 400. When applied to an interface, this Layer-2 ACL table will deny all ARP, IPv6, and MPLS multicast traffic; and permit all other traffic in VLAN 100.

```
device(config)# access-list 1399 permit any any 100 etype any dscp-marking 54
Warning: this ACL will have unexpected results on non-IP packets. Make sure the traffic on the interfaces
are IP packets.
```

The above configuration creates a Layer-2 ACL with an ID of 1399 and matches VLAN 100 and mark DSCP to 54.

> **NOTE**
> A warning message is displayed, if the incoming packet is a non-IP packet and without an L3 header. The warning message is displayed in the above configuration example.

For more examples of valid Layer-2 ACL clauses, see Filtering and priority manipulation based on 802.1p priority on page 88.

The ACL functionality for filtering traffic is enhanced with sequence numbers that enable users to insert, modify or delete rules at any position, without having to remove and reapply the entire ACL. A sequence number is assigned to each ACL entry and ACL rules are applied in the order of lowest to highest sequence number. Therefore, you can insert a new filter rule at any position you want in the ACL table by specifying the sequence number. If you do not specify a sequence number a default sequence number is applied to each ACL entry. The default value is 10+ the sequence number of the last ACL entry provisioned in the ACL table. Therefore, when you do not specify a sequence number, the rule is added to the end of the ACL table. The default value for the first entry in a Layer 2 ACL table is "10".

The following example creates a numbered Layer-2 ACL table "401" with two ACL entries.

```
device(config)# access-list 401 permit 0000.1111.1111 ffff.ffff.ffff any any etype any
device(config)# access-list 401 sequence 23 permit 0000.1111.1121 ffff.ffff.ffff any 23 etype any
```

The first entry in this example does not specify an ACL entry sequence number. Therefore the system assigns the default sequence number "10". In the second entry, the sequence number is specified as "23". The output from the **show access-list** command for the ACL table is:

```
device(config)# show access-list 401
L2 MAC Access List 401:
10: permit 0000.1111.1111 ffff.ffff.ffff any any etype any
23: sequence 23 permit 0000.1111.1121 ffff.ffff.ffff any 23 etype any
```

The **show access-list** command only displays user-configured sequence numbers. In this example, "sequence 23" is shown for the second ACL entry because this is a user-specified sequence number. ACL entry sequence numbers that are generated by the system are not displayed.

> **NOTE**
> If you specify a sequence number that is already used by another ACL filter rule, the following error message is displayed."Error: Entry with sequence 23 already exists!"

> **NOTE**
> If you specify a sequence number which is greater than the limit (214748364) the following error message is displayed."Error: Valid range for sequence is 1 to 214748364"

Re-sequencing a numbered Layer-2 ACL table

To allow new ACL entries to be inserted between ACL entries that have consecutive sequence numbers, you can create space between sequence numbers of adjacent filters by regenerating the ACL table.

To re-sequence ACL table "407", use the following command.

```
device(config)# access-list 407 regenerate-seq-num
```

This command regenerates the filter sequence numbers in steps of 10, assigning the default sequence number "10" to the first entry in the table.

> **NOTE**
> If sequence numbers generated by the **regenerate-seq-num** command cross the limit (214748364), then re-sequencing of ACL filters will not take place and the following error message is displayed."Error: Valid range for sequence is 1 to 214748364".

Deleting a numbered Layer-2 ACL entry

You can delete an ACL filter rule by providing the sequence number or without providing the sequence number. To delete an ACL filter rule without providing a sequence number you must specify the filter rule attributes. To delete an ACL filter rule providing a sequence number you can provide the sequence number alone or the sequence number and the other filter rule attributes.

To delete a filter rule with the sequence number "23" from Layer-2 access list "401" by specifying the sequence number alone, enter the following command.

```
device(config)# no access-list 401 sequence 23
```

You can also delete this entry by specifying both the entry sequence number and filter rule attributes. For example:

```
device(config)# no access-list 401 sequence 23 permit 0000.1111.1121 ffff.ffff.ffff any 23 etype any
```

Alternatively, you can delete this rule by providing the filter rule attributes only. For example:

```
device(config)# no access-list 401 permit 0000.1111.1121 ffff.ffff.ffff any 23 etype any
```

> **NOTE**
> If you try to delete an ACL filter rule using the sequence number, but the sequence number that you specify does not exist, the following error message will be displayed."Error: Entry with sequence 20 does not exist!"

Parameters to configure numbered Layer-2 ACL statements

Syntax: [no] access-list *num* [sequence *num*] permit | deny { *src-mac mask* | any } { *dest-mac mask* | any } [{ *vlan-id* | any }] [0180.c200.0002 ffff.ffff.ffff] [etype { *etype-str* | *etype-hex* }] [priority *802.1p-value* | priority-force *802.1p-value* | priority-mapping *802.1p-value* | mark-flow-id | dscp-marking *number*]

Syntax: access-list *num* regenerate-seq-num [*num*]

The *num* parameter specifies the Layer-2 ACL table that the clause belongs to. The table ID can range from 400 to 1399. You can define a total of 1000 Layer-2 ACL tables.

> **NOTE**
> If users configure the maximum L2 ACL of 1399, the other ACL types, such as IP and IPv6 ACL, will have limited space. It may affect memory usage in CES 2000 Seriesor CER 2000 Series and MLX Series or XMR Series devices.

The **sequence** parameter specifies where the conditional statement is to be added in the access list. You can add a conditional statement at particular place in an access list by specifying the entry number using the **sequence** keyword. The range is from 1 through 214748364. If the **sequence***num* is specified for the first ACL clause "clause-1" and not specified for the second ACL clause "clause-2", then the system rounds off the sequence number of the last ACL filter rule provisioned to the next 10th digit. The default value for the first clause in a Layer-2 ACL table is "10".

The **permit | deny** argument determines the action to be taken when a match occurs.

The *src-macmask* | **any** parameter specifies the source MAC address. You can enter a specific address and a comparison mask or the keyword any to filter on all MAC addresses. Specify the mask using Fs and zeros. For example, to match on the first two bytes of the address aabb.ccdd.eeff, use the mask ffff.0000.0000. In this case, the clause matches all source MAC addresses that contain "aabb" as the first two bytes and any values in the remaining bytes of the MAC address. If you specify any, you do not need to specify a mask and the clause matches on all MAC addresses.

The *dest-macmask* | **any** parameter specifies the destination MAC address. The syntax rules are the same as those for the *src-macmask* | **any** parameter.

The optional *vlan-id* | **any** parameter specifies the vlan-id to be matched against the VLAN ID of the incoming packet. You can specify **any** to ignore the vlan-id match.

The optional **etype**{ *etype-str* | *etype-hex* } argument specifies the Ethernet type field of the incoming packet in order for a match to occur.

LACP traffic can be filtered using MAC destination address **0180.c200.0002 ffff.ffff.ffff** in a Layer 2 ACL.

The *etype-str* variable can be one of the following keywords:

- **IPv4-l5** (Etype=0x0800, IPv4, HeaderLen 20 bytes)
- **ARP** (Etype=0x0806, IP ARP)
- **IPv6** (Etype=0x86dd, IP version 6)
- **ANY** – specify etype any to ignore Ethernet type field match.

The *etype-hex* variable can be a hex number that indicates a supported Ethernet type (etype). For example:

- **8847** (MPLS unicast)
- **8848** (MPLS multicast)
- **88A8** (MAC in MAC)
- **0806** (ARP)

    NOTE
    Filtering based on etype value is only supported for Layer-2 inbound ACLs. It is not supported for Layer-2 outbound ACLs. The *etype-hex* option is supported only for Gen2 and Gen2+ cards.

Use **dscp-marking** *number* to mark the DSCP value in the incoming packet with the value you specify.

Parameters for regenerating Layer-2 ACL table sequence numbers

num	Specifies the number of the Layer-2 ACL table to resequence
regenerate-seq-num [*num* **]**	(Optional) Specifies the initial sequence number for the access list after regeneration. The valid range is from 1 through 214748364. The default value is 10. ACL filter rule sequence numbers are regenerated in steps of 10.

Filtering and priority manipulation based on 802.1p priority

Layer-2 ACL support has been provided for filtering and priority manipulation based on a packet's 802.1p priority using the following keywords.

The following priority options can be configured following the etype argument.

    NOTE
    The keywords **priority** and **priority-force** cannot be used together in an ACL entry.

The **priority** option assigns outgoing traffic that matches the ACL to a hardware forwarding queue based on the incoming 802.1p value. If the incoming packet priority is lower than the specified value, the outgoing packet priority is set to the specified value. Should the incoming packet priority have a higher priority than the specified value, the priority is not changed. This option is applicable for inbound ACLs only.

The **priority-force** option sets the outgoing priority of the matching packet to the specified value, regardless of the incoming packet priority value. This option is applicable for inbound ACLs only.

The **priority-mapping** option matches on the incoming packet's 802.1p value. This option does not change the packet's forwarding internal forwarding queue or change the outgoing 802.1p value. This keyword is applicable for both inbound and outbound ACLs.

On the CER 2000 Series and CES 2000 Series series, the **exclude-pcp-marking** option is available if you include the **priority** or **priority-force** keywords in ACL permit or deny statements. The **exclude-pcp-marking** option prevents PCP modification both for regular Layer 2/Layer 3 traffic and for VPLS packets.

The *802.1p-value* variable specifies one of the following QoS queues for use with the priority, priority-force options

- 0 - qosp0
- 1 - qosp1
- 2 - qosp2
- 3 - qosp3
- 4 - qosp4
- 5 - qosp5
- 6 - qosp6
- 7 - qosp7

Use the **no** parameter to delete the Layer-2 ACL clause from the table. When all clauses are deleted from a table, the table is automatically deleted from the system.

The following shows some examples of valid Layer-2 ACL clauses.

```
device(config)# access-list 501 permit 0025.0113.0101 ffff.ffff.ffff 0021.3113.0101 ffff.ffff.ffff any
etype any priority 2
device(config)# access-list 501 deny 0025.0113.0102 ffff.ffff.ffff 0021.3113.0101 ffff.ffff.ffff any etype
any log
device(config)# access-list 501 permit any 0021.3121.0101 ffff.ffff.ffff any etype any priority-mapping 1
device(config)# access-list 501 deny 0025.0122.010a ffff.ffff.ffff any any etype arp log
device(config)# access-list 501 permit 0025.0123.010a ffff.ffff.ffff 0021.3113.0101 ffff.ffff.ffff any
etype ipv4-l5 mirror
device(config)# access-list 501 permit 0025.0124.010a ffff.ffff.ffff 0021.3113.0101 ffff.ffff.ffff any
etype ipv6 mirror priority-force 5
device(config)# access-list 501 permit 0025.0124.010c ffff.ffff.ffff 0021.3113.0101 ffff.ffff.ffff any
etype any
device(config)# access-list 501 deny any any 1618 etype any priority-mapping 0
device(config)# access-list 501 deny any any 1615 etype any priority-force 5
device(config)# access-list 501 deny any any 1613 etype any priority 3
device(config)# access-list 401 sequence 23 permit 0000.1111.1121 ffff.ffff.ffff any 23 etype any
```

Inserting and deleting Layer-2 ACL clauses

You can make changes to the Layer-2 ACL table definitions without unbinding and rebinding the table from an interface. For example, you can add a new clause to the ACL table, delete a clause from the table, delete the ACL table, etc.

When making changes to a Layer-2 ACL that has already been bound to an interface, the ACL is automatically unbound and rebound as CAM entries for the ACL and applied interfaces are removed and recreated. This can create a brief period of time lasting several milliseconds during which the ACL is effectively unbound from all applied interfaces. The duration of this period can change, depending on the number of lines in the ACL.

Increasing the maximum number of clauses per Layer-2 ACL table

You can increase the maximum number of clauses configurable within a Layer-2 (L2) ACL table.

To increase the maximum number of clauses per L2 ACL table, enter a command such as the following at the Global CONFIG level of the CLI. The system supports 64 (default) to 256 ACL table entries per L2 ACL and a system reload is required after changing this value.

```
device(config)# system-max l2-acl-table-entries 200
```

Syntax: [no] **system-max l2-acl-table-entries** *max*

> NOTE
> The **l2-acl-table-entries** controls the maximum number of filters supported on one Layer-2 ACL. The named Layer-2 ACL is also subject to the configuration of this **system-max** value.

The *max* parameter specifies the maximum number of clauses per Layer-2 ACL.

Binding a numbered Layer-2 ACL table to an interface

To enable Layer-2 ACL filtering, bind the Layer-2 ACL table to an interface. Enter a command such as the following at the Interface level of the CLI to bind an inbound Layer-2 ACL.

```
device(config)# int e 4/12
device(config-int-e100-4/12)# mac access-group 400 in
```

Enter a command such as the following at the Interface level of the CLI to bind an outbound Layer-2 ACL.

```
device(config)# int e 4/12
device(config-int-e100-4/12)# mac access-group 400 out
```

Syntax: [no] **mac access-group** *num* **in** | **out**

Filtering by MAC address

In the following example, an ACL is created that denies all traffic from the host with the MAC address 0000.0056.7890 being sent to the host with the MAC address 0000.0033.4455.

```
device(config)# access-list 401 deny 0012.3456.7890 ffff.ffff.ffff 0000.0033.4455 ffff.ffff.ffff
device(config)# access-list 401 permit any any
```

Using the mask, you can make the access list apply to a range of addresses. For instance if you changed the mask in the previous example from 0012.3456.7890 to ffff.ffff.fff0, all hosts with addresses from 0000.0056.7890 to 0000.0056.789f would be blocked. This configuration for this example is shown in the following.

```
device(config)# access-list 401 deny 0000.0056.7890 ffff.ffff.fff0 0000.0033.4455 ffff.ffff.ffff
device(config)# access-list 401 permit any any
```

The *num* parameter specifies the Layer-2 ACL table ID to bind to the interface.

Filtering broadcast traffic

To define an Layer-2 ACL that filters Broadcast traffic, enter commands such as the following.

```
device(config)#access-list 401 deny any ffff.ffff.ffff ffff.ffff.ffff
device(config)#access-list 401 permit any any any
```

To bind an Layer-2 ACL that filters Broadcast traffic, enter commands such as the following.

```
device(config)#int eth 14/1
device(config-if-e10000-14/1)#mac access-gr 401 in
```

Using the priority option

In the following example, Access-list 401 assigns ARP packets with any source and destination addresses from VLAN 10 to internal priority queue 5. Access-list 401 then maps the ARP packets to the 802.1p value 5 when outbound on an 802.1q interface and when an 802.1p priority is lower than 5. Incoming packets with an 802.1p priority value greater than 5 are unchanged.

```
device(config)# access-list 401 permit any any 10 etype arp priority 5
```

On the CER 2000 Series and CES 2000 Series series, the **exclude-pcp-marking** option is available if you include the **priority** or **priority-force** keywords in ACL permit or deny statements. The **exclude-pcp-marking** option prevents PCP modification both for regular Layer 2/Layer 3 traffic and for VPLS packets.

Using the priority force option

In the following example, access list 401 assigns IPv4 packets with any source and destination addresses from VLAN 10 to the internal priority queue 6 and changes the outgoing 802.1p value to 6.

```
device(config)# access-list 401 permit any any 10 etype ipv4-l5 priority-force 6
```

Using the priority mapping option

In the following example, access list 401 permits IPv6 packets with any source and destination addresses from VLAN 10 that have an 802.1p priority of 3. The outgoing packet is not modified.

```
device(config)# access-list 401 permit any any 10 etype ipv6 priority-mapping 3
```

Using the drop-precedence keyword option

In the following example, access list 410 assigns IPv4 packets with any source and destination addresses from VLAN 10 to drop-precedence 0.

```
device(config)# access-list 410 permit any any 10 etype ipv4-l5 drop-precedence 0
```

The CES 2000 Series and CER 2000 Series devices treat the **drop-precedence** (DP) value internally, and do not mark any packets on DP explicitly.

For example the following ACL is accepted but will not change the DP value of any packet going through CES 2000 Series and CER 2000 Series devices:

```
access-list 1300 permit any any 11 etype any drop-precedence 2
access-list 1300 permit any any 12 etype any drop-precedence-force 1
```

The above configuration CLI specifies DP from 0 to 3, but CES 2000 Series and CER 2000 Series devices map them to 0 to 2 as follows:

Configuration CLI CES Internal Process

0 0

1 1

2 1

3 2

Using the drop-precedence-force keyword option

In the following example, access list 411 assigns packets with any source and destination addresses from VLAN 11 to drop-precedence 1.

```
device(config)# access-list 411 perm an an 11 etype an drop-precedence-force 1
```

Using the mirror keyword option

In the following example, access list 413 permits IPv6 packets with any source and destination addresses from VLAN 10 having an 802.1p priority of 3 and sends a copy of the matching packet to the specified mirror port.

```
device(config)# access-list 413 permit any any 10 etype ipv6 priority-mapping 3
```

Using the mark flow ID keyword option

> NOTE
> The mark-flow-id keyword option is available for CES 2000 Series and CER 2000 Series devices only.

The mark-flow-id option balances traffic coming from a LAG port and going to another LAG port. By applying the **mark-flow-id** command to the inbound LAG port of an ACL, the matching traffic is marked with a flow ID and will be distributed over different physical ports on the outbound LAG interface.

In the following example, access list 414 permits IPv6 packets with any source and destination addresses from VLAN 10 having an 802.1p priority of 2 and marks the flow ID for load-balancing on LAG ports.

```
device(config)#access-list 414 permit 1425.0124.010c ffff.ffff.ffff any 14 etype ipv4-l5 priority-mapping 2
mark-flow-id
```

In the following example, access list 414 permits IPv4 packets from source mac 1425.0124.010c and any destination addresses from VLAN 14 having an 802.1p priority of 2 and marks the flow ID for load-balancing on LAG ports.

```
device(config)#access-list 414 permit 1425.0124.010c ffff.ffff.ffff any 14 etype ipv4-l5 priority-mapping 2
mark-flow-id
```

Creating a named Layer-2 ACL table

To create for example a named Layer-2 ACL called example_l2_acl, enter the following commands.

```
device(config)#mac access-list example_l2_acl
device(config-mac-nacl)#deny 0000.0000.0001 ffff.ffff.ffff any
device(config-mac-nacl)#permit any 0000.0000.0002 ffff.ffff.ffff
device(config-mac-nacl)#exit
```

Following is an example of how a named Layer-2 ACL "example_l2_acl" is displayed in the configuration file.

```
!
mac access-list example_l2_acl
 deny 0000.0000.0001 ffff.ffff.ffff any
 permit any 0000.0000.0002 ffff.ffff.ffff
!
```

The following example displays the output of the **show access-list** command for "l2_ACL".

```
L2 MAC Access List l2_acl:
21: sequence 21
 permit 0000.3333.3333 ffff.ffff.ffff any any etype any
31: deny any any any etype any log
```

In this example, the display of "**sequence 21** " for the first entry indicates that the sequence number is user-configured. In the second entry, the sequence number is not displayed; this indicates that the sequence number was not specified by the user but generated by the system.

To re-sequence a named Layer-2 ACL table, enter the following command:

```
device(config)# mac access-list l2_acl
device(config-std-nacl-l2_acl)# regenerate-seq-num
```

Syntax: **[no] mac access-list** *acl_name*

Syntax: **[no] sequence** *num* **permit** | **deny** *src-mac mask* | **any** *dest-mac mask* | **any** [*vlan-id* | **any**] [**0180.c200.0002 ffff.ffff.ffff**] [**etype** { *etype-str* | *etype-hex* }] [**priority** *802.1p-value* | **priority-force** *802.1p-value* | **priority-mapping** *802.1p-value* | **mark-flow-id** | **dscp-marking** *number*]

On the CER 2000 Series and CES 2000 Series series, the **exclude-pcp-marking** option is available if you include the **priority** or **priority-force** keywords in ACL permit or deny statements. The **exclude-pcp-marking** option prevents PCP modification both for regular Layer 2/Layer 3 traffic and for VPLS packets.

Syntax: **regenerate-seq-num** [*num*]

> NOTE
> Filtering based on etype value is only supported for Layer-2 inbound ACLs. It is not supported for Layer-2 outbound ACLs. The *etype-hex* option is supported only for Gen2 and Gen2+ cards.

Binding a named Layer-2 ACL table to an interface

Following is an example of the named Layer-2 ACL "example_l2_acl" applied to the inbound of port 2/2.

```
device(config)# interface e 2/2
device(config-if-e1000-2/2)#mac access-group example_l2_acl in
```

Syntax: [**no**] **mac access-group** *acl_name in* | *out*

ACL accounting

Multi-Service devices may be configured to monitor the number of times an ACL is used to filter incoming or outgoing traffic on an interface. The **show access-list accounting** command displays the number of "hits" or how many times ACL filters permitted or denied packets that matched the conditions of the filters. For more detailed information about ACL accounting, please refer to "ACL accounting".

Enabling and disabling ACL accounting on XMR Series and MLX Series devices

ACL accounting is disabled by default on XMR Series and MLX Series devices. To enable ACL accounting, enter the following command:

```
device(config)# acl-policy
device(config-acl-policy)# enable-acl-counter
```

Syntax: **[no] enable-acl-counter**

> NOTE
> Enabling or disabling ACL accounting affects the gathering of statistics from all ACL types (Layer-2, IPv4 and IPv6).

When ACL accounting is enabled, use the **accounting-no-sort** command to present the access-list entries in the configured order when displaying ACL accounting data.

```
device(config)# acl-policy
device(config-acl-policy)# accounting-no-sort
```

Syntax: [no] accounting-no-sort

The **no** version of the command displays the access-list entries in sorted order based on the number of ACL hits.

High CPU utilization and ACL accounting

High CPU utilization may be seen for ACL accounting on a line card when ACL accounting is enabled along with large number (exceeds 10000) of ACL entries. Communication between the management module and the line card module may also get affected which may result in failure of some operational and configuration commands.

High CPU utilization may also impact time-sensitive protocols such as BFD, LACP (with short timer), and others resulting in erroneous traffic forwarding.

The following line card modules are affected with this issue:

- MLX 2-port 100-GbE (M) CFP2 module (BR-MLX-100GX2-CFP2-M)

- MLXe 4-port 40-GbE (M) module (BR-MLX-40Gx4-M)

- BR-MLX-10Gx4-M-IPSEC 4-port 1/10GbE (BR-MLX-10Gx4-M-IPSEC)

- MLX 20-port 10-GbE/ 1GbE (M) combo module (BR-MLX-10GX20-M)

To recover from the issue, disable ACL accounting using the following command:

```
device(config)# acl-policy
device(config-acl-policy)# no enable-acl-counter
```

ACL accounting on CES 2000 Series and CER 2000 Series devices

The following special considerations affect how ACL accounting is configured on CES 2000 Series and CER 2000 Series devices.

- On CES 2000 Series and CER 2000 Series devices you enable ACL accounting at the filter level by adding an **enable-accounting** keyword in each clause of an ACL for which you want to gather statistics.

- CAM resources are shared on CES 2000 Series and CER 2000 Series devices between ACL accounting and ACL rate-limiting. This limits the number of ACL accounting instances available on the system.

- If ACL deny logging and ACL accounting are enabled on the same ACL clause, deny logging takes precedence and ACL accounting statistics will not be available for that clause.

- You can bind both an inbound L2 ACL and an inbound IP ACL to the same port on CES 2000 Series and CER 2000 Series devices. Refer to "Configuration considerations for dual inbound ACLs on CES 2000 Series and CER 2000 Series devices" and "ACL Accounting interactions between L2 ACLs and IP ACLs" for further information.

For detailed information about ACL accounting considerations for CES 2000 Series and CER 2000 Series devices, please refer to "ACL accounting".

Configuring ACL Deny Logging for Layer-2 inbound ACLs

Configuring ACL Deny Logging for Layer-2 ACLs requires the following:

- Enabling the Log Option on a filter.

- Enabling ACL Deny Logging on an Interface

Enabling the log option on a filter

ACL Logging of Layer-2 ACLs requires that you add the **log** option to an ACL statement as shown.

```
device(config)#access-list 401 deny any any any log
```

The **log** option enables logging for the Layer-2 ACL being defined.

Enabling ACL Deny Logging on an interface

The **mac access-group enable-deny-logging** command must be configured as shown on each interface that you want ACL Deny Logging for Layer-2 ACLs to function.

```
device(config)# interface ethernet 5/1
device(config-if-e1000-5/1)# mac access-group enable-deny-logging
```

Syntax: [no] mac access-group enable-deny-logging [hw-drop]

The **hw-drop** option specifies that Layer-2 ACL Log packets be dropped in hardware. This is implemented to reduce the CPU load. In practice this means that the packet counts for denied traffic will only account for the first packet in each time cycle. The **no mac access-group enable-deny-logging hw-drop** command only removes the **hw-drop** keyword.

> NOTE
> Using this command, ACL logging can be enabled and disabled dynamically and does not require you to rebind the ACLs using the **ip rebind-acl** command

> NOTE
> When configuring the **mac access-group enable-deny-logging** command on VPLS, VLL, and VLL-Local endpoints, please refer to Configuration considerations for VPLS, VLL, and VLL-Local endpoints on page 84 for configuration guidelines.

Displaying Layer-2 ACLs

Use the **show access-list** command to display named and numbered Layer 2 (L2) ACL tables.

To display the total number of Layer-2 and IPv4 access lists and the number of filters configured for each list, use the **show access-list count** command.

```
device(config)#show access-list count
Total
 4 ACLs exist.
ACL 102, total 10 clauses
ACL 105, total 15 clauses
ACL 400, total 100 clauses
ACL 401, total 2 clauses
```

> NOTE
> Empty ACLs that are applied to interfaces are included in the total ACL count but are not displayed.

To display a L2 numbered ACL table, use the **show access-list** *num* command.

```
device(config)# show access-list 598
L2 MAC Access List 598:
10: deny 0000.0030.0313 ffff.ffff.ffff 0000.0030.0313 ffff.ffff.ffff any etype 20: any log permit any any
any etype any priority-force 4
```

To display a Layer-2 named ACL table use the **show access-list** *l2_acl_name* command.

```
device(config)# show access-list example
L2 MAC Access List example:
10: deny 0000.0030.0310 ffff.ffff.ffff 0000.0030.0110 ffff.ffff.ffff any etype ipv4-l5 log
```

```
20: deny 0000.0030.0311 ffff.ffff.ffff 0000.0030.0111 ffff.ffff.ffff any etype arp log
30: deny 0000.0030.0312 ffff.ffff.ffff 0000.0030.0112 ffff.ffff.ffff any etype ipv6 log
40: deny 0000.0030.0313 ffff.ffff.ffff 0000.0030.0313 ffff.ffff.ffff any etype any log
50: permit any any any etype any priority-force 4
```

Syntax: **show access-list** { **count** | *num* | *l2_acl _ name* }

The **count** parameter specifies displaying the total number of Layer-2 and IPv4 access lists and the number of filters configured for each list. Empty ACLs that are applied to interfaces are included in the total ACL count but are not displayed.

The *num* variable specifies the Layer-2 ACL table ID.

The *l2_acl_name* variable specifies the Layer-2 ACL name.

To display all Layer-2 named ACL tables, use the following command.

```
device(config)# show access-list l2
L2 MAC Access List example:
10: deny 0000.0030.0310 ffff.ffff.ffff 0000.0030.0010 ffff.ffff.ffff any etype 20: permit any any any etype
any
L2 MAC Access List mac-access-list-481-12345678901234567890123456789012345678901234567890:
10: permit 0025.0113.0101 ffff.ffff.ffff 0021.3113.0101 ffff.ffff.ffff any etype any
20: permit any 0021.3121.0101 ffff.ffff.ffff any etype any
30: deny 0025.0122.010a ffff.ffff.ffff any any etype arp log
40: deny any any any etype any
```

Syntax: **show access-list l2**

The **l2** parameter specifies the display of all Layer-2 ACL tables.

Displaying Layer-2 ACL statistics on XMR Series and MLX Series devices

To display Layer 2 inbound ACL statistics on XMR Series and MLX Series devices, enter commands such as the following.

```
(config-if-e10000-14/1)#show access-list acc eth 14/1 in l2
Collecting L2 ACL accounting for 400 on port 14/1  ...  Completed successfully.
L2 ACL Accounting Information:
Inbound: ACL 400
    0:   permit any any 100 etype ipv4-l5
       Hit count: (1 sec)                 0   (1 min)                  0
                  (5 min)                 0   (accum)                  0
    1:   deny any any any etype arp
       Hit count: (1 sec)                 0   (1 min)                  0
                  (5 min)                 0   (accum)                  0
```

To display Layer 2 outbound ACL statistics on XMR Series and MLX Series devices, enter commands such as the following.

```
device(config-if-e10000-14/1)#show access-list acc eth 14/1 out l2
Collecting L2 ACL accounting for 400 on port 14/1  ...  Completed successfully.
L2 ACL Accounting Information:
Outbound: ACL 400
    0:   permit any any 100 etype ipv4-l5
       Hit count: (1 sec)                 0   (1 min)                  0
                  (5 min)                 0   (accum)                  0
    1:   deny any any any etype arp
       Hit count: (1 sec)                 0   (1 min)                  0
                  (5 min)                 0   (accum)                  0
```

Syntax: **show access-list accounting** *int_type slot/port in* | *out* **l2**

To display the **show access-list** command output in the configuration format, use the **display-config-format** command.

```
device(config)# acl-policy
device(config-acl-policy)# display-config-format
```

Output example with **display-config-format** command enabled.

```
device(config)#show access-list name xGW_Filter2
ip access-list extended xGW_Filter2
 permit vlan 2405 ip host 10.33.44.55 any
 permit vlan 3000 ip any any
```

Syntax: [no] display-config-format

The**no** version of the **display-config-format** command will be present the **show access-list** command in standard form.

There is an SNMP table that supports this command. Refer to the *Unified IP MIB Reference* for more information.

Displaying Layer-2 ACL statistics on CES 2000 Series and CER 2000 Series devices

To display Layer 2 inbound ACL statistics on CES 2000 Series and CER 2000 Series devices, enter commands such as the following.

```
(config-if-e10000-14/1)#show access-list acc eth 14/1 in l2
Collecting L2 ACL accounting for 400 on port 14/1  ...  Completed successfully.
L2 ACL Accounting Information:
Inbound: ACL 400
    0:   permit enable-accounting any any 100 etype ipv4-l5
       Hit count: (1 sec)                 0   (1 min)               0
                  (5 min)                 0   (accum)               0
    1:   deny any any any etype arp
       Hit count: Accounting is not enabled
    2:   deny enable-accounting tcp any any log
       Hit count: Accounting is not available due to deny logging
```

To display Layer 2 outbound ACL statistics on CES 2000 Series and CER 2000 Series devices, enter commands such as the following.

```
(config-if-e10000-14/1)#show access-list acc eth 14/1 out l2
Collecting L2 ACL accounting for 400 on port 14/1  ...  Completed successfully.
L2 ACL Accounting Information:
Outbound: ACL 400
    0:   permit enable-accounting any any 100 etype ipv4-l5
       Hit count: (1 sec)                 0   (1 min)               0
                  (5 min)                 0   (accum)               0
    1:   deny any any any etype arp
       Hit count: Accounting is not enabled
    2:   deny enable-accounting tcp any any log
       Hit count: Accounting is not available due to deny logging
```

show access-list accounting *int_type slot/port in | out* l2

IPv4 ACLs

This section discusses the IPv4 Access Control List (ACL) feature, which enables you to filter traffic based on the information in the IP packet header.

You can use IPv4 ACLs to provide input to other features such as route maps, distribution lists, rate limiting, and BGP. When you use an ACL this way, use permit statements in the ACL to specify the traffic that you want to send to the other feature. If you use deny statements, the traffic specified by the deny statements is not supplied to the other feature. Refer to the chapters for a specific feature for information on using ACLs as input to those features.

IPv4 ACL overview and guidelines

This section describes how ACLs are processed and includes configuration guidelines.

How the device processes ACLs

The Extreme device processes traffic that ACLs filter in hardware. The Extreme device creates an entry for each ACL in the Content Addressable Memory (CAM) at startup or when the ACL is created. The Extreme device uses these CAM entries to permit or deny packets in the hardware, without sending the packets to the CPU for processing.

Default ACL action

The default action when no ACLs is applied or binded on an Extreme interface is to permit all traffic, if the ACL is applied on the interface, which is not configured, then the default action is deny all traffic that is not explicitly permitted on the port:

- If you want to tightly control access, configure ACLs consisting of permit entries for the access you want to permit. The ACLs implicitly deny all other access.
- If you want to secure access in environments with many users, you might want to configure ACLs that consist of explicit deny entries, then add an entry to permit all access to the end of each ACL. The software permits packets that are not denied by the deny entries.
- If dual inbound ACLs (both L2 and IP) are bound to a single port on a CES 2000 Series or CER 2000 Series device, consider ending the IP ACL with a "permit any any" filter to ensure that the L2 ACL is also applied to incoming packets. (See also Configuration considerations for dual inbound ACLs on CES 2000 Series and CER 2000 Series devices on page 99.)

  NOTE
  Do not apply an empty ACL (an ACL ID without any corresponding entries) to an interface. If you accidentally do this, the software applies the default ACL action, deny all, to the interface and thus denies all traffic.

Types of IPv4 ACLs

IPv4 ACLs can be configured as standard or extended ACLs. A standard ACL permits or denies packets based on source IP address. An extended ACL permits or denies packets based on source and destination IP address and also based on IP protocol information.

Standard or extended ACLs can be numbered or named. Standard numbered ACLs have an ID of 1 – 99. Extended numbered ACLs are numbered 100 – 199. IDs for standard or extended ACLs can be a character string. In this document, an ACL with a string ID is called a named ACL.

General IPv4 ACL configuration guidelines

Consider the following configuration guidelines for IPv4 ACLs:

- On physical interfaces and LAG groups, IPv4 ACLs are supported for routed traffic only.
- On virtual Ethernet (VE) interfaces, IPv4 ACLs are supported both for routed and for switched traffic. Switched traffic requires that the VE-traffic flag be enabled.
- Both inbound and outbound ACLs are supported.
- The maximum number of IPv4 ACLs that you can create varies with the platform.
- By default, you can create up to 4096 IPv4 ACL rules in all the IPv4 ACLs on the device. To modify the maximum, refer to Modifying the Layer 3 ACL rule maximum on page 120.

- On XMR Series and MLX Series devices, you cannot bind Layer 2 ACLs and IPv4 or IPv6 ACLs for inbound traffic to the same port. However, you can perform the following configuration:
  - On one port, bind a Layer 2 ACL.
  - On a second port, bind one or both of the following:
    - › An IPv4 ACL (standard or extended)
    - › An IPv6 ACL (standard or extended)
- CES 2000 Series and CER 2000 Series devices enable you to bind a Layer 2 ACL, an IPv4 ACL, and an IPv6 ACL for inbound traffic on the same port.
- On all platforms, for outbound traffic, you can bind only one ACL—L2 or IPv4 or IPv6.
- You cannot enable any of the following features on the interface if an ACL is already applied to that interface:
  - ACL-based rate limiting
  - Policy-based routing (PBR)
  - VLAN ID Translation or Inner VLAN ID translation feature
- Support for ACLs on MPLS VPN Endpoints - ACLs can be supported on the following endpoints:
  - IPv4 and IPv6 inbound ACLs are not supported on VPLS, VLL, or VLL-Local endpoints and vice-versa.
  - PBR route-map cannot be applied on VPLS, VLL, or VLL-Local endpoints and vice-versa.
  - The **ip access-group redirect-deny-to-inter** and **ip access-group enable-deny-logging** commands cannot be applied on VPLS, VLL, or VLL-local endpoints and vice versa.
  - IPv4 ACL-based rate limiting is not supported on VPLS and VLL endpoints.
  - Layer-2 ACLs and Layer-2 ACL-based rate limiting is not supported on Layer-3 VPNs.
  - PBR policies are not supported on Layer-3 VPNs.
- For all NetIron devices, if a port has an IPv4 or IPv6 ACL applied, you must remove the ACL bindings before adding that port to a VLAN that has a VE interface.
- IPv4 ACL-based rate limiting on a port that belongs to a VLAN is not supported on a VLAN without a VE configured.
- IPv4 ACL-based rate limiting is not supported on a port that belongs to a VLAN where in Layer-3 Interface(VE) is configured for MPLS.
- To disable IPv4 and IPv6 ACLs on the terminating node of a GRE tunnel, see the "Bypassing ACLs in a GRE tunnel" topic in the *Extreme NetIron Layer 3 Routing Configuration Guide*.
- If you need to downgrade to a version earlier than 5.6, refer to .

Configuration considerations for dual inbound ACLs on CES 2000 Series and CER 2000 Series devices

CES 2000 Series and CER 2000 Series devices enable you to bind a Layer 2 ACL, an IPv4 ACL, and an IPv6 ACL to the same port, as follows:

- A Layer 2 ACL is bound to the port.
- An IPv4 and an IPv6 ACL, or one of the two, are bound to the same port.

The filtering algorithm is as follows:

1. An incoming packet is first examined by the IPv4 ACL or the IPv6 ACL, depending on the protocol.

2. If the packet is denied by the IPv4/IPv6 ACL, the packet is dropped, without being examined by the L2 ACL.

3. If the packet is permitted by the IPv4/IPv6 ACL, it is then examined by the L2 ACL.

However, there is an implicit "deny" at the end of any ACL. To enable the above algorithm, IPv4/IPv6 ACLs intended for this scenario must include a "permit any" filter as the last rule. Such a rule ensures that even packets not explicitly permitted by the IPv4/IPv6 ACL are passed to the L2 ACL.

Dual inbound ACLs can also affect the behavior of ACL accounting. For details, refer to ACL Accounting interactions between L2 ACLs and IP ACLs on page 150.

Configuration considerations for IPv4 outbound ACLs on VPLS, VLL, and VLL-Local endpoints

IPv4 outbound ACLs are supported on VPLS, VLL, and VLL-local endpoints with the following configuration considerations:

- First configure the port as a VPLS, VLL, or VLL-local endpoint and then bind the IPv4 outbound ACL on it.
- First remove the IPv4 outbound ACL from a VPLS, VLL, or VLL-local endpoint before removing the port from the VPLS, VLL, or VLL-local instance or corresponding VLAN.
- First remove the IPv4 outbound ACL from a VPLS, VLL, or VLL-local endpoint(s) before deleting the VPLS, VLL, or VLL-local instance or corresponding VLAN.
- If the VPLS, VLL, or VLL-local endpoint is a LAG port, you must first remove the IPv4 outbound ACL from the primary LAG port before deleting the LAG. This restriction is applicable even if you attempt to delete the lag using **force** keyword.
- If a VLL or VLL-local endpoint is a LAG port with a IPv4 outbound ACL, you have to first remove the IPv4 outbound ACL from the primary LAG port before dynamically removing a port from the LAG.
- Ensure that no VPLS, VLL, or VLL-local endpoint exists with an IPv4 outbound ACL before entering the command: **no router mpls**.

Disabling outbound ACLs for switching traffic

By default, when an outbound ACL is applied to a virtual interface, the Extreme device always filters traffic that is switched from one port to another within the same virtual routing interface. Additional commands have been added that allow you to exclude switched traffic from outbound ACL filtering. This exclusion can be configured globally or on per-port basis. This feature applies to IPv4 and IPv6 ACLs only.

All global and interface level command for disabling outbound ACLs for Switching Traffic are mutually exclusive. If the global command is configured, the interface command is not accepted. If the interface command has already been configured, configuring the global command will remove all individual port commands from the Extreme device's configuration.

> NOTE
> This feature is not recommended for MPLS interfaces.

CAM considerations for CES 2000 Series and CER 2000 Series devices

CAM entries are shared between ingress and egress ACLs. An ACL clause applied to the inbound consumes one CAM entry and an egress ACL clause consumes four CAM entries. The maximum number of egress ACL clauses is 2000 and the maximum number of ingress clauses is 8000.

CES 2000 Series and CER 2000 Series devices have a total of 8000 CAM entries per PPCR (packet processor). The total number of CAM entries in CES 2000 Series and CER 2000 Series devices depends on the number of PPCR (packet processors) in the system. See the table below for the types of ports, the number of PPCR (packet processors), and the total number of CAM entries available:

TABLE 7

CES 2000 Series and CER 2000 Series devices	PPCR (packet processor)	Total CAM entries
24-1G	1	8000
48-1G	2	16000
24-1G & 2-10G	2	16000
24-1G & 4-10G	3	24000
48-1G & 2-10G	3	24000

Globally enabling outbound ACLs for switching traffic

Configuring the **acl-outbound exclude-switched-traffic** command at the general configuration level, allows you to globally exclude all switched traffic from outbound ACL filtering. This feature is configured as shown in the following.

```
device(config)# acl-outbound exclude-switched-traffic ipv4
```

Syntax: [no] acl-outbound exclude-switched-traffic ipv6 | ipv4

The **ipv6** option limits the traffic excluded to IPv6 traffic only.

The **ipv4** option limits the traffic excluded to IPv4 traffic only.

The **ipv4** and **ipv6** options are mutually exclusive within the same command. If you want to configure this command to exclude both IPv4 and IPv6 traffic, you must use two separate commands.

Enabling outbound ACLs for switching traffic per port

Configuring the **if-acl-outbound exclude-switched-traffic** command at the interface configuration level, allows you to exclude all switched traffic from outbound ACL filtering on a per-port basis. With this command, one or more physical ports (for instance all ports within a VLAN) can be configured to exclude switched traffic from outbound ACL filtering.

This feature is configured as shown in the following.

```
device(config)# interface ethernet 3/1
device(config-if-e10000-3/1)# if-acl-outbound exclude-switched-traffic
```

Syntax: [no] if-acl-outbound exclude-switched-traffic [ipv6 | ipv4]

The **ipv6** option limits the traffic excluded to IPv6 traffic only.

The **ipv4** option limits the traffic excluded to IPv4 traffic only.

The **ipv4** and **ipv6** options are mutually exclusive within the same command. If you want to configure this command to exclude both IPv4 and IPv6 traffic, you must use two separate commands.

CES and CER Internal ACLs

By default, CER 2000 Series and CES 2000 Series devices program system-wide ACLs to trap Layer 2 and Layer 3 protocol packets to the CPU. All packets received in CER 2000 Series and CES 2000 Series devices is subjected to internal ACLs, if these packets match the internal ACL policies set then it is forwarded to CPU. In the CPU on further processing, depending upon whether the protocol is configured on the system or an interface, these packets are consumed by the router and switch, or will be flooded on its VLAN domain. In case if the packets received from MPLS uplink for Layer 2 VPN(VPLS/VLL) applications, then the MPLS header will be terminated and internal ACL lookup will be performed for the inner payload.

NOTE
System ACLs have higher precedence than user ACLs, and cannot be changed by the user.

For the following protocols Internal ACLs are created as MAC ACL entries:

1. LACP and 802.1x MAC

2. FDP

3. CDP and PVST

4. Superspan or SpanningTree

5. UDLD

6. MRP

7. Layer 2 Trace

8. MCT control

9. LLDP

10. ARP

11. ISIS

12. CFM

Following example is the output of Internal MAC ACL entries.

```
device# show cam l2acl 1/1
    LP Index  VLAN Src MAC         Dest MAC         Port  Action PRAM
       (Hex)                                                     (Hex)
    1  00001e 0    0000.0000.0000 748e.f811.6340 0     CPU    0001e
    1  000011 0    0000.0000.0000 0000.0000.0000 0     CPU    00011
    1  000010 0    0000.0000.0000 0000.0000.0000 0     CPU    00010
    1  10000b 0    0000.0000.0000 0000.0000.0000 0     CPU    0000b
    1  00000a 0    0000.0000.0000 0180.c200.000e 0     CPU    0000a
    1  000009 0    0000.0000.0000 0304.8000.0500 0     CPU    00009
    1  000008 0    0000.0000.0000 0304.8000.0400 0     CPU    00008
    1  000007 0    0000.0000.0000 0304.8000.0000 0     CPU    00007
    1  000006 0    0000.0000.0000 00e0.5200.0000 0     CPU    00006
    1  000005 0    0000.0000.0000 0380.c200.0000 0     CPU    00005
    1  000004 0    0000.0000.0000 0100.0ccc.cccc 0     CPU    00004
    1  000003 0    0000.0000.0000 01e0.52cc.cccc 0     CPU    00003
    1  000002 0    0000.0000.0000 0180.c200.0002 0     CPU    00002
```

For the following protocols and address Internal ACLs are created as IP ACL entries:

1. DHCP

2. Broadcast IP address

3. OSPF

4. RIPv2

5. VRRPE

6. Multicast IP address

- 224.0.0.0/27

- 224.0.0.2/32

Following example is the output of Internal IP ACL entries.

```
device# show cam l4 1/1
     LP Index  Src IP          SPort Dest IP         DPort Pro Age VLAN Out IF PRAM
                                                                      Action
      1  200031 0.0.0.0         0     224.0.0.0        0     0   N/A      CPU     00031
      1  200030 0.0.0.0         0     224.0.0.18       0     0   N/A      CPU     00030
      1  20002f 0.0.0.0         0     224.0.0.2        0     0   N/A      CPU     0002f
      1  20002e 0.0.0.0         0     224.0.0.9        0     0   N/A      CPU     0002e
      1  20002d 0.0.0.0         0     0.0.0.0          0     89  N/A      CPU     0002d
      1  20002c 0.0.0.0         0     255.255.255.255 0     0   N/A      CPU     0002c
      1  200027 0.0.0.0         0     0.0.0.0          68    17  N/A      CPU     00027
      1  200026 0.0.0.0         0     0.0.0.0          67    17  N/A      CPU     00026
```

Numbered and named IPv4 ACLs

When you configure IPv4 ACLs, you can refer to the ACL by a numeric ID or by an alphanumeric name.

The commands to configure numbered ACLs are different from the commands for named ACLs:

- If you refer to the ACL by a numeric ID, you can use 1 - 99 for a standard ACL or 100 - 199 for an extended ACL. This document refers to this ACL as *numbered ACL*.
- If you refer to the ACL by a name, you specify whether the ACL is a standard ACL or an extended ACL, then specify the name. This document refers to this ACL type as *named ACL*.

You can configure up to 99 standard numbered IPv4 ACLs and 100 extended numbered IPv4 ACLs. The maximum number of named IPv4 ACLs varies with the platform. For details, refer to ACL and rule limits on page 82.

Configuring standard numbered IPv4 ACLs

The following section describes how to configure standard numbered IPv4 ACLs with numeric IDs:

- For configuration information on extended ACLs, refer to Configuring extended numbered IPv4 ACLs on page 106.
- For configuration information on named ACLs, refer to Named IPv4 ACLs on page 114.

Standard ACLs permit or deny packets based on source IP address. You can configure up to 99 standard numbered ACLs. There is no limit to the number of ACL entries an ACL can contain except for the system-wide limitation. For details, refer to ACL and rule limits on page 82

To configure a standard ACL and apply it to inbound traffic on port 1/1, enter the following commands.

```
device(config)# access-list 1 deny host 10.157.22.26
device(config)# access-list 1 deny 10.157.29.12
device(config)# access-list 1 deny host IPHost1
device(config)# access-list 1 permit any
device(config)# int eth 1/1
device(config-if-e10000-1/1)# ip access-group 1 in
device(config)# write memory
```

The commands in this example configure an ACL to deny incoming packets from three source IP addresses from being forwarded on port 1/1. The last ACL entry in this ACL permits all packets that are not explicitly denied by the first three ACL entries.

The ACL functionality for filtering traffic is enhanced with sequence numbers that enable users to insert, modify or delete rules at any position, without having to remove and reapply the entire ACL. A sequence number is assigned to each ACL entry and ACL rules are applied in the order of lowest to highest sequence number. Therefore, you can insert a new filter rule at any position you want in the ACL table by specifying the sequence number. If you do not specify a sequence number a default sequence number is applied to each ACL entry. The default value is 10+ the sequence number of the last ACL entry provisioned in the ACL table. Therefore, when you do not specify a sequence number, the rule is added to the end of the ACL table. The default value for the first entry in an IPv4 ACL table is "10".

To configure an ACL filter rule with the sequence number "4" for ACL "1", enter the following command:

```
device(config)# access-list 1 sequence 4 permit any any
```

If the sequence number "4" is already used by another ACL filter rule, the following error message is displayed.

```
"Error: Entry with sequence 4 already exists!"
```

If you specify a sequence number which is greater than the limit (214748364) the following error message is displayed.

```
"Error: Valid range for sequence is 1 to 214748364"
```

Re-sequencing a standard numbered ACL table

To allow new ACL entries to be inserted between ACL entries that have consecutive sequence numbers, you can create space between sequence numbers of adjacent filters by regenerating the ACL table.

To re-sequence ACL table "1", use the following command.

```
device(config)# access-list 1 regenerate-seq-num
```

This command regenerates the filter sequence numbers in steps of 10, assigning the default sequence number "10" to the first entry in the table.

> **NOTE**
> If sequence numbers generated by the **regenerate-seq-num** command cross the limit (214748364), then re-sequencing of ACL filters will not take place and the following error message is displayed."Error: Valid range for sequence is 1 to 214748364".

> **NOTE**
> The **regenerate-seq-num** command is not allowed while **tftp copy** in progress.

Deleting a standard numbered ACL entry

You can delete an ACL filter rule by providing the sequence number or without providing the sequence number.

To delete an ACL filter rule without providing a sequence number you must specify the filter rule attributes. To delete an ACL filter rule providing a sequence number you can provide the sequence number alone or the sequence number and the other filter rule attributes.

1. To delete a filter rule with the sequence number "20" from access list "100" by specifying the sequence number alone, enter the following command.

   ```
   device(config)# no access-list 100 sequence 20
   ```

2. To delete a filter rule without specifying the sequence number, specify the filter rule attributes.

   ```
   device(config)# no access-list 100 permit any any
   ```

Standard numbered IPv4 ACL syntax

This section presents the syntax for creating and re-sequencing a standard IPv4 ACL and for binding the ACL to an interface.

Use the **access-list regenerate-seq-num** command to re-sequence the ACL table. Use the **ip access-group** command in the interface level to bind the ACL to an interface.

Syntax: [no] **access-list** *num* [**sequence** *num*] **deny** | **permit** [**vlan** *vlan-id*] { **host** { *source-ip* | *hostname* } | *hostname wildcard* | *source-ip/mask-bits* | **any** } [**log**]

Syntax: access-list *num* **regenerate-seq-num** [*num*]

Syntax: [no] ip access-group num in

Parameters for standard numbered ACL statements

num	Enter 1 - 99 for a standard ACL.
sequence *num*	The **sequence** parameter specifies where the conditional statement is to be added in the access list. You can add a conditional statement at particular place in an access list by specifying the entry number using the **sequence** keyword. The range is from 1 through 214748364. If the **sequence***num* is specified for the first ACL clause "clause-1" and not specified for the second ACL clause "clause-2", then the system rounds off the sequence number of the last ACL filter rule provisioned to the next 10th digit. The default value for the first clause in an IPv6 ACL table is "10".
deny | **permit**	Enter **deny** if the packets that match the policy are to be dropped; **permit** if they are to be forwarded.
vlan *vlan-id*	Specifies the vlan-id for the ACL filter rule.
host *source-ip* | *hostname*	Specify a host IP address or name. When you use this parameter, you do not need to specify the mask. A mask of all zeros (0.0.0.0) is implied. **NOTE** To specify the host name instead of the IP address, the DNS server must be configured using the **ip dns server-address***ip-addr* command at the global configurationlevel.
hostname	Specifies the host name for the policy.
wildcard	Specifies the portion of the source IP host address to match against. The *wildcard* is a four-part value in dotted-decimal notation (IP address format) consisting of ones and zeros. Zeros in the mask mean the packet's source address must match the *source-ip*. Ones mean any value matches. For example, the *source-ip* and *wildcard* values 10.157.22.26 0.0.0.255 mean that all hosts in the Class C subnet 10.157.22.x match the policy. If you prefer to specify the mask value in Classless Inter domain Routing (CIDR) format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of "10.157.22.26 0.0.0.255" as "10.157.22.26/24". The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros. For example, if you specify 10.157.22.26/24 or 10.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 10.157.22.0/24 (if you have enabled display of subnet lengths) or 10.157.22.0 0.0.0.255 in the startup-config file. If you enable the software to display IP subnet masks in CIDR format, the mask is saved in the file in "*mask-bits* " format. You can use the CIDR format to configure the ACL entry regardless of whether the software is configured to display the masks in CIDR format. **NOTE** If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but they are shown with subnet mask in the display produced by the **show access-list** command.
any	Use this parameter to configure the policy to match on all host addresses.

log	Deny rules—Enables inbound logging for the rule. An additional requirement for deny logging is that the **ip access-group enable-deny-logging** command be in effect.
	Permit rules—(Not supported on CER 2000 Series or CES 2000 Series devices) Enables inbound logging for the rule. An additional requirement for permit logging is that the **ip access-group enable-permit-logging** command be in effect.

Parameters for regenerating IPv4 ACL table sequence numbers

num	Specifies the number of the ACL table to re-sequence
regenerate-seq-num [num]	(Optional) Specifies the initial sequence number for the access list after regeneration. The valii range is from 1 through 214748364. The default value is 10. ACL filter rule sequence numbers are regenerated in steps of 10.

Parameters to bind standard ACLs to an interface

Use the **ip access-group** command to bind the ACL to an inbound interface and enter the ACL number for *num*.

Configuring extended numbered IPv4 ACLs

This section describes how to configure extended numbered IPv4 ACLs.

- For configuration information on standard ACLs, refer to Configuring standard numbered IPv4 ACLs on page 103.
- For configuration information on named ACLs, refer to Named IPv4 ACLs on page 114.

Extended ACLs let you permit or deny packets based on the following information:

- IP protocol
- Source IP address or host name
- Destination IP address or host name
- Source TCP, UDP, or SCTP port
- Destination TCP, UDP, or SCTP port

The IP protocol can be one of the following well-known names or any IP protocol number from 0 - 255:

- Internet Control Message Protocol (ICMP)
- Internet Group Management Protocol (IGMP)
- Internet Gateway Routing Protocol (IGRP)
- Internet Protocol (IP)
- Open Shortest Path First (OSPF)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- Stream Control Transmission Protocol (SCTP)

    NOTE
    ACL do not account IGMP packets when multicast is enabled.

For TCP, UDP, or SCTP, you also can specify a comparison operator and port name or number. For example, you can configure a policy to block web access to a specific web site by denying all TCP port 80 (HTTP) packets from a specified source IP address to the web site's IP address.

To configure an extended access list that blocks all Telnet traffic received on port 1/1 from IP host 10.157.22.26, create the ACL with permit and deny rules, then bind the ACL to port 1/1 using the **ip access-group** command. Enter the following commands.

```
device(config)# access-list 101 deny tcp host 10.157.22.26 any eq telnet
device(config)# access-list 101 permit ip any any
device(config)# int eth 1/1
device(config-if-e10000-1/1)# ip access-group 101 in
device(config)# write memory
```

Here is another example of commands for configuring an extended ACL and applying it to an interface. These examples show many of the syntax choices.

```
device(config)# access-list 102 perm icmp 10.157.22.0/24 10.157.21.0/24
device(config)# access-list 102 deny igmp host rkwong 10.157.21.0/24
device(config)# access-list 102 deny igrp 10.157.21.0/24 host rkwong
device(config)# access-list 102 deny ip host 10.157.21.100 host 10.157.22.1
device(config)# access-list 102 deny ospf any any
device(config)# access-list 102 permit ip any any
```

The first entry permits ICMP traffic from hosts in the 10.157.22.x network to hosts in the 10.157.21.x network.

The second entry denies IGMP traffic from the host Extreme device named "rkwong" to the 10.157.21.x network.

The third entry denies IGRP traffic from the 10.157.21.x network to the host Extreme device named "**rkwong** ".

The fourth entry denies all IP traffic from host 10.157.21.100 to host 10.157.22.1.

The fifth entry denies all OSPF traffic.

The sixth entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

The following commands apply ACL 102 to the incoming traffic on port 1/2 and to the outgoing traffic on port 4/3.

```
device(config)# int eth 1/2
device(config-if-e10000-1/2)# ip access-group 102 in
device(config-if-e10000-1/2)# exit
device(config)# int eth 4/3
device(config-if-e10000-4/3)# ip access-group 102 out
device(config)# write memory
```

Here is another example of an extended ACL.

```
device(config)# access-list 103 deny tcp 10.157.21.0/24 10.157.22.0/24
device(config)# access-list 103 deny tcp 10.157.21.0/24 eq ftp 10.157.22.0/24
device(config)# access-list 103 deny tcp 10.157.21.0/24 10.157.22.0/24 lt telnet neq 5
device(config)# access-list 103 deny udp any range 5 6 10.157.22.0/24 range 7 8
device(config)# access-list 103 permit ip any any
```

The first entry in this ACL denies TCP traffic from the 10.157.21.x network to the 10.157.22.x network.

The second entry denies all FTP traffic from the 10.157.21.x network to the 10.157.22.x network.

The third entry denies TCP traffic from the 10.157.21.x network to the 10.157.22.x network if the TCP port number of the traffic is less than the well-known TCP port number for Telnet (23) and if the TCP port is not equal to 5. Thus, TCP packets with a TCP port number equal to 5 or greater than 23 are allowed.

The fourth entry denies UDP packets from any source to the 10.157.22.x network, if the UDP port number from the source network is 5 or 6 and the destination UDP port is 7 or 8.

The fifth entry permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL would deny all incoming or outgoing IP traffic on the ports to which you assign the ACL.

The following commands apply ACL 103 to the incoming and outgoing traffic on ports 2/1 and 2/2.

```
device(config)# int eth 2/1
device(config-if-e10000-2/1)# ip access-group 103 in
device(config-if-e10000-2/1)# ip access-group 103 out
device(config-if-e10000-2/1)# exit
device(config)# int eth 2/2
device(config-if-e10000-2/2)# ip access-group 103 in
device(config-if-e10000-2/2)# ip access-group 103 out
device(config)# write memory
```

The following example shows how sequence numbers are assigned to ACL entries. This example configures filter rules for the extended numbered IPv4 ACL "100".

```
device(config)# access-list 100 permit udp any any
device(config)# access-list 100 sequence 11 permit tcp any any
device(config)# access-list 100 permit icmp any any
```

The first entry in this example permits all UDP traffic. As this is the first entry in the ACL table and a sequence number is not specified, the system assigns the default sequence number "10". The second entry, which specifies the sequence number "11", permits all TCP traffic. The third entry permits all ICMP traffic. Again, the sequence number is not specified and the system assigns the default sequence number "21" (10+ the sequence number of the last ACL filter rule provisioned in the table) to this entry. The output from the **show access-list** command for the ACL table is:

```
10: access-list 100 permit udp any any
11: access-list 100 sequence 11
 permit tcp any any
21: access-list 100 permit icmp any any
```

And the output from the **show running-config** command is:

```
access-list 100 permit udp any any
access-list 100 sequence 11
 per tcp any any
access-list 100 permit icmp any any
```

The **show access-list** or **show running-config** commands only display user-configured sequence numbers. In these examples, the display of "**sequence 11** " after the access list number indicates a user-configured sequence number for the ACL entry. When the ACL entry sequence number is system-generated it is not displayed.

To insert more rules between adjacent sequence numbers "10" and "11", you need to re-sequence the ACL table first. The **regenerate-seq-num** command generates new sequence numbers for ACL table entries creating space between the sequence numbers of adjacent filters. To re-sequence the ACL table "100", enter the following command.

```
device(config)# access-list 100 regenerate-seq-num
```

This command resequences entries in the ACL table in steps of 10 so that the output from the **show access-list** command is:

```
10: access-list 100 permit udp any any
20: access-list 100 sequence 20 permit tcp any any
30: access-list 100 permit icmp any any
```

And the output from the **show running-config** command is:

```
access-list 100 permit udp any any
access-list 100 sequence 20 per tcp any any
access-list 100 permit icmp any any
```

Extended numbered IPv4 ACL syntax

This section presents the syntax for creating and re-sequencing an extended IPv4 ACL and for binding the ACL to an interface. Use the **access-list regenerate-seq-num** command to re-sequence the ACL table**.** Use the **ip access-group** command in the interface level to bind the ACL to an interface.

Syntax: [no] access-list *num* [**sequence** *num*] **deny** | **permit** [**vlan** *vlan-id*] *ip-protocol* { *source-ip* | *hostnamewildcard* | **any** } [*operatorsource-tcp/udp/sctp-port*] { *destination-ip* | *hostnamewildcard* | **any** } [*operatordestination-tcp/udp/sctp-port*] [*icmp-type*] [**established**] [**precedence** { *name* | *num* }] [**tos** { *name* | *number* }] [**dscp-mapping** *number*] [**dscp-marking** *number*] | [{ **fragment** } | **non-fragment** }] [**option value** | **name** | **keyword**] [**priority** *priority-value* | **priority-force** *priority-value* | **priority-mapping** *priority-value*] [**log**] [**mirror**]

Syntax: access-list *num* regenerate-seq-num [*num*]

Syntax: [no] ip access-group num in | out

General parameters for extended numbered IPv4 ACLs

The following parameters apply to any extended ACL you are creating.

num	Enter 100 - 199 for an extended ACL.
sequence *num*	The **sequence** parameter specifies where the conditional statement is to be added in the access list. You can add a conditional statement at particular place in an access list by specifying the entry number using the **sequence** keyword. The range is from 1 through 214748364. If the **sequence***num* option is specified for the first ACL clause "clause-1" and not specified for the second ACL clause "clause-2", then the system rounds off the sequence number of the last ACL filter rule provisioned to the next 10th digit. The default value for the first clause in an IPv6 ACL table is "10".
deny	**permit**	Enter **deny** if the packets that match the policy are to be dropped; **permit** if they are to be forwarded.
ip-protocol	Indicate the type of IP packet you are filtering. You can specify a well-known name for any protocol whose number is less than 255. For other protocols, you must enter the number. Enter "?" instead of a protocol to list the well-known names recognized by the CLI.
source-ip	*hostname*	Specify the source IP host for the policy. If you want the policy to match on all source addresses, enter **any**.
wildcard	Specifies the portion of the source IP host address to match against. The *wildcard* is a four-part value in dotted-decimal notation (IP address format) consisting of ones and zeros. Zeros in the mask mean the packet's source address must match the *source-ip*. Ones mean any value matches. For example, the *source-ip* and *wildcard* values 10.157.22.26 0.0.0.255 mean that all hosts in the Class C subnet 10.157.22.x match the policy. If you prefer to specify the wildcard (mask value) in Classless Inter domain Routing (CIDR) format, you can enter a forward slash after the IP address, then enter the number of significant bits in the mask. For example, you can enter the CIDR equivalent of "10.157.22.26 0.0.0.255" as "10.157.22.26/24". The CLI automatically converts the CIDR number into the appropriate ACL mask (where zeros instead of ones are the significant bits) and changes the non-significant portion of the IP address into zeros. For example, if you specify 10.157.22.26/24 or 10.157.22.26 0.0.0.255, then save the changes to the startup-config file, the value appears as 10.157.22.0/24 (if you have enabled display of subnet lengths) or 10.157.22.0 0.0.0.255 in the startup-config file. The IP subnet masks in CIDR format is saved in the file in "/*mask-bits* " format.

	If you use the CIDR format, the ACL entries appear in this format in the running-config and startup-config files, but are shown with subnet mask in the display produced by the **show access-list** command.
destination-ip | *hostname*	Specifies the destination IP host for the policy. If you want the policy to match on all destination addresses, enter **any.**
fragment	Enter this keyword if you want to filter fragmented packets. Refer to Enabling ACL filtering of fragmented or non-fragmented packets on page 125. NOTE The **fragmented** and **non-fragmented** parameters cannot be used together in an ACL entry.
non-fragment	Enter this keyword if you want to filter non-fragmented packets. Refer to Enabling ACL filtering of fragmented or non-fragmented packets on page 125. NOTE The **fragmented** and **non-fragmented** parameters cannot be used together in an ACL entry.
priority | **priority-force** | **priority-mapping**	The **Priority** option assigns internal priority to traffic that matches the ACL. In addition to changing the internal forwarding priority, if the outgoing interface is an 802.1q interface, this option maps the specified priority to its equivalent 802.1p (QoS) priority and marks the packet with the new 802.1p priority. This option is applicable for inbound ACLs only. The **Priority-force** option assigns internal priority to packets of traffic that match the ACL, even though the incoming packet may be assigned a higher priority. This option is applicable for inbound ACLs only. The **priority-mapping** option matches on the packet's 802.1p value. This option does not change the packet's forwarding priority through the device or mark the packet. This keyword is applicable for both inbound and outbound ACLs.
priority-value	The **priority-value** variable specifies one of the following QoS queues for use with the **priority, priority-force** or **priority-mapping** options: 0 - qosp0 1 - qosp1 2 - qosp2 3 - qosp3 4 - qosp4 5 - qosp5 6 - qosp6 7 - qosp7
log	**Deny** rules—Enables inbound logging for the rule. An additional requirement for deny logging is that the **ip access-group enable-deny-logging** command be in effect. **Permit** rules—(Not supported on CER 2000 Series or CES 2000 Series devices) Enables inbound logging for the rule. An additional requirement for permit logging is that the **ip access-group enable-permit-logging** command be in effect.
mirror	Specifies mirror packets matching ACL permit clause. For more information on configuring the **acl-mirror-port** command, refer to *Extreme NetIron Layer 2 Switching Configuration Guide.*

Parameters to filter TCP, UDP, or SCTP packets

Use the parameters below if you want to filter traffic with the TCP, UDP, or SCTP packets. These parameters apply only if you entered **tcp**, **udp**, or **sctp** for the *ip-protocol* parameter. For example, if you are configuring an entry for HTTP, specify **tcp eq http**.

operator	Specifies a comparison operator for the TCP, UDP, or SCTP port number. You can enter one of the following operators: • **eq** - The policy applies to the TCP, UDP, or SCTP port name or number you enter after **eq**. • **gt** - The policy applies to TCP, UDP, or SCTP port numbers greater than the port number or the numeric equivalent of the port name you enter after **gt**. • **lt** - The policy applies to TCP, UDP, or SCTP port numbers that are less than the port number or the numeric equivalent of the port name you enter after **lt**. • **neq** - The policy applies to all TCP, UDP, or SCTP port numbers except the port number or port name you enter after **neq**. • **range** - The policy applies to all TCP, UDP, or SCTP port numbers that are between the first TCP, UDP, or SCTP port name or number and the second one you enter following the **range** parameter. The range includes the port names or numbers you enter. For example, to apply the policy to all ports between and including 23 (Telnet) and 53 (DNS), enter the following: **range 23 53**. The first port number in the range must be lower than the last number in the range. • **established** - This operator applies only to TCP packets. If you use this operator, the policy applies to TCP packets that have the ACK (Acknowledgment) or RST (Reset) bits set on (set to "1") in the Control Bits field of the TCP packet header. Thus, the policy applies only to established TCP sessions, not to new sessions. Refer to Section 3.1, "Header Format", in RFC 793 for information about this field. **NOTE** This operator applies only to destination TCP ports, not source TCP ports.
source-tcp/udp/sctp-port	Enter the source TCP, UDP, or SCTP port number.
destination-tcp/udp/sctp-port	Enter the destination TCP, UDP, or SCTP port number.

Filtering traffic with ICMP packets

Use the following parameters if you want to filter traffic that contains ICMP packets. These parameters apply only if you specified **icmp** as the *ip-protocol* value.

icmp-type	Enter one of the following values, depending on the software version the Extreme device is running: • any-icmp-type • echo • echo-reply • information-request • mask-reply • mask-request • parameter-problem • redirect

	NOTE The redirect parameter is not supported on the CES 2000 Series or CER 2000 Series devices. • source-quench • time-exceeded **NOTE** The time-exceeded parameter is not supported on CES 2000 Series or CER 2000 Series devices. • timestamp-reply • timestamp-request • unreachable • *num* **NOTE** If the ACL is for the inbound traffic direction on a virtual routing interface, you also can specify a subset of ports within the VLAN containing that interface when assigning an ACL to the interface. Refer to Numbered and named IPv4 ACLs on page 103.

Using ACL QoS options to filter packets

You can filter packets based on their QoS values by entering values for the following parameters:

tos *name* | *num*	Specify the IP ToS name or number.
	
	    NOTE
	    This parameter is not supported on CES 2000 Series or CER 2000 Series devices.
	
	You can specify one of the following:
	• **max-reliability** or **2** - The ACL matches packets that have the maximum reliability ToS. The decimal value for this option is 2.
	• **max-throughput** or **4** - The ACL matches packets that have the maximum throughput ToS. The decimal value for this option is 4.
	• **min-delay** or **8** - The ACL matches packets that have the minimum delay ToS. The decimal value for this option is 8.
	• **normal** or **0** - The ACL matches packets that have the normal ToS. The decimal value for this option is 0.
	• *num* - A number from 0 - 15 that is the sum of the numeric values of the options you want. The ToS field is a four-bit field following the Precedence field in the IP header. You can specify one or more of the following. To select more than one option, enter the decimal value that is equivalent to the sum of the numeric values of all the ToS options you want to select. For example, to select the **max-reliability** and **min-delay** options, enter number 10. To select all options, select 15.

dscp-mapping *number*	The ACL matches packets on the DSCP value. This option does not change the packet's forwarding priority through the device or mark the packet.

Parameters to mark the DSCP value in a packet

Specify the DSCP value to a packet by entering the following parameter:

Use **dscp-marking** *number* to mark the DSCP value in the incoming packet with the value you specify.**Dscp-marking** is not supported on outbound ACLs.

Parameters for regenerating IPv4 ACL table sequence numbers

num	Specifies the standard ACL number
regenerate-seq-num [*num*]	Specifies the initial sequence number for the access list after regeneration. The valid range is from 1 through 214748364. The default value is 10. ACL filter rule sequence numbers are regenerated in steps of 10.

Parameters to bind standard ACLs to an interface

Use the **ip access-group** command to bind the ACL to an interface and enter the ACL number for *num*.

Parameters to filter IP option packets

You can filter IP Option traffic based upon the content of the IP option field in the IP header.

> **NOTE**
> This feature is not supported on CES 2000 Series or CER 2000 Series devices.

- value
- You can match based upon a specified IP Option value. Values between 1 - 255 can be used.
- keyword
- You can use the any keyword to match packets with IP Options or use the ignore keyword to match packets with or without IP Options. If you are configuring a filter to permit or deny rsvp or igmp packets, it will ignore IP options within the packet by default.
- name
- You can match by using any of the following well-known options by name: eol - Matches IP Option packets that contain the eol option. extended-security - Matches IP Option packets that contain the extended security option. loose-source-route - Matches IP Option packets that contain the loose source route option. no-op - Matches IP Option packets that contain the no-op option. record-route - Matches IP Option packets that contain the record route option. router-alert - Matches IP Option packets that contain the router alert option. security - Matches IP Option packets that contain the security option. streamid - Matches IP Option packets that contain the stream id option. strict-source-route - Matches IP Option packets that contain the strict source route option. timestamp - Matches IP Option packets that contain the timestamp option.

Please note, the behavior of an implicit **deny ip any any** ACL filter is different than that of an explicit **deny ip any any** filter as described in the following:

- Explicit **deny ip any any** will only apply to non-option packets.
- Explicit **deny ip any any option ignore** will apply to both option and non-option packets
- Implicit **deny ip any any** will apply to both option and non-option packets

Named IPv4 ACLs

The commands for configuring named ACL entries are different from the commands for configuring numbered ACL entries.

The command to configure a numbered ACL is **access-list** . The command for configuring a named ACL is **ip access-list** . In addition, when you configure a numbered ACL entry, you specify all the command parameters on the same command. When you configure a named ACL, you specify the ACL type (standard or extended) and the ACL name with one command, which places you in the configuration level for that ACL. Once you enter the configuration level for the ACL, the command syntax is similar as the syntax for numbered ACLs.

The following topics show how to configure a named standard ACL entry and a named extended ACL entry.

Standard named IPv4 ACL syntax and configuration

To configure a standard named IPv4 ACL, enter commands such as the following.

```
device(config)# ip access-list standard Net1
device(config-std-nacl-Net1)# deny host 10.157.22.26
device(config-std-nacl-Net1)# deny 10.157.29.12
device(config-std-nacl-Net1)# deny host IPHost1
device(config-std-nacl-Net1)# permit any
device(config-std-nacl-Net1)# exit
device(config)# int eth 1/1
device(config-if-e10000-1/1)# ip access-group Net1 in
```

The commands in this example configure a standard ACL named "Net1". The entries in this ACL deny packets from three source IP addresses from being forwarded on port 1/1. Since the implicit action for an ACL is "deny", the last ACL entry in this ACL permits all packets that are not explicitly denied by the first three ACL entries. For an example of how to configure the same entries in a numbered ACL, refer to Configuring standard numbered IPv4 ACLs on page 103.

The command prompt changes after you enter the ACL type and name. The "std" in the command prompt indicates that you are configuring entries for a standard ACL. For an extended ACL, this part of the command prompt is "ext". The "nacl" indicates that are configuring a named ACL.

Syntax: [no] ip access-list standard { *string* | *num* }

Syntax: [no] [**sequence** *num*] { **deny** | **permit** } [**vlan** *vlan-id*] **host** { *source-ip* | *hostname* } | { *hostnamewildcard* | *source-ip/mask-bits* } | **any** [**log**]

Syntax: **regenerate-seq-num** [*num*]

Syntax: [no] ip access-group num in

The **standard** parameter indicates the ACL type.

The *string* parameter is the ACL name. You can specify a string of up to 256 alphanumeric characters. The string "All" cannot be used to form creation of a named standard ACL. You can use blanks in the ACL name if you enclose the name in quotation marks (for example, "ACL for Net1"). The *num* parameter allows you to specify an ACL number if you prefer. If you specify a number, you can specify from 1 - 99 for standard ACLs or 100 - 199 for extended ACLs.

> **NOTE**
> For convenience, the software allows you to configure numbered ACLs using the syntax for named ACLs. The software also still supports the older syntax for numbered ACLs. Although the software allows both methods for configuring numbered ACLs, numbered ACLs are always formatted in the startup-config and running-config files in using the older syntax, as follows.
> ```
> access-list 1 deny host 10.157.22.26access-list 1 deny 10.157.22.0 0.0.0.255access-list 1
> permit any access-list 101 deny tcp any any eq http
> ```

The options at the ACL configuration level and the syntax for the **ip access-group** command are the same for numbered and named ACLs and are described in Configuring standard numbered IPv4 ACLs on page 103.

To re-sequence a named standard ACL table, enter the following command:

```
device(config)# ip access-list standard Net1
device(config-std-nacl-Net1)# regenerate-seq-num
```

Extended named IPv4 ACL syntax and configuration

To configure an extended named IPv4 ACL, enter commands such as the following.

```
device(config)# ip access-list extended "block Telnet"
device(config-ext-nacl-block telnet)# deny tcp host 10.157.22.26 any eq telnet
device(config-ext-nacl-block telnet)# permit ip any any
device(config-ext-nacl-block telnet)# exit
device(config)# int eth 1/1
device(config-if-e10000-1/1)# ip access-group "block Telnet" in
```

> **NOTE**
> The command prompt changes after you enter the ACL type and name. The "ext" in the command prompt indicates that you are configuring entries for an extended ACL. The "nacl" indicates that are configuring a named ACL.

To re-sequence a named extended ACL table, enter the following command:

```
device(config-ext-nacl-block telnet)# regenerate-seq-num
```

Syntax: [no] ip access-list extended *string* | *num*

Syntax: [no] [**sequence** *num*] **deny** | **permit** [**vlan** *vlan-id*] *ip-protocol* { *source-ip* | *hostnamewildcard* | **any** } [*operatorsource-tcp/udp/ sctp-port*] { *destination-ip* | *hostnamewildcard* | **any** } [*operatordestination-tcp/udp/sctp-port*] [*icmp-type*] [**established**] [**precedence** { *name* | *num* }] [**tos** { *name* | *number* }] [**dscp-mapping** *number*] [**dscp-marking** *number*] | [{ **fragment** | **non-fragment** }] [**option** *value* | *name* | *keyword*] [**priority** *priority-value* | **priority-force** *priority-value* | **priority-mapping** *priority-value*] [**log**] [**mirror**]

Syntax: regenerate-seq-num [*num*]

Syntax: [no] ip access-group *string* | *num* **in** | **out**

The options at the ACL configuration level and the syntax for the **ip access-group** command are the same for numbered and named ACLs and are described in Configuring extended numbered IPv4 ACLs on page 106.

> **NOTE**
> The string **"all"** cannot be used as ACL name while defining
> ACLs.

The **precedence** option for of an IP packet is set in a three-bit field following the four-bit header-length field of the packet's header. This parameter is not supported on CES 2000 Series or CER 2000 Series devices. You can specify one of the following name or number: critical or 5 - The ACL matches packets that have the critical precedence. If you specify the option number instead of the name, specify number 5. flash or 3 - The ACL matches packets that have the flash precedence. If you specify the option number instead of the name, specify number 3. flash-override or 4 - The ACL matches packets that have the flash override precedence. If you specify the option number instead of the name, specify number 4. immediate or 2 - The ACL matches packets that have the immediate precedence. If you specify the option number instead of the name, specify number 2. internet or 6 - The ACL matches packets that have the internetwork control precedence. If you specify the option number instead of the name, specify number 6. network or 7 - The ACL matches packets that have the network control precedence. If you specify the option number instead of the name, specify number 7. priority or 1 - The ACL matches packets that have the priority precedence. If you specify the option number instead of the name, specify number 1. routine

or 0 - The ACL matches packets that have the routine precedence. If you specify the option number instead of the name, specify number 0.

Deleting rules from named IPv4 ACLs

You can delete an ACL filter rule by providing the sequence number or without providing the sequence number.

To delete an ACL filter rule without providing a sequence number you must specify the filter rule attributes. To delete an ACL filter rule providing a sequence number you can provide the sequence number alone or the sequence number and the other filter rule attributes.

1. Enter the following command to display the contents of the ACL list.

```
device#show access-list name entry
Standard IP access list entry
10: deny host 10.2.4.5
20: deny host 10.1.1.1
30: deny host 10.6.7.8
40: permit any
```

2. To delete the second ACL entry from the list by specifying the sequence number only, enter the following commands.

```
device(config)#ip access-list standard entry
device(config-std-nacl-entry)# no sequence 20
```

3. Enter the **show access-list** command to display the updated list.

```
device(config)# ip show access-list name entry
Standard IP access list entry
10: deny host 10.2.4.5
30: deny host 10.6.7.8
40: permit any
```

Displaying ACL definitions

To display the ACLs configured on an Extreme device, use the **show access-list** command.

To display the total number of Layer-2 and IPv4 access lists and the number of filters configured for each list, use the **show access-list count** command.

```
device(config)#show access-list count
Total 4 ACLs exist.
ACL 102, total 10 clauses
ACL 105, total 15 clauses
ACL 400, total 100 clauses
ACL 401, total 2 clauses
```

> **NOTE**
> Empty ACLs that are applied to interfaces are included in the total ACL count but are not displayed.

For a numbered ACL, you can enter a command such as the following.

```
device(config)#show access-list 99
ACL configuration:
!
Standard IP access list 10
10: access-list 99 deny host 10.10.10.1
20: access-list 99 permit any
```

For a named ACL, enter a command such as the following.

```
device(config)#show access-list name entry
Standard IP access list entry
```

```
10: deny host 5.6.7.8
20: deny host 192.168.12.3
30: permit any
```

Syntax: **show access-list** { **count** | **number** | **name** acl-name | **all** }

The **count** parameter specifies displaying the total number of Layer-2 and IPv4 access lists and the number of filters configured for each list. Empty ACLs that are applied to interfaces are included in the total ACL count but are not displayed.

The *number* variable specifies displaying information for a specific numbered ACL:

- 1 - 99 for standard ACLs
- 100 - 199 for extended ACLs

The **name** *acl-name* option specifies displaying information for a specific named ACL.

Enter **all** if you want to display all the ACLs configured on the device.

Impact of Layer 2 ACL limits on Layer 3 ACL resources

The Layer 2 **system-max l2-acl-table-entries** setting can effect Layer 3 ACL functionality.

CES 2000 Series and CER 2000 Series devices have 8192 CAM entries, and 1000 ingress Layer-2 numbered ACL takes 1000 CAM entries, while egress Layer-2 numbered ACL needs 2000 CAM entries. If users configure the maximum Layer-2 ACL, the other types of ACL, such as IP and IPv6 ACL, will have limited space.

The change may also impact memory use in MLX Series and XMR Series series. The memory increase can be from 2.5M to 10M, depending on **system-max l2-acl-table-entries** configurations:

```
device(config)#system-max l2-acl-table-entries
   DECIMAL   Valid range 64 to 256 (default: 64)
```

Once the above is set to 256, and the user configures one Layer-2 ACL with 256 entries, then each of other Layer-2 ACL will take memory of 256 entries, even though each of these ACL has a single entry only.

This does not affect CAM occupation, that is, a single entry Layer-2 ACL still take a CAM entry, even though **system-max l2-acl-table-entries** is configured to 256.

VLAN Accounting

VLAN accounting already exists in previous release. Now it works with the increased ACL infrastructure on CES 2000 Series and CER 2000 Series devices as well.

Syntax: **[no] vlan-accounting**

```
device(config)#vlan 100
device(config-vlan-100)# vlan-accounting
```

Following command will display the VLAN accounting.

```
device(config)#show vlan 100
```

byte-accounting command is deprecated in CES 2000 Series and CER 2000 Series devices. Similar to MLX Series, CES 2000 Series and CER 2000 Series devices use the **vlan-accounting** command.

Simultaneous per VLAN rate limit and QoS

Simultaneous per-VLAN Rate Limit and QoS and add DSCP-marking to the Layer-2 ACL are added to the CES 2000 Series and CER 2000 Series platforms only. VLAN accounting works with the increased ACL infrastructure on CES 2000 Series and CER 2000 Series platforms only.

Currently Layer-2 ACL does not provide an action of **DSCP-marking** , since DSCP belongs to Layer-3. Simultaneous per-VLAN rate limit and QoS requires Layer-2 ACL to mark DSCP. This is available only on CES 2000 Series and CER 2000 Series platforms.

This assumes the packets have both Layer-2 and Layer-3 headers, so that matching Layer-2 will mark Layer-3 parameters. For pure Layer-2 packets without Layer-3 header or non-IP packets, the result is unpredictable, and the ACL may give wrong data. For this reason, a warning message will display once a user configure a DSCP-marking on Layer-2 ACL.

Syntax: [no] **access-list** *num* [**sequence** *num*] **permit** | **deny** *src-macmask* | **any** *dest-macmask* | **any** [*vlan-id* | **any**] [**etype** *etype-str*] [**priority** *queue-value* | **priority-force** *queue-value* | **priority-mapping** *queue-value*] [**log**] **mirror**] [**mark-flow-id**] [**dscp-marking** *number*]

The following example matches VLAN 100 and mark DSCP to 54:

```
device(config)# access-list 1399 permit any any 100 etype any dscp-marking 54
```

> **NOTE**
> This ACL will have unexpected results on non-IP packets. Make sure the traffic on the interfaces are IP packets.

Modifying ACLs

When you configure any ACL, a sequence number is assigned to each ACL entry. If you do not specify the sequence number, the software assigns a sequence number to each entry. The default value is 10+ the sequence number of the last ACL entry provisioned in the ACL table. Therefore, when you do not specify a sequence number, the rule is added to the end of the ACL table. The default value for the first entry in an IPv4 ACL table is "10". The software always applies the ACL entries to traffic in the order of lowest to highest sequence number. The following example configures two entries for ACL "1".

```
device(config)#access-list 1 deny 10.157.22.0/24
device(config)#access-list 1 permit 10.157.22.26
```

The system assigns the sequence number "10" to the first entry and "20" to the second entry so that the output from the **show access-list** command will be:

```
10: access-list 1 deny 10.157.22.0/24
20: access-list 1 permit 10.157.22.26
```

Thus, if a packet matches the first ACL entry in this ACL and is therefore denied, the software does not compare the packet to the remaining ACL entries. In this example, packets from host 10.157.22.26 will always be dropped, even though packets from this host match the second entry.

By specifying the ACL entry sequence number you can insert the entry at any position that you want in an ACL table. For example, enter the following command:

```
device(config)#access-list 1 sequence 15 permit 10.157.22.24
```

The output from the show access-list command is now:

```
10: access-list 1 deny 10.157.22.0/24
15: access-list 1 sequence 15 permit 10.157.22.24
20: access-list 1 permit 10.157.22.26
```

NOTE
Modifications done in the ACL, will be effective in the hardware only after the execution of an explicit rebind command. For more information, refer to Applying ACLs to interfaces on page 123.

There is an alternative method for modifying ACLs on a device. The alternative method lets you upload an ACL list from a TFTP server and replace the ACLs in the Extreme device's running-config file with the uploaded list. Thus, to change an ACL, you can edit the ACL on the file server, then upload the edited ACL to the Extreme device. You then can save the changed ACL to the Extreme device's startup-config file.

ACL lists contain only the ACL entries themselves, not the assignments of ACLs to interfaces. You must assign the ACLs on the Extreme device itself.

NOTE
The only commands that are valid in the ACL list are the **access-list** and **end** commands; other commands are ignored.

Modify an ACL by configuring an ACL list on a file server.

1. Use a text editor to create a new text file. When you name the file, use 8.3 format (up to eight characters in the name and up to three characters in the extension).

    NOTE
    Make sure the Extreme device has network access to the TFTP server.

2. Optionally, clear the ACL entries from the ACLs you are changing by placing commands such as the following at the top of the file.

    ```
    device(config)#no access-list 1
    device(config)#no access-list 101
    ```

    When you load the ACL list into the Extreme device, the software adds the ACL entries in the file after any entries that already exist in the same ACLs. Thus, if you intend to entirely replace an ACL, you must use the **no access-list** *num* command to clear the entries from the ACL before the new ones are added.

3. Place the commands to create the ACL entries into the file. The order of the separate ACLs does not matter, but the order of the entries within each ACL is important. The software applies the entries in an ACL in the order they are listed within the ACL. Here is an example of some ACL entries.

    ```
    device(config)#access-list 1 deny host 10.157.22.26
    device(config)#access-list 1 deny 10.157.22.0 0.0.0.255
    device(config)#access-list 1 permit any
    device(config)#access-list 101 deny tcp any any eq http
    ```

    The software will apply the entries in ACL 1 in the order shown and stop at the first match. Thus, if a packet is denied by one of the first three entries, the packet will not be permitted by the fourth entry, even if the packet matches the comparison values in this entry.

4. Enter the command "**end** " on a separate line at the end of the file. This command indicates to the software that the entire ACL list has been read from the file.

5. Save the text file.

6. On the Extreme device, enter the following command at the Privileged EXEC level of the CLI:**copy tftp running-config***tftp-ip-addrfilename*

    NOTE
    This command will be unsuccessful if you place any commands other than **access-list** and **end** (at the end only) in the file. These are the only commands that are valid in a file you load using the **copy tftp running-config...** command.

7. To save the changes to the Extreme device's startup-config file, enter the following command at the Privileged EXEC level of the CLI:**write memory**

> NOTE
> Do not place other commands in the file. The Extreme device reads only the ACL information in the file and ignores other commands, including **ip access-group** commands. To assign ACLs to interfaces, use the CLI.

Modifying the Layer 3 ACL rule maximum

By default, you can include a total of 4096 rules in Layer 3 (IPv4 and IPv6) ACLs on your device. You can increase this limit to 102,400.

To enable 102,400 IPv4 ACL rules, enter the following command at the Global CONFIG level of the CLI.

```
device(config)# system-max ip-filter-sys 102400
```

Syntax: **[no] system-max ip-filter-sys** *num*

You can load ACLs dynamically by saving them in an external configuration file on flash card or TFTP server, then loading them using one of the following commands:

- **copy { slot1 | slot2 } running** *from-name*
- **ncopy { slot1 | slot2 }** *from-name* **running**
- **copy tftp running-config** *ip-addr filename*
- **ncopy tftp ip-addr** *from-name* **running-config**

In this case, the ACLs are added to the existing configuration.

Deleting ACL entries using the entry sequence number

ACL entries can be deleted by specifying the sequence number only. In the following example, a filter rule is deleted by specifying its sequence number.

```
device(config)# show access-list name v4_acl
10: permit 1.1.1.1 0.0.0.0
20: permit 2.2.2.2 0.0.0.0
21: sequence 21 permit 3.3.3.3 0.0.0.0
30: deny any
device(config)# ip access-list standard v4_acl
device(config-std-nacl-v4_acl)# no sequence 20
device(config-std-nacl-v4_acl)# exit
device(config)# show access-list name v4_acl
10: permit 1.1.1.1 0.0.0.0
21: sequence 21 permit 3.3.3.3 0.0.0.0
30: deny any
```

Adding or deleting a comment

You can add or delete comments to an IP ACL entry.

Numbered ACLs: Adding a comment

To add a comment to an ACL entry in a numbered IPv4 ACL, perform the tasks listed below.

1. Use the **show access-list** to display the entries in an ACL.

   ```
   device(config-std-nacl)# show access-list 99
   Standard IP access-list 99
   deny host 10.2.4.5
   permit host 10.6.7.8
   ```

2. To add the comment "Permit all users" to filter "permit any" (the ACL remark is attached to the filter "permit any" as instructed in Step 4). Enter a command such as the following.

   ```
   device(config)# access-list 99 remark Permit all users
   ```

3. Entering a **show access-list** command displays the following:

   ```
   device(config-std-nacl)# show access-list 99
   Standard IP access-list 99
   deny host 10.2.4.5
   permit host 10.6.7.8
   ACL Remarks: Permit all users
   ```

4. Enter the filter "permit any".

   ```
   device (config-std-nacl)# permit any
   ```

5. Entering a **show access-list** command displays the following.

   ```
   device(config-std-nacl)# show access-list 99
   Standard IP access-list 99
   deny host 10.2.4.5
   permit host 10.6.7.8
   ACL Remarks: Permit all users
   permit any
   ```

   Syntax: [no] access-list acl-num remark comment-text

   Simply entering **access-list** *acl-num* **remark** *comment-text* adds a remark to the next ACL entry you create.

   The **remark** *comment-text* adds a comment to the ACL entry. The remark can have up to 128 characters in length. The comment must be entered separately from the actual ACL entry; that is, you cannot enter the ACL entry and the ACL comment with the same command. Also, in order for the remark to be displayed correctly in the output of **show** commands, the comment must be entered immediately before the ACL entry it describes.

   > **NOTE**
   > An ACL remark is attached to each individual filter only, not to the entire ACL (ACL 199).

   Complete the syntax by specifying any options you want for the ACL entry. Options you can use to configure standard or extended numbered ACLs are discussed in Named IPv4 ACLs on page 114.

Numbered ACLs: deleting a comment

For example, if the remark "Permit all users" has been defined for ACL 99, remove the remark by entering the following command.

```
device(config)# no access-list 99 remark Permit all users
```

Syntax: [no] access-list number remark comment-text

Named ACLs: adding a comment to a new ACL

You can add a comment to an ACL by performing the tasks listed below.

1. Use the **show access-list** command to display the contents of the ACL. For example, you may have an ACL named "entry" and a **show access-list** command shows that it has only one entry.

   ```
   device(config)# show access-list name entry
   Standard IP access-list 99
   deny host 10.2.4.5
   ```

2. Add a new entry with a remark to this named ACL by entering commands such as the following:

   ```
   device(config)#ip access-list standard entry
   device(config-std-nacl-entry)# remark Deny traffic from Marketing
   device(config-std-nacl-entry)# deny 10.6.7.8
   ```

3. Enter a **show access-list** command displays the new ACL entry with its remark.

   ```
   device(config)# show access-list name entry
   Standard IP access-list entry
   deny host 10.2.4.5
   ACL remark: Deny traffic from Marketing
   deny host 10.6.7.8
   ```

Syntax: [no] ip access-list standard | extended acl-name

Syntax: [no] remark string

Syntax: [no] deny options | permit options

The **standard | extended** parameter indicates the ACL type.

The *acl-name* parameter is the IPv4 ACL name. You can specify a string of up to 256 alphanumeric characters. You can use blanks in the ACL name if you enclose the name in quotation marks (for example, "ACL for Net1"). The *acl-num* parameter allows you to specify an ACL number if you prefer. If you specify a number, enter a number from 1 - 99 for standard ACLs or 100 - 199 for extended ACLs.

The **remark***string* adds a comment to the ACL entry that you are about to create. The comment can have up to 128 characters in length. The comment must be entered separately from the actual ACL entry; that is, you cannot enter the ACL entry and the ACL comment with the same command. Also, in order for the remark to be displayed correctly in the output of show commands, the comment must be entered immediately before the ACL entry it describes.

Enter **deny** to deny the specified traffic or **permit** to allow the specified traffic. Complete the configuration by specifying *options* for the standard or extended ACL entry. Options you can use to configure standard or extended named ACLs are discussed in the section Named IPv4 ACLs on page 114.

Named ACLs: deleting a comment

To delete a remark from a named ACL, enter the following command.

```
device(config)#ip access-list standard entry
device(config-std-nacl-entry)#no remark Deny traffic from Marketing
```

Syntax: no remark string

Applying ACLs to interfaces

Configuration examples in the section Numbered and named IPv4 ACLs on page 103 show that you apply ACLs to interfaces using the **ip access-group** command. This section present additional information about applying ACLs to interfaces.

Reapplying modified ACLs

If you make an ACL configuration change, you must reapply the ACLs to their interfaces to place the change into effect.

An ACL configuration change includes any of the following:

- Adding, changing, or removing an ACL or an entry in an ACL
- Changing ToS-based QoS mappings

To reapply ACLs following an ACL configuration change, enter the following command at the global CONFIG level of the CLI.

```
device(config)#
ip rebind-acl all
```

Syntax: [no] ip rebind-acl num | name | all

> NOTE
> When an ACL rebinds with a VE, the member ports are checked, and if there is any physical port associated with that VE, then the entire ACL entries are refreshed by an inbound ACL timer in the packet processor(PPCR).

Applying ACLs to a virtual routing interface

The virtual interface is used for routing between VLANs and contains all the ports within the VLAN. If the ACL is for the inbound traffic direction, you also can specify a subset of ports within the VLAN containing a specified virtual interface when assigning an ACL to that virtual interface. But if the ACL is for the outbound traffic direction, then it is not possible to specify a subset of ports within the VLAN on the devices.

Use this feature when you do not want the ACLs to apply to all the ports in the virtual interface's VLAN or when you want to streamline ACL performance for the VLAN.

To apply an ACL to a subset of ports within a virtual interface, enter commands such as the following.

```
device(config)# vlan 10 name IP-subnet-vlan
device(config-vlan-10)# untag ethernet 1/1 to 1/20 ethernet 2/1 to 2/12
device(config-vlan-10)# router-interface ve 1
device(config-vlan-10)# exit
device(config)# access-list 1 deny host 10.157.22.26
device(config)# access-list 1 deny 10.157.29.12
device(config)# access-list 1 deny host IPHost1
device(config)# access-list 1 permit any
device(config)# interface ve 1
device(config-vif-1)# ip access-group 1 in ethernet 1/1 ethernet 1/3 ethernet 2/1 to 2/4
```

The commands in this example configure port-based VLAN 10, add ports 1/1 - 2/12 to the VLAN, and add virtual routing interface 1 to the VLAN. The commands following the VLAN configuration commands configure ACL 1. Finally, the last two commands apply ACL 1 to a subset of the ports associated with virtual interface 1.

Syntax: [no] ip access-group num in [ethernet slot/portnum] [slot/portnum...] to slot/portnum

The **ethernet** *slot*/*portnum* option allow you to limit the ACL to a subset of ports within the virtual interface. You can also use the **to***slot*/*portnum* option to specify a range of ports. A maximum of 4 port ranges are supported.

Deletion of ACLs bound to an interface

To delete an ACL bound to an interface, use the **force-delete-bound-acl** command.

Initially **force-delete-bound-acl** is disabled.

```
device(config)#acl-policy
device(config-acl-policy)# force-delete-bound-acl
```

The **no force-delete-bound-acl** command does not allow the ACLs bound to an interface to be deleted.

```
device(config-acl-policy)# no force-delete-bound-acl
```

Syntax: [no] force-delete-bound-acl

When **force-delete-bound-acl** is enabled, it allows deletion of ACLs bound to one or more interfaces. After enabling this command for the deletion of the ACLs, however the binding of the ACL to an interface still remains. On rebinding this will be an empty ACL and will have no affect on traffic forwarding. On rebinding the CAM entries are reprogrammed appropriately, so no ACL filtering takes place after the ACL is deleted. This command is available as a sub-command of **acl-policy** command. However like any other ACL modification the CAM is only reprogrammed during rebind. Without a rebind the old filters are still present in the CAM.

> **NOTE**
> This command is also supported on CES 2000 Series and CER 2000 Series devices.

An example of the command is as below.

```
device(config-acl-policy)# force-delete-bound-acl
device(config-acl-policy)# exit
device(config)# show access-list all

ACL configuration:
!
mac access-list SampleACL
 permit any any 10 etype any
!
device(config)# show access-list bindings
L4 configuration:
!
interface ethe 2/1
 mac access-group SampleACL in
!
device(config)#show cam l2acl
  SLOT/PORT   Interface number
device(config)# sh cam l2acl 2/1
LP Index  VLAN Src MAC        Dest MAC        Port  Action   PRAM
   (Hex)                                                     (Hex)
2  0a3800 10   0000.0000.0000 0000.0000.0000 0     Pass     0009c
2  0a3802 0    0000.0000.0000 0000.0000.0000 0     Drop     0009d
device(config)#
device(config)#no mac acc SampleACL
device(config)#sh cam l2acl 2/1
LP Index  VLAN Src MAC        Dest MAC        Port  Action   PRAM
   (Hex)                                                     (Hex)
device(config)#show access-list all ACL configuration:
!
device(config)#show access-list bindings
L4 configuration:
!
!
interface ethe 2/1 mac access-group SampleACL in
!
device(config)#
```

> **NOTE**
> Rebinding of an ACL is explicitly required for IPv4 and IPv6 ACLs.

Enabling ACL duplication check

If desired, you can enable software checking for duplicate ACL entries. To do so, enter the following command at the **config-acl-policy** level of the CLI.

```
device(config)# acl-policy
device(config-acl-policy)# acl-duplication-check
device(config-acl-policy)# access-list 173 permit ip host 1.1.6.197 198.6.1.0 0.0.0.255
device(config-acl-policy)# access-list 173 permit ip host 1.1.6.197 198.6.1.0 0.0.0.255
Error: Duplicate entry in ACL 173
permit ip host 1.1.6.197 198.6.1.0 0.0.0.255
```

The above example generates an error message from the system as access-list 173 has a duplicate entry. For **no** command, enter the following command at the **config-acl-policy** level of the CLI.

```
device(config)# acl-policy
device(config-acl-policy)# no acl-duplication-check
device(config-acl-policy)# access-list 173 permit ip host 1.1.6.195 198.6.1.0 0.0.0.255
device(config-acl-policy)# access-list 173 permit ip host 1.1.6.195 198.6.1.0 0.0.0.255
device(config-acl-policy)# sh acc 173
Extended IP access list 173
    0: permit ip host 1.1.6.195 198.6.1.0 0.0.0.255
    1: permit ip host 1.1.6.195 198.6.1.0 0.0.0.255
```

Syntax: [no] acl-duplication-check

Enabling ACL conflict check

If desired, you can enable software checking for conflicting ACL entries. To do so, enter the following command at the **config-acl-policy** level of the CLI.

```
device(config)# acl-policy
device(config-acl-policy)# acl-conflict-check
device(config-acl-policy)# access-list 173 permit ip host 1.1.6.203 198.6.1.0 0.0.0.255
device(config-acl-policy)# access-list 173 deny ip host 1.1.6.203 198.6.1.0 0.0.0.255
Warning: Conflicting entry in ACL 173: permit ip host 1.1.6.203 198.6.1.0 0.0.0.255
device(config-acl-policy)# acc 174 deny ip host 1.1.6.203 198.6.1.0 0.0.0.255
device(config-acl-policy)#
```

The above example generates an error message from the system as access-list 173 has a conflicting entry. For **no** command, enter the following command at the **config-acl-policy** level of the CLI.

```
device(config)# acl-policy
device(config-acl-policy)# no acl-conflict-check
device(config-acl-policy)# access-list 173 permit ip host 1.1.6.201 198.6.1.0 0.0.0.255
device(config-acl-policy)# access-list 173 deny ip host 1.1.6.201 198.6.1.0 0.0.0.255
device(config-acl-policy)#
```

Syntax: [no] acl-conflict-check

> NOTE
> This command checks for conflicts across multiple ACL clauses, except for the permit or deny keyword.

Enabling ACL filtering of fragmented or non-fragmented packets

To define an extended IPv4 ACL to deny or permit traffic with fragmented or unfragmented packets, enter a command such as those shown in one of the methods below.

Numbered ACLs

```
device(config)# access-list 111 deny ip any any fragment
device(config)# int eth 1/1
device(config-if-e10000-1/1)# ip access-group 111 in
device(config)# write memory
```

The first line in the example defines ACL 111 to deny any fragmented packets. Other packets will be denied or permitted, based on the next filter condition.

Next, after assigning the ACL to Access Group 111, the access group is bound to port 1/1. It will be used to filter incoming traffic.

Refer to Extended numbered IPv4 ACL syntax on page 109 for the complete syntax for extended ACLs.

Named ACLs

```
device(config)# ip access-list extended entry
device(config-ext-nacl)# deny ip any any fragment
device(config)# int eth 1/1
device(config-if-e10000-1/1)# ip access-group entry in
device(config)# write memory
```

The first line in the example defines ACL entry to deny any fragmented packets. Other packets will be denied or permitted, based on the next filter condition.

Next, after assigning the ACL to Access Group entry, the access group is bound to port 1/1. It will be used to filter incoming traffic.

Syntax: ip access-list extended acl–name | **acl–num deny** | **permit ip–protocol source–ip** | **hostname wildcard** [**operator source–tcp/ udp–port**] **destination–ip** | **hostname** [**icmp–type** | **num**] **wildcard** [**operator destination–tcp/udp–port**] [**precedence name** | **num**] [**tos name** | **num**] [**fragment**] | [**non–fragmented**]

Enter **extended** to indicate the named ACL is an extended ACL.

The *acl–name* | *acl–num* parameter allows you to specify an IPv4 ACL name or number. If using a name, specify a string of up to 256 alphanumeric characters. You can use blanks in the ACL name, if you enclose the name in quotation marks (for example, "ACL for Net1"). The *acl–num* parameter allows you to specify an ACL number if you prefer. If you specify a number, enter a number from 100 - 199 for extended ACLs.

Enter the **fragment** keyword to allow the ACL to filter fragmented packets. Use the **non–fragmented** keyword to filter non-fragmented packets.

> NOTE
> The **fragmented** and **non–fragmented** parameters cannot be used together in an ACL entry.

Complete the configuration by specifying options for the ACL entry. Options you can use are discussed in the appropriate sections for configuring ACLs in this chapter.

Configuring the conservative ACL fragment mode

The **acl-frag-conservative** command enables you to change the operation of ACLs on fragmented packets.

> NOTE
> This command is supported for outbound ACLs only when the outbound physical port is on a 10Gx8 module. Furthermore, use of **fragment**-keyword ACL lines is not supported for any module type other than 10Gx8 when **acl-frag-conservative** is configured. When it is necessary to handle fragments in outbound ACLs against traffic sent out of any module type other than 10Gx8, **acl-frag-conservative** should not be configured and **fragment**-keyword ACL lines should be used instead.

When a packet exceeds the maximum packet size, the packet is fragmented into a number of smaller packets that contain portions of the contents of the original packet. This packet flow begins with an initial packet that contains all of the Layer-3 and Layer-4 header information contained in the original packet and is following by a number of packets that contain only the Layer-3 header information. This packet flow contains all of the information contained in the original packet distributed through the packet flow into packets that are small enough to avoid the maximum packet size limit. This provides a particular problem for ACL processing. If the ACL is filtering based on Layer-4 information, the non-initial packets within the fragmented packet flow will not match the Layer-4 information even if the original packet that was fragmented would have matched the filter. Consequently, packets that the ACL was designed to filter for are not processed by the ACL.

This can be a particular problem for Deny ACLs because packets can be dropped that should be forwarded. For this reason, the conservative ACL fragment mode has been created to treat fragmented packets differently both when the **fragment** keyword is and is not used. While under normal operation, fragmented packets are treated the same as all other packets, when the **acl-frag-conservative** command is enabled, the device only applies Layer-4 information within an ACL to non-fragmented packets and to the initial packet within a fragmented packet flow.

Layer 4 filters in a Layer 3 ACL

An ACL entry with one or more of the following keywords is consider to have Layer-4 information:

- TCP, UDP, or SCTP source or destination port.
- ICMP—matching based on ICMP type or code values
- TCP SYN flag
- TCP Established flag

ACL operation with the device configured in conservative ACL fragment mode

Operation of ACLs with Fragmented packets when the device is configured in Conservative ACL Fragment mode, through use of the **acl-frag-conservative** command, can be described to follow one these four procedures:

- ACL entries with Layer-3 Information only that do not contain the **fragment** keyword.
- ACL entries with Layer-3 Information only that do contain the **fragment** keyword.
- ACL entries with Layer-3 and Layer-4 Information that do not contain the **fragment** keyword.
- ACL entries with Layer-3 and Layer-4 Information that do contain the **fragment** keyword.

Detailed operation of ACLs under each of these conditions are described as follows.

ACL entries with Layer-3 information only that do not contain the fragment keyword

In this situation, the operation of the ACL is exactly like it is during normal operation (**acl-frag-conservative** command not configured). Any packet or fragment that matches the Layer-3 information specified in the ACL will be matched as described in Table 8.

TABLE 8 ACL entry with Layer-3 information only and no fragment keyword

	Packet matches AND is either a non-fragmented or the 1st packet within a fragmented packet flow	Packet matches AND is a non-initial packet within a fragmented packet flow
permit	Yes - Matches because the packet matches the Layer-3 Information in the ACL.	Yes - Matches because the packet matches the Layer-3 Information in the ACL.
deny	Yes - Matches because the packet matches the Layer-3 Information in the ACL.	Yes - Matches because the packet matches the Layer-3 Information in the ACL.

ACL entries with Layer-3 information only that do contain the fragment keyword

In this situation, any packet that is not fragmented will not match because the fragment keyword is configured in the ACL. Non-initial packets within a fragmented packet flow that contain the Layer-3 information specified in the ACL will be matched as described in Table 9.

TABLE 9 ACL entry with Layer-3 information only and fragment keyword in ACL

	Packet matches AND is either a non-fragmented or the 1st packet within a fragmented packet flow	Packet matches AND is a non-initial packet within a fragmented packet flow
permit	No - Does not match because fragment keyword is in ACL and packet is either non-fragmented or the 1st packet within a fragmented packet flow.	Yes - Matches because fragment keyword is in ACL and packet is a non-initial packet within a fragmented packet flow and the packet matches the Layer-3 information in the ACL.
deny	No - Does not match because fragment keyword is in ACL and packet is either non-fragmented or the 1st packet within a fragmented packet flow.	Yes - Matches because fragment keyword is in ACL and packet is a non-initial packet within a fragmented packet flow and the packet matches the Layer-3 information in the ACL.

ACL entries with Layer-3 and Layer-4 information that do not contain the fragment keyword

In this situation, any packet that is not fragmented or is the 1st packet within a fragmented packet flow and also contains the Layer-3 and Layer-4 information specified in the ACL will be matched. Packets that are non-initial packets within a fragmented packet flow and match the Layer-3 information will be matched for the permit clause because in conservative ACL fragment mode, Layer-4 information is disregarded for non-initial packets. Also, non-initial packets within a fragmented packet flow will not be matched for the deny clause because in conservative ACL fragment mode, the deny clause is not invoked for non-initial packets within a fragmented packet flow. Refer to Table 10 for operation in this scenario.

TABLE 10 ACL entry with Layer-3 and Layer-4 information and no fragment keyword in ACL

	Packet matches AND is either a non-fragmented or the 1st packet within a fragmented packet flow	Packet matches AND is a non-initial packet within a fragmented packet flow
permit	Yes - Matches because the packet matches the Layer-3 and Layer-4 Information in the ACL.	Yes - Matches because the packet matches the Layer-3 Information in the ACL and in conservative mode, Layer-4 information is disregarded for non-initial packets within a fragmented packet flow.
deny	Yes - Matches because the packet matches the Layer-3 and Layer-4 Information in the ACL.	No - Does not match because in conservative mode, the deny clause is not invoked for non-initial packets within a fragmented packet flow.

ACL entries with Layer-3 and Layer-4 information that contains the fragment keyword

In this situation, any packet that is not fragmented or is the 1st packet within a fragmented packet flow will not be matched because the fragment keyword is specified in the ACL. Packets that are non-initial packets within a fragmented packet flow, match the **fragment** keyword and match the Layer-3 information will be matched for the permit clause because in conservative ACL fragment mode, Layer-4 information is disregarded for non-initial packets. Also, non-initial packets within a fragmented packet flow will not be matched for the deny clause because in conservative ACL fragment mode, the **deny** clause is not invoked for non-initial packets within a fragmented packet flow. Refer to Table 11 for operation in this scenario.

TABLE 11 ACL entry with Layer-3 and Layer-4 information and fragment keyword in ACL

	Packet matches AND is either a non-fragmented or the 1st packet within a fragmented packet flow	Packet matches AND is a non-initial packet within a fragmented packet flow
permit	No - Does not match because fragment keyword is in ACL and packet is either non-fragmented or the 1st packet within a fragmented packet flow.	Yes - Matches because the packet matches the Layer-3 Information in the ACL and in conservative mode, Layer-4 information is disregarded for non-initial packets within a fragmented packet flow.
deny	No - Does not match because fragment keyword is in ACL and packet is either non-fragmented or the 1st packet within a fragmented packet flow.	No - Does not match because in conservative mode, the deny clause is not invoked for non-initial packets within a fragmented packet flow.

Configuring the conservative ACL fragment mode

The Conservative ACL Fragment mode is configured using the **acl-frag-conservative** command as shown in the following.

```
device(config)# acl-policy
device(config-acl-policy)# acl-frag-conservative
```

Syntax: [no] acl-frag-conservative

Examples of ACL filtering in normal and conservative ACL fragment modes

The following examples illustrate how an ACL with the fragment keyword operates for filtering applications in both the normal and conservative mode:

- ACL Configuration Example with Fragment Keyword and Permit Clause
- ACL Configuration Example with Fragment Keyword and Deny Clause

ACL configuration example with fragment keyword and permit clause

In the following example, ACL 100 is configured to process fragmented IP packets in Normal and Conservative ACL modes as described.

```
device(config)# access-list 100 permit tcp 10.1.0.0.0.0.0.255 any fragment
device(config)# access-list 100 deny ip any any
```

Behavior In Normal ACL Fragment Mode - In the normal Extreme device mode, fragmented and non-fragmented packets will be dropped or forwarded as described in the following:

All TCP fragments (both initial and subsequent fragments) from the specified IP address, will match the first ACL entry. Because this is a **permit** ACL entry, the matching packets are forwarded.

Non-fragmented packets will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**deny**) ACL entry and consequently will be dropped.

Behavior In Conservative ACL Fragment Mode - If the Extreme device is configured for Conservative ACL Fragment mode using the **acl-frag-conservative** command, fragmented and non-fragmented packets will be dropped or forwarded as described in the following:

The initial fragment will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**deny**) ACL entry and consequently will be dropped.

Non-initial TCP fragments from the specified IP address, will match the first ACL entry based on Layer-3 information. Because this is a **permit** ACL entry, the matching packets are forwarded.

Non-fragmented packets will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**deny**) ACL entry and consequently will be dropped.

ACL configuration example with fragment keyword and deny clause

In the following example, ACL 101 is configured to process fragmented IP packets in Normal and Conservative ACL modes as described.

```
device(config)# access-list 101 deny tcp 10.1.0.0.0.0.0.255 any fragment
device(config)# access-list 101 permit ip any any
```

Behavior In Normal ACL Fragment Mode - In the normal Extreme device mode, fragmented and non-fragmented packets will be dropped or forwarded as described in the following:

All TCP fragments (both initial and subsequent fragments) from the specified IP address will match the first ACL entry. Because this is a **deny** ACL entry, the matching packets are dropped.

Non-fragmented packets will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**permit**) ACL entry and consequently will be forwarded.

Behavior In Conservative ACL Fragment Mode - If the Extreme device is configured for Conservative ACL Fragment mode using the **acl-frag-conservative** command, fragmented and non-fragmented packets will be dropped or forwarded as described in the following:

The initial fragment will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**permit**) ACL entry and consequently will be forwarded.

Non-initial TCP fragments will match the first ACL entry based on Layer-3 information. Because this is a **deny** ACL entry with Layer-3 information only, the matching packets are dropped.

Non-fragmented packets will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**permit**) ACL entry and consequently will be forwarded.

Examples of ACL-based rate limiting in normal and conservative ACL fragment modes

The following examples illustrate how an ACL with the fragment keyword operates for rate limiting applications in both the normal and conservative mode:

- ACL-based Rate Limiting Configuration Example with Fragment Keyword and Deny Clause
- ACL-based Rate Limiting Configuration Example with Fragment Keyword and Permit Clause

ACL-based rate limiting configuration example with fragment keyword and deny clause

In the following example, ACL 102 is configured to process fragmented IP packets in Normal and Conservative ACL modes as described.

```
device(config)# interface ethernet 3/1
device(config-if-e1000-3/1)# enable
device(config-if-e1000-3/1)# rate-limit strict-acl
device(config-if-e1000-3/1)# rate-limit input access-group 102 499992736 750000000
device(config-if-e1000-3/1)# no spanning-tree
device(config-if-e1000-3/1)# exit
device(config)# access-list 102 deny ip any any fragment
device(config)# access-list 102 permit ip any any
```

Behavior In Normal ACL Fragment Mode - In the normal Extreme device mode, fragmented and non-fragmented packets will be dropped or forwarded as described in the following:

All IP fragments (both initial and subsequent fragments) will match the first ACL entry. Because this is a **deny** ACL entry, and **rate-limit strict-acl** is configured, the matching packets are dropped.

Non-fragmented packets will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**permit**) ACL entry and consequently will be forwarded and rate-limited.

Behavior In Conservative ACL Fragment Mode - If the Extreme device is configured for Conservative ACL Fragment mode using the **acl-frag-conservative** command, fragmented and non-fragmented packets will be dropped or forwarded as described in the following:

The initial fragment will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**permit**) ACL entry and consequently will be forwarded and rate-limited.

Non-initial IP fragments will match the first ACL entry based on Layer-3 information. Because this is a **deny** ACL entry with Layer-3 information only, and **rate-limit strict-acl** is configured, the matching packets are dropped.

Non-fragmented packets will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**permit**) ACL entry and consequently will be forwarded and rate-limited.

ACL-based rate limiting configuration example with fragment keyword and permit clause

In the following example, ACL 103 is configured to process fragmented IP packets in Normal and Conservative ACL modes as described.

```
device(config)# interface ethernet 3/1
device(config-if-e1000-3/1)# enable
device(config-if-e1000-3/1)# rate-limit strict-acl
device(config-if-e1000-3/1)# rate-limit input access-group 103 499992736 750000000
device(config-if-e1000-3/1)# no spanning-tree
device(config-if-e1000-3/1)# exit
device(config)# access-list 103 permit ip any any fragment
device(config)# access-list 102 deny ip any any
```

Behavior In Normal ACL Fragment Mode - In the normal Extreme device mode, fragmented and non-fragmented packets will be dropped or forwarded as described in the following:

All IP fragments (both initial and subsequent fragments) will match the first ACL entry. Because this is a **permit** ACL entry, the matching packets are forwarded and rate-limited.

Non-fragmented packets will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**deny**) ACL entry and consequently will be dropped.

Behavior In Conservative ACL Fragment Mode - If the Extreme device is configured for Conservative ACL Fragment mode using the **acl-frag-conservative** command, fragmented and non-fragmented packets will be dropped or forwarded as described in the following:

The initial fragment will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**deny**) ACL entry and consequently will be dropped.

Non-initial IP fragments will match the first ACL entry based on L3 information. Because this is a **permit** ACL entry, the matching packets are forwarded and rate-limited.

Non-fragmented packets will not match the first ACL entry because the **fragment** keyword is present. The packet will then match the second (**deny**) ACL entry and consequently will be dropped.

ACL filtering for traffic switched within a virtual routing interface

By default, an Extreme device does not filter traffic that is switched from one port to another within the same virtual routing interface, even if an ACL is applied to the interface. You can enable the Extreme device to filter switched traffic within a virtual routing interface. When you enable the filtering, the Extreme device uses the ACLs applied to inbound traffic to filter traffic received by a port from another port in the same virtual routing interface. This feature does not apply to ACLs applied to outbound traffic.

To enable filtering of traffic switched within a virtual routing interface, enter the following command at the configuration level for the interface.

```
device(config-vif-1)# ip access-group ve-traffic
```

Syntax: [no] ip access-group ve-traffic

Filtering and priority manipulation based on 802.1p priority

Filtering and priority manipulation based on a packet's 801.1p priority is supported in the Extreme devices through the following QoS options:

- **priority** - Assigns traffic that matches the ACL to a hardware forwarding queue. In addition to changing the internal forwarding priority, if the outgoing interface is an 802.1q interface, this option maps the specified priority to its equivalent 802.1p (QoS) priority and marks the packet with the new 802.1p priority.

- **priority-force** - Assigns packets of outgoing traffic that match the ACL to a specific hardware forwarding queue, even though the incoming packet may be assigned to another queue. Specify one of the following QoS queues:

    - 0 - qosp0
    - 1 - qosp1
    - 2 - qosp2
    - 3 - qosp3
    - 4 - qosp4
    - 5 - qosp5
    - 6 - qosp6
    - 7 - qosp7

If a packet's 802.1p value is forced to another value by its assignment to a lower value queue, it will retain that value when it is sent out through the outbound port.

The default behavior on previous revisions of this feature was to send the packet out with the higher of two possible values: the initial 802.1p value that the packet arrived with or the new (higher) priority that the packet has been "forced" to.

- **priority-mapping** - Matches on the packet's 802.1p value. This option does not change the packet's forwarding priority through the device or mark the packet.

- **drop-precedence** - Assigns traffic that matches the ACL to a drop precedence value between 0 -3.

drop-precedence-force - This keyword applies in situations where there are conflicting priority values for packets on an Ingress port, that conflict can be resolved by performing a priority merge (the default) or by using a **force** command to direct the device to use a particular value above other values. The **drop-precedence-force** keyword specifies that if a drop precedence is applied on the port the ACL keyword will override existing or default mappings, however, if forced at the ingress port, the port value will prevail over the acl value. Assigns traffic that matches the ACL to a drop precedence value between 0 -3.

Example using the priority option (IPv4)

In the following IPv4 example, access list 100 assigns TCP packets with the source and destination addresses specified to internal priority 2 and maps them to the 802.1p value 2 when outbound.

```
device(config)#access-list 100 permit tcp 10.1.1.0/24 10.23.45.0/24 priority 2
```

The **priority** parameter specifies one of the 8 internal priorities of the Extreme device. Possible values are between 0 and 7. If the outgoing interface is an 802.1q interface, the packet will have its 802.1p (QoS) priority marked with the new priority defined in this ACL.

Example using the priority force option

In the following IPv4 ACL example, access list 100 assigns UDP packets with the source and destination addresses specified to the internal priority 3.

```
device(config)#access-list 100 permit udp 10.1.1.0/24 10.23.45.0/24 priority-force 3
```

The **priority-force** parameter specifies one of the 8 internal priorities of the Extreme device. Possible values are between 0 and 7.

For limitations when using the **priority-force** parameter, please see Configuration considerations for IPv4 outbound ACLs on VPLS, VLL, and VLL-Local endpoints on page 100.

Example using the priority mapping option

In the following IPv4 ACL example, access list 100 permits UDP packets with the source and destination addresses specified and the 802.1p priority 7.

```
device(config)# access-list 100 permit udp 10.1.1.0/24 10.75.34.0/24 priority-mapping 7
```

The **priority-mapping** parameter specifies one of the eight possible 802.1p priority values. Possible values are between 0 and 7.

> **NOTE**
> When the priority configured for a physical port and the 802.1p priority of an arriving packet differ, the higher of the two priorities is used.

ICMP filtering for extended ACLs

Extended IPv4 ACL policies can be created to filter traffic based on its ICMP message type. You can either enter the description of the message type or enter its type and code IDs. All packets matching the defined ICMP message type or type number and code number are processed in hardware.

Numbered ACLs

For example, to deny the echo message type in a numbered, extended ACL, enter commands such as the following when configuring a numbered ACL.

```
device(config)# access-list 109 deny icmp any any echo
```

or

```
device(config)# access-list 109 deny icmp any any 8 0
```

Syntax: [no] access-list num deny | permit [vlan *vlan_id*] icmp any any icmp-type | type-number code-number

The **deny | permit** parameter indicates whether packets that match the policy are dropped or forwarded.

You can either enter the name of the message type for *icmp-type* or the message's *type number* and *code number* of the message type. Refer to the table below for valid values.

Named ACLs

For example, to deny the administratively-prohibited message type in a named ACL, enter commands such as the following.

```
device(config)# ip access-list extended entry
device(config-ext-nacl)# deny ICMP any any administratively-prohibited
```

or

```
device(config)# ip access-list extended entry
device(config-ext-nacl)#deny ICMP any any 3 13
```

Syntax: [no] ip access-list extended acl-name deny | permit host icmp any any icmp-type | type-number code-number

The **extended** parameter indicates the ACL entry is an extended ACL.

The *acl-name | acl-num* parameter allows you to specify an ACL name or number. If using a name, specify a string of up to 256 alphanumeric characters. You can use blanks in the ACL name if you enclose the name in quotation marks (for example, "ACL for Net1"). The *acl-num* parameter allows you to specify an ACL number if you prefer. If you specify a number, enter a number from 100 - 199 for extended ACLs.

The **deny | permit** parameter indicates whether packets that match the policy are dropped or forwarded.

You can either use the *icmp-type* and enter the name of the message type or use the *type-number* parameter to enter the type number and code number of the message. Refer to the table below for valid values.

TABLE 12

ICMP message type	Type	Code
administratively-prohibited	3	13
any-icmp-type	x	x
destination-host-prohibited	3	10
destination-host-unknown	3	7
destination-net-prohibited	3	9
destination-network-unknown	3	6
echo	8	0
echo-reply	0	0
general-parameter-problem NOTE This message type indicates that required option is missing.	12	1
host-precedence-violation	3	14
host-redirect	5	1
host-tos-redirect	5	3
host-tos-unreachable	3	12
host-unreachable	3	1
information-reply	16	0
information-request	15	0
mask-reply	18	0

TABLE 12 (continued)

ICMP message type	Type	Code
mask-request	17	0
net-redirect	5	0
net-tos-redirect	5	2
net-tos-unreachable	3	11
net-unreachable	3	0
packet-too-big	3	4
parameter-problem **NOTE** This message includes all parameter problems	12	0
port-unreachable	3	3
precedence-cutoff	3	15
protocol-unreachable	3	2
reassembly-timeout	11	1
redirect **NOTE** This includes all redirects.This option is not available in CES 2000 Series or CER 2000 Series devices.	5	x
router-advertisement	9	0
router-solicitation	10	0
source-host-isolated	3	8
source-quench	4	0
source-route-failed	3	5
time-exceeded **NOTE** This option is not available in CES 2000 Series or CER 2000 Series devices.	11	x
timestamp-reply	14	0
timestamp-request	13	0
ttl-exceeded	11	0
unreachable **NOTE** This includes all unreachable messages. This option is not available in CES 2000 Series or CER 2000 Series devices.	3	x

Binding IPv4 inbound ACLs to a management port

You can bind a small number of IPv4 inbound ACLs to the Ethernet port on the Management Module for filtering IP traffic sent to the Management module's CPU. These ACLs are processed in software only and are not programmed in CAM. Outbound IPv4 ACLs are not supported on the Management module's Ethernet port.

The default size of IPv4 Inbound ACLs on a management port is 20 filters. This number can be set from 1 to 100 using the following command.

```
device(config)# system-max mgmt-port-acl-size 100
```

Syntax: system mgmt-port-acl-size acls-supported

The *acls-supported* variable allows you set a maximum number of filters that are supported for the IPv4 ACL bound to the Management Module's Ethernet port.

The possible values are 1 - 100.

The default value is 20.

> **NOTE**
> For IPv4 inbound ACL applied to management port, the user can log traffic matching both "permit" and "deny" ACL filters that have the log keyword. The command **ip access-group enable-deny-logging** is not be required to turn on logging on a management port.

> **NOTE**
> On CES 2000 Series or CER 2000 Series devices you can bind an ACL with accounting clauses to the management port. However, no ACL counters will be incremented by packets permitted or denied by those clauses.

IP broadcast ACL

The IP broadcast Access Control List (ACL) enables filtering of IP subnet-based directed broadcast traffic. The IP broadcast ACL is configured by creating an ACL (standard or extended) and then binding that ACL to the IP interface on the router for which filtering needs to be enabled. The IP broadcast ACLs identify directed broadcast traffic based on the subnets configured on the interfaces, and filter all the traffic for the respective VRF of an interface. An ACL entry is programmed in CAM for each interface. Thereby, the need to add a filter for each trusted source and destination subnet combination is eliminated.

As an example, suppose you define the standard ACL clause access-list 1 permit host 10.1.5.1 and bind the ACL to the IP interface on the router using the **ip subnet-broadcast-acl** command. Multiple ACL CAM entries are programmed for such a binding, as shown in the following example.

For example, a router has the following three interface IP addresses configured in the same VRF:

- 2.2.2.2/24
- 10.10.10.1/24
- 10.10.20.1/24

The ACL CAM is then programmed with the following three entries:

- permit host 10.1.5.1 host 10.2.2.255
- permit host 10.1.5.1 host 10.10.10.255
- permit host 10.1.5.1 host 10.10.20.255

The ACL CAM is then implicitly programmed with the following three deny any entries:

- deny host any host 10.2.2.255

- deny host any host 10.10.10.255
- deny host any host 10.10.20.255

Configuration considerations for IP broadcast ACL

The configuration considerations for binding an IP directed-broadcast ACL to an interface are as follows:

- If a physical port is a member of a virtual interface, then ACL binding is permitted only at the VE level and not at the physical port level.
- For LAG ports, all ports within the LAG are required to have the same IP broadcast ACL applied to them before the LAG is created. On deleting the LAG, the IP broadcast ACL binding is replicated on all individual LAG ports.
- IP directed-broadcast ACL binding is not be permitted on VPLS and VLL endpoints.
- For interface-level inbound IPv4 ACL or Rate Limiting-ACLs (RL-ACLs) - Traffic matching IP broadcast ACLs is not subject to interface-level ACLs or RL-ACLs. You must configure an IP broadcast ACL so that only directed broadcast traffic matches the IP broadcast ACL clauses.
- For interface-level inbound Layer 2 ACLs or RL-ACLs - For XMR Series and MLX Series devices, either an IPv4 inbound or Layer 2 inbound ACL can be configured on an interface, but not both. But for CER 2000 Series and CES 2000 Series devices, both the IPv4 inbound and Layer 2 inbound ACL can be configured on an interface.
- IP broadcast ACLs do not support ACL-based logging, Sample Flow (sFlow), and mirroring features.
- Traffic matching IP broadcast ACLs is not subject to policy-based routing.

Configuring IP broadcast ACL and establishing the sequence of IP broadcast ACL commands

You can enable filtering of directed broadcast traffic using ACLs at the IP interface level and the global configuration level, with the interface-level command taking precedence over the global configuration level command.

To enable filtering of directed broadcast traffic using ACLs globally, enter the following commands.

```
device(config)# access-list 5 permit host 10.1.1.2
device(config)# ip global-subnet-broadcast-acl 5
```

Syntax: [no] ip global-subnet-broadcast-acl *acl-num*

The *acl-num* parameter can be a standard or extended access list number. Enter a number from 1 through 99 for a standard ACL, and a number from 100 through 199 for an extended ACL.

The **no** option is used to disable filtering of directed broadcast traffic globally.

> NOTE
> The binding of global subnet broadcast ACLs filter only the traffic belonging to default VRF interfaces.

> NOTE
> Only numbered IPv4 ACLs are supported.

To enable filtering of directed broadcast traffic on an individual interface, enter the following commands.

```
device(config)# access-list 5 permit host 10.1.1.2
device(config)# interface ethernet 2/1
device(config-if-e10000-2/1)# ip subnet-broadcast-acl 5
```

Syntax: [no] ip subnet-broadcast-acl *acl-num*

The *acl-num* parameter can be a standard or extended access list number. Enter a number from 1 through 99 for a standard ACL, and a number from 100 through 199 for an extended ACL.

The **no** option is used to disable filtering of directed broadcast traffic on an individual interface.

> **NOTE**
> IP tunnel interfaces are not supported.

> **NOTE**
> Upon dynamically configuring or removing the **ip broadcast-zero** command, you must manually rebind the subnet broadcast ACLs.

Configuration example for IP broadcast ACL

Figure 1 illustrates how filtering of IP directed broadcast traffic is enabled on the Router 3 interface. For example, to enable filtering of IP directed broadcast traffic on the Router 3 interface 1/2, you must configure an ACL with source IP address 10.1.1.2 and ACL action **permit**, and then bind that ACL to interface 1/2 on Router 3 as the IP broadcast ACL. As a result of enabling the IP broadcast ACL filter on the Router 3 interface, Router 3 allows the IP broadcast packet from the source IP address 10.1.1.2 and drops the IP broadcast packet from any other source on port 1/2.

FIGURE 1 Filtering of IP directed broadcast traffic on the Router 3 interface
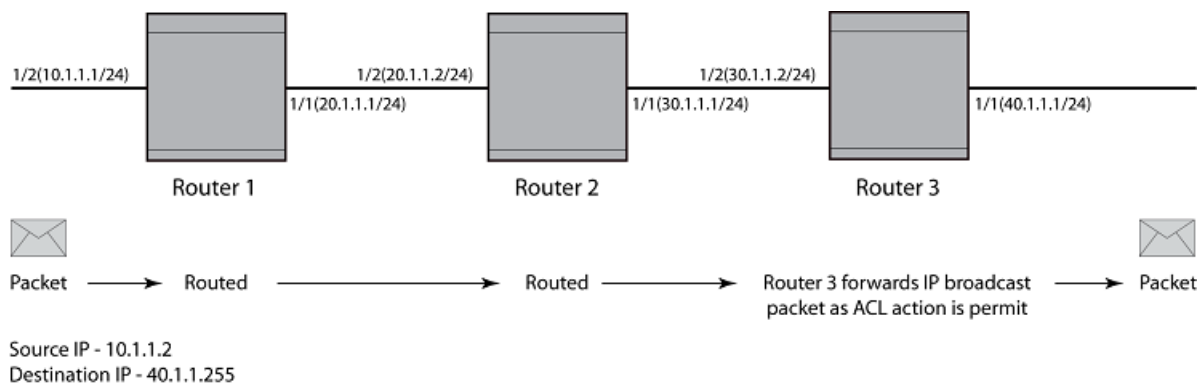
To configure an IP broadcast ACL on Router 3 interface 1/2, enter the following commands.

```
device(config)# access-list 1 permit host 10.1.1.2
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# enable
device(config-if-e10000-1/1)# ip address 10.1.1.1/24
device(config-if-e10000-1/1)# exit
device(config)# interface ethernet 1/2
device(config-if-e10000-1/2)# enable
device(config-if-e10000-1/2)# ip address 10.1.1.1/24
device(config-if-e10000-1/2)# ip subnet-broadcast-acl 1
device(config-if-e10000-1/2)# ip directed-broadcast
device(config-if-e10000-1/2)# exit
```

Displaying accounting information for IP broadcast ACL

To display the accounting information for an IP broadcast ACL at the IP interface level, enter the following command.

```
device(config-if-e1000-4/1)# show access-list subnet-broadcast accounting ethernet 4/1
Subnet broadcast ACL 120
    0: permit udp host 10.10.10.1 host 10.20.20.255
        Hit count: (1 sec)                    0    (1 min)                    0
```

```
          (5 min)                          0   (accum)                          0
    1: permit tcp host 10.10.10.1 host 10.20.20.255
       Hit count: (1 sec)                 10   (1 min)                         67
                  (5 min)                  0   (accum)                         67
    2: deny ip any any
       Hit count: (1 sec)                  0   (1 min)                          0
                  (5 min)                  0   (accum)                          0
```

Syntax: show access-list subnet-broadcast accounting [ethernet | ve] *port-id orvid*

The **ethernet**, and **ve** options specify the interfaces for which you can display the accounting information. If you specify an Ethernet interface, you must also specify the port number associated with the interface. If you specify a VE interface, you must specify the VE number associated with the interface.

The *port-id* parameter specifies the port for which you want to display the accounting information.

The *vid* parameter specifies the VE interface ID.

The table below describes the output parameters of the **show access-list subnet-broadcast accounting** command.

TABLE 13

Field	Description
Subnet broadcast ACL ID	The ID of the IP broadcast ACL.
#	The index of the IP broadcast ACL entry, starting with 0, followed by the permit or deny condition defined for that ACL entry. (The first entry created for an ACL is assigned the index 0. The index of the subsequent entries created are incremented by 1.)
permit udp host	The UDP packets are permitted from a specific source address to a specific destination address.
permit tcp host	The TCP packets are permitted from a specific source address to a specific destination address.
deny ip any	The IP packets are denied for all the host addresses.
Hit count	The number of hits for each counter.

To display the accounting information for an IP broadcast ACL globally, enter the following command.

```
device# show access-list subnet-broadcast accounting global
Subnet broadcast ACL 12
    0: permit enable-accounting host 10.1.103.217
       Hit count: (1 sec)                  2   (1 min)                        150
                  (5 min)                  0   (accum)                        384
    1: deny enable-accounting host 172.24.103.217
       Hit count: (1 sec)                  4   (1 min)                        298
                  (5 min)                  0   (accum)                        764
    2: permit enable-accounting host 10.100.103.217
       Hit count: (1 sec)                 10   (1 min)                        600
                  (5 min)                  0   (accum)                       1540
```

Syntax: show access-list subnet-broadcast accounting global

The table below describes the output parameters of the **show access-list subnet-broadcast accounting global** command.

TABLE 14

Field	Description
Subnet broadcast ACL ID	The ID of the IP broadcast ACL.
#	The index of the IP broadcast ACL entry, starting with 0, followed by the permit or deny condition defined for that ACL entry. (The first entry created for an ACL is assigned the index 0. The index of the subsequent entries created are incremented by 1.)

TABLE 14 (continued)

Field	Description
permit enable-accounting host	The ACL accounting is enabled for a specific host IP address.
deny enable-accounting host	The ACL accounting is disabled for a specific host IP address.
Hit count	The number of hits for each counter.

Clearing accounting information for IP broadcast ACL

To clear the accounting information for an IP broadcast ACL at the IP interface level, enter the following command.

```
device# clear access-list subnet-broadcast accounting ve 10
```

Syntax: clear access-list subnet-broadcast accounting [**ethernet** | **ve**] *port-id orvid*

The **ethernet**, and **ve** options specify the interfaces for which you can remove the accounting information. If you specify an Ethernet interface, you must also specify the port number associated with the interface. If you specify a VE interface, you must specify the VE number associated with the interface.

The *port-id* parameter specifies the port for which you want to clear the accounting information.

The *vid* parameter specifies the VE interface ID.

To clear the accounting information for an IP broadcast ACL globally, enter the following command.

```
device# clear access-list subnet-broadcast accounting global
```

Syntax: clear access-list subnet-broadcast accounting global

IP broadcast ACL CAM

For XMR Series and MLX Series devices, a separate sub-partition is created for IP broadcast ACLs in the in-bound ACL partition. You can configure the maximum CAM size that you want for an IP broadcast ACL by comparing it with the available CAM resources. If there are not enough CAM resources available, then the CAM profile can be changed to specify the maximum allowable CAM size for the IP broadcast ACL.

> **NOTE**
> Hitless upgrade support for the IP broadcast ACL CAM entries is supported only on the XMR Series and MLX Series devices.

Considerations for implementing IP broadcast ACL

The considerations that must be observed while implementing IP broadcast ACL are listed as follows:

- If filtering of IP directed broadcast traffic using an ACL is enabled at the IP interface level, then the IP broadcast ACL CAM entry matching is based on ACL group ID and VLAN ID for the physical IP interface and virtual IP interface, respectively.
- If filtering of IP directed broadcast traffic using an ACL is enabled globally, then the IP broadcast ACL CAM entry matching is completed for the default VRF.
- Physical ports must undergo the VRF membership check only if the ports have implemented IP broadcast ACL.

Specifying the maximum CAM size for IP broadcast ACL

To configure the maximum allowable number of ACL CAM entries assigned to the IP broadcast ACL CAM sub-partition, enter the following command.

```
device(config)# system-max subnet-broadcast-acl-cam 2000
```

Syntax: [no] **system-max subnet-broadcast-acl-cam** *max-cam-entries*

The **max-cam-entries** parameter specifies the maximum CAM size that you want for an IP broadcast ACL. On the XMR Series and MLX Series devices, the minimum value supported is 0 and the maximum value supported is 4096. The default value is 0.

The **no** option is used to reset the maximum allowable CAM value to the default value.

> **NOTE**
> The system maximum value for the IP broadcast ACL CAM entries is configurable only on the XMR Series and MLX Series devices.

Upon configuration, the system verifies the input value with the amount of CAM resources available. If the system is unable to allocate the requested space, the following error message is displayed.

```
Error - IPV4 subnet-broadcast-acl-cam roundup value (4096 - power of 2) exceeding available CAM resources
Total IPv4 ACL CAM:                                49152(Raw Size)
IPv4 Multicast CAM:                                       32768(Raw Size)
IPv4 Receive ACL CAM:                                      8192(Raw Size)
IPv4 Source Guard CAM:                          4096
Reserved IPv4 Rule ACL CAM:                               1024(Raw Size)
Available Subnet Broadcast ACL CAM: 3072(Raw Size) 1536(User Size)
```

If there are not enough CAM resources available, you can change the CAM profile and configure the sub-partition size before doing a reload. The change is permitted only if the new CAM profile can support the currently defined system maximum values for the various CAM partitions.

Rebinding of IP broadcast ACL CAM entries

To rebind IP broadcast ACL CAM entries, enter the following command.

```
device(config)# ip rebind-subnet-broadcast-acl
```

Syntax: [no] **ip rebind-subnet-broadcast-acl**

The **no** option is used to disable rebinding of IP broadcast ACL CAM entries.

> **NOTE**
> The **ip rebind-subnet-broadcast-acl** command is applicable only for XMR Series and MLX Series devices. For CES 2000 Series and CER 2000 Series devices, rebinding of an IP broadcast ACL is done using the **ip rebind-acl all** command.

Warning message for rebinding IP broadcast ACL

When binding an IP broadcast ACL globally or at the IP interface level, if you make a configuration change to the default VRF by adding or deleting the IP address from an interface, the following warning message is triggered asking you to rebind the IP broadcast ACL.

```
Warning: IP Address configuration change detected, rebind IP subnet broadcast ACLs to update the CAM
```

The warning message is not triggered if you make an ACL configuration change to the default VRF by adding or deleting or modifying the ACL definition.

IP receive ACLs

The IP receive access-control list feature (rACL) provides hardware-based filtering capability for IPv4 traffic destined for the CPU in all the VRFs such as management traffic. Its purpose is to protect the management module's CPU from overloading due to large amounts of traffic sent to one of the Extreme device's IP interfaces. Using the **rACL** command, the specified ACL is applied to every interface on the Extreme device. This eliminates the need to add an ACL to each interface on an Extreme device.

The rACL feature is configured by creating an ACL to filter traffic and then specifying that ACL in the rACL command. This applies the ACL to all interfaces on the device. The destination IP address in an ACL specified by the rACL command is interpreted to apply to all interfaces in the default VRF of the device. This is implemented by programming an ACL entry in CAM that applies the ACL clause for each interface.

For example there are the following three interfaces defined on a device:

- loopback 1 = 10.2.2.2
- ethernet 4/1 = 10.10.10.1
- virtual ethernet interface 1 = 10.10.20.1

The access list defined in the following command will act to deny ICMP traffic to each of the defined interfaces.

```
device(config)# access-list 170 deny icmp host 10.1.1.1 any
```

The ACL CAM would then be programmed with the following three entries:

- deny icmp host 10.1.1.1 host 10.2.2.2
- deny icmp host 10.1.1.1 host 10.10.10.1
- deny icmp host 10.1.1.1 host 10.10.20.1

> **NOTE**
> You must rebind an rACL whenever it is changed, as described in Rebinding a rACL definition or policy-map on page 146, otherwise now invalid entries will still be in CAM.

> **NOTE**
> For more information on configuring the **acl-mirror-port** command for IP Receive ACLs, refer to *Extreme NetIron Layer 2 Switching Configuration Guide*

Configuration guidelines for IPv4 receive ACLs

Use the following considerations when configuring IPv4 receive ACLs (rACLs):

- (Interface-level inbound IPv4 ACL or RL-ACLs) Traffic matching rACLs is not subject to interface-level ACL or RL-ACLs. You must take care to configure an rACL such that only management traffic matches the rACL clauses.
- (Interface-level inbound IPv4 ACL or RL-ACLs) On an interface, we support either launching an IPv4 inbound or L2 inbound ACL CAM lookup, but not both. For interfaces with L2 inbound ACLs, rACL filtering will be performed by software. Therefore, only traffic permitted by L2 inbound ACL will be processed by rACLs. Note that rate-limiting using rACLs will not be applicable for such traffic.
- (MLX Series and XMR Series devices only) For traffic destined to VRRP/E router instances, enable the **racl-vrrp-vrip-filter ip-packets** command. For details, refer to Summary of ACL system and policy parameters on page 181.
- VLAN ID translation or Inner VLAN ID translation program L2 inbound ACL CAM entries, and therefore, for ports in VLAN or Inner VLAN translation group, rACL filtering is performed in software. Note that rate limiting using rACLs will not be applicable for traffic incoming on such interfaces.
- Global DOS attack policies are supported in software. The order of precedence is:
  - rACL filtering (either in hardware or software)
  - Global DOS attack policies (only in software)
- IPv4 rACLs are applicable only for line card interfaces, not for management ethernet interfaces.
- rACLs are not supported for non default VRF.

Interfaces on which no UDA or Layer 2 ACL is applied

If neither an ingress Layer 2 ACL nor a UDA is applied on an interface, even if a Layer 3 rACL is configured on the device globally, rACL CPU filtering is not triggered by default on such an interface. You need to explicitly enable the **racl-cpu-filtering ip-packets** command at device level:

```
device(config)# acl-policy
device(config-acl-policy)# racl-cpu-filtering ip-packets
```

If **racl-cpu-filtering ip-packets** is enabled and rACLs are applied on a device, guidelines for Layer 3 Unicast packets are as follows:

- If the destination address in the incoming packet is a connected IPv4/IPv6 address (physical/VE Interface or loopback), the decision to permit or deny the packet is taken in hardware—under CAM programming—subject to the following conditions:
  - Packets matching a configured rACL permit rule are forwarded to the M-CPU for processing.
  - Packets that do not match any rACL permit rules are dropped in hardware, by a default deny rule applicable if the destination is a connected IPv4/IPv6 address.
- If the destination is not a connected IPv4/IPv6 address, the packet is not destined for M-CPU, and is not subject to rACL CPU filtering.

If **racl-cpu-filtering ip-packets** is enabled and rACLs are applied on a device, guidelines for Layer 3 Broadcast or Multicast packets are as follows:

- Packets that match an rACL permit rule are forwarded to the M-CPU.
- Packets that match an rACL deny rule are dropped.
- The default implicit rACL rule is permit; packets that do not match any rACL deny rule are permitted. (A default implicit permit enables protocol packets to be processed by the M-CPU. Multicast-based protocols such as OSPF, LDP, RIP, VRRP, and PIM work seamlessly, even when rACL CPU filtering is configured.)
- If you add an explicit "deny ip/ipv6 any any" rule to an rACL, make sure to add permit rules for the relevant multicast-protocol destination addresses.

Interfaces on which a UDA or Layer 2 ACL is also applied

M-CPU-bound Layer 3 Broadcast, Unicast, or Multicast packets permitted by the UDA or Layer 2 ACL are subject to the following conditions:

- Even if you do not explicitly configure **racl-cpu-filtering** at device level, it is configured automatically on ingress interfaces with a Layer 2 ACL or UDA applied.
- If the packet matches an rACL permit rule, the packet is forwarded to the M-CPU.
- If the packet matches an rACL deny rule, the packet is dropped.
- If a Layer 2 ACL or UDA is applied, the default implicit rACL rule for CPU filtering is deny. Packets that do not match any of the permit statements are dropped.

IPv4 rACL example for multicast protocol

In the following example, racl-cpu-filtering is enabled. If you conclude the rACL with a "deny any any" rule (in order for IPv4 OSPFv2 peering to work seamlessly) you also need to include the following permit rules:

- Unicast peer IPv4 address
- Destination multicast addresses: 224.0.0.5 and 224.0.0.6

Configuring IPv4 rACLs

Configuring IPv4 rACLs requires the following steps:

- Configuring an rACL and establishing the sequence of rACL commands.
- Applying rate limiting on rACLs defined traffic.
- Specifying the maximum number of rACL entries.
- Rebinding a rACL definition or policy map.

You can bind multiple IPv4 rACLs, up to a maximum of 199. You must, however, ensure that no explicit **permit ip any any** or **deny ip any any** clause exists in any of the rACLs except the last one.

> **NOTE**
> An implicit **deny ip any any** rule is programmed by default as the final rACL rule. This implicit rule drops traffic destined for connected IPv4/IPv6 addresses that did not match any permit rules.

> **NOTE**
> An explicit **deny ip any any** filter does not match any multicast traffic. Since the destination field in the ACL contains "any", rACLs program interface IP addresses in the CAM instead of the destination's IP address. To match the multicast traffic, you should specify the multicast group address in the destination field in the ACL.

> **NOTE**
> For details of rACL logging, refer to Enabling L3 rACL logging on page 186.

Configuring rACL to apply a defined ACL and establishing the sequence of rACL commands

To configure rACL to apply ACL number "101" with a sequence number of "15" to all interfaces within the default VRF for all CPU-bound traffic, enter the following command:

```
device(config)# ip receive access-list 101 sequence 15
```

If you are using loopback interfaces for all BGP peering sessions, you can define an ACL that only permits BGP traffic from a specified source IP address. Where the peer source has an IP address of 10.1.1.1 and the loopback IP address on the device is 10.2.2.2, the access list command is configured as shown in the following.

```
device(config)# access-list 106 permit tcp host 10.1.1.1 host 10.2.2.2 eq bgp
```

The rACL command that implements ACL "106" is configured as shown in the following.

```
device(config)# ip receive access-list 106 sequence 10
```

To configure rACL to apply the named ACL "**acl_stand1**" with a sequence number of "10" to all interfaces within the default VRF for all CPU-bound traffic, enter the following command:

```
device(config)# ip receive access-list acl_stand1 sequence 10
```

Syntax: **[no] ip receive access-list** { **acl-num** | *acl-name* } **sequence seq-num**

The {*acl-num* | *acl-name*} variable identifies the ACL (standard or extended) that you want to apply to all interfaces within the default VRF for all CPU-bound traffic.

The sequence *seq-num* option defines the sequence in which the rACL commands will be applied. The valid range is from 1 through 200. Commands are applied in order of the lowest to highest sequence numbers. For example, if the following rACL commands are entered.

```
device(config)# ip receive access-list 100 sequence 10
device(config)# ip receive access-list 101 sequence 25
device(config)# ip receive access-list 102 sequence 15
```

The effective binding of the commands will be in the following order.

```
ip receive access-list 100 sequence 10
ip receive access-list 102 sequence 15
ip receive access-list 101 sequence 25
```

Using the [no] option removes the rACL access list defined in the command.

Applying rate limiting on rACL defined traffic

The rACL feature allows you to apply rate limiting to CPU-bound traffic using the **policy-map** and **strict-acl** options of the **ip receive access-list** command.

To configure rACL to apply the named ACL "acl_stand1" with a **policy-map** "m1", enter the following command.

device(config)# **ip receive access-list acl_stand1 sequence 10 policy-map m1**

Syntax: [no] ip receive access-list { acl-num | *acl-name* } sequence seq-num [policy-map policy-map-name [strict-acl]]

By default, traffic matching the "permit" clause in the specified ACL is permitted and traffic matching the "deny" clause in the ACL is dropped.

When the **policy-map** option is used, traffic matching the permit clause of the specified ACL is rate-limited as defined in the policy map specified by the *policy-map-name* variable and traffic matching the "deny" clause in the ACL is permitted but not rate limited. Using the [no] option removes the policy map defined in the command.

When the **policy-map** option is used with the **strict-acl** option, traffic matching the permit clause of the specified ACL is rate-limited as defined in the policy map specified by the *policy-map-name* variable and traffic matching the "deny" clause in the ACL is dropped. Using the [no] option removes the **strict-acl** option for the rACL command defined in the command.

Specifying the maximum number of rACLs supported in CAM

You can configure the number of software ACL CAM entries available for rACLs using the following command.

```
device(config)# system-max receive-cam 2048
```

Syntax: [no] system-max receive-cam *number*

The *number* variable is the maximum number of ACL CAM entries that are allowed. Acceptable values are powers of 2 from 512 through 16384. Examples of powers of 2 are 512, 1024, 2048, and so on. The default value is 1024.

> **NOTE**
> You must reload the device for this command to take
> effect.

If you enter a value that is not a power of 2, the system rounds off the entry to a number less than the input value. For example, if you enter 16383, which is not a power of 2, the system rounds it off to 8192 and displays a warning.

```
device(config)# system-max receive-cam 16383
Warning - Receive ACL CAM size requires power of 2, round down to 8192
Reload required. Please write memory and then reload or power cycle the system.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

The **system-max receive-cam** command also checks if there is enough space in the IPv4 ACL partition before allocating more space to rACL sub-partition. The following error is displayed when there is less space to increase rACL.

```
device(config)# system-max receive-cam 16384
Error - IPv4 Receive ACL CAM (16384) exceeding available CAM resources
Total IPv4 ACL CAM:              49152(Raw Size)
```

```
IPv4 Multicast CAM:            32768(Raw Size)
IPV4 Broadcast ACL CAM:        0(Raw Size)
IPv4 Source Guard CAM:         0(Raw Size)
Reserved IPv4 Rule ACL CAM:    1024(Raw Size)
Available IPv4 Receive ACL CAM: 15360(Raw Size) 7680(User Size)
```

NOTE

The following limitations apply when the *number* variable has a maximum limit of 16384.

- The 16K Receive ACL CAM partition is not supported on the cam profiles such as IPv6, Multi-service 3, and Multi-service 4.

- Depending on the configuration, any of the IPv4 ACL sub-partitions such as IP Source Guard, Broadcast ACL, IP Multicast, and Open Flow should be decreased to allow the creation of the 16K rACL partition.

Rebinding a rACL definition or policy-map

If a change is made to the definition of an IP rACL or to a rate-limiting, policy map that is specified for an rACL, you must perform a rebind using either of the following commands:

```
device(config)# ip rebind-receive-acl all
```

or

```
device(config)# ip receive rebind-acl-all
```

Syntax: ip rebind-receive-acl all

Syntax: ip receive rebind-acl-all

NOTE

If you add or delete an IP address to or from a device interface, you need to rebind the IP receive ACLs.

Deactivating the rACL configuration

To deactivate the IPv4 rACL configuration and remove all the rules from CAM, enter the following command.

```
device(config)# ip receive deactivate-acl-all
```

Syntax: [no] ip receive deactivate-acl-all

The **no** form of this command reactivates the rACL configuration.

NOTE

To prevent ACL binding to CAM after reload, use the **write memory** command to save this configuration change permanently.

Deleting the rACL configuration

To delete the rACL configuration and remove all IPv4 rACL rules from the system, use the following command.

```
device(config)# ip receive delete-acl-all
This command deletes all IP Receive ACLs from system.
Are you sure? (enter 'y' or 'n'): y
```

Syntax: ip receive delete-acl-all

Displaying accounting information for rACL

To display rACL accounting information for ACL number "102", use the following command.

```
device# show access-list receive accounting 102
```

To display rACL accounting information for the ACL named "acl_ext1", use the following command.

```
device(config)# show access-list receive accounting name acl_ext1
IP Receive ACL Accounting Information:
IP Receive ACL acl_ext1
ACL hit count for software processing (accum)              0
HW counters:
    0: permit tcp any host 10.10.10.14
        Hit count: (1 sec)               0  (1 min)                   0
                   (5 min)               0  (accum)                   0
```

Syntax: show access-list receive accounting { acl-num | name acl-name }

The *acl-num* variable specifies, in number format, the ACL that you want to display rACL statistics for.

The **name** *acl-name* variable specifies, in name format, the ACL that you want to display rACL statistics for.

Clearing accounting information for rACL

To clear rACL accounting information for ACL number "102", use the following command.

```
device(config)# clear access-list receive 102
```

To clear rACL accounting information for a rACL named "acl_ext1", use the following command.

```
device(config)# clear access-list receive name acl_ext1
```

Syntax: clear access-list receive { acl-num | name *acl-name* }

The *acl-num* variable specifies the ACL that you want to clear rACL statistics for in number format. The **name**acl-name variable specifies clearing accounting statistics for a named IPv4 rACL.

To clear rACL accounting statistics for all configured IPv4 rACLs, enter the following command.

```
device(config)# clear access-list receive all
```

Syntax: clear access-list receive all

The **all** parameter specifies clearing accounting statistics for all configured IPv4 rACLs.

ACL CAM sharing for inbound IPv4 ACLs (XMR Series and MLX Series devices)

ACL CAM sharing allows you to conserve CAM by sharing it between ports that are supported by the same packet processor (PPCR). If this feature is enabled globally, you can share CAM space that is allocated for inbound ACLs between instances on ports that share the same packet processor (PPCR). For example, if you have bound-inbound ACL 101 to ports 1/1 and 1/5, the ACL is stored in a single location in CAM and used by both ports. The table below describes which ports share PPCRs and can participate in ACL CAM sharing.

Module type	PPCR number	Ports supported by PPCR
20 x 1G	PPCR 1	1 - 20
4 x 10G	PPCR 1	1 - 2
	PPCR 2	3 - 4
2 x 10G	PPCR 1	1 - 2

Considerations when implementing this feature

The following consideration apply when implementing this feature:

- If you enable ACL CAM sharing, ACL statistics will be generated per-PPCR instead of per-port. If you require the statistics per-port granularity for your application, you cannot use this feature.
- This feature is only applicable for inbound IPv4 ACLs, IPv6 ACLs, VPNv4 ACLs, Layer-2 ACLs, and Global PBR policies.
- This feature is not applicable for ACL-based rate-limiting and interface-level PBR policies.
- This feature cannot be applied to a virtual interface.
- CAM entry matching within this feature is based on the ACL group ID.
- This feature cannot co-exist with IP Multicast Routing or IP Multicast Traffic Reduction.

Configuring ACL CAM sharing for IPv4 ACLs

> NOTE
> The **enable-acl-cam-sharing** command is not supported on CES 2000 Series or CER 2000 Series devices.

When enabled, ACL CAM sharing for IPv4 ACLs is applied across all ports in a system. To apply ACL CAM sharing for IPv4 ACLs on an Extreme device, use the following command.

```
device(config)# acl-policy
device(config-acl-policy)# enable-acl-cam-sharing
```

Syntax: **[no] enable-acl-cam-sharing**

ACL CAM sharing is disabled by default.

Matching on TCP header flags for IPv4 ACLs

You can match packets for one additional TCP header flag using named or numbered IPv4 extended ACLs.

The following command implements the additional TCP parameter for IPv4 ACLs.

Syntax:

[no] access-list num permit | deny tcp any any syn

Example

```
device(config)# ip access-list extended 133
device(config-ext-nacl-)# permit tcp 172.24.33.0 0.0.0.255 198.124.133.0 0.0.0.255 syn
device(config-ext-nacl-)# permit tcp host 172.124.133.10 eq 1024 host 198.124.133.10 syn eq 1025
device(config-ext-nacl-)# permit tcp any any syn
```

Extended named or numbered ACL IDs 100 -199 support the syn keyword.

The **tcp** parameter indicates that you are filtering the TCP protocol.

The **syn** parameter directs the ACL to permit or deny based upon the status of the syn flag in the TCP header. If the contents of the flag is "1" the condition is met.

The **syn** key word may be used in combination with other TCP filter options.

ACL accounting

Multi-Service devices monitor the number of times an ACL is used to filter incoming or outgoing traffic on an interface. The **show access-list accounting** command displays the number of "hits" or how many times ACL filters permitted or denied packets that matched the conditions of the filters.

> **NOTE**
> ACL accounting does not tabulate nor display the number of implicit denials by an ACL.

Counters, stored in hardware, keep track of the number of times an ACL filter is used.

The counters that are displayed on the ACL accounting report are:

- 1s - Number of hits during the last second. This counter is updated every second.
- 1m - Number of hits during the last minute. This counter is updated every one minute.
- 5m - Number of hits during the last five minutes. This counter is updated every five minutes.
- ac - Accumulated total number of hits. This counter begins when an ACL is bound to an interface and is updated every one minute. This total is updated until it is cleared.

The accumulated total is updated every minute. For example, a minute after an ACL is bound to a port, it receives 10 hits per second and continues to receive 10 hits per second. After one minute, the accumulated total hits is 600. After 10 minutes, there will be 6000 hits.

The counters can be cleared when the device is rebooted, when an ACL is bound to or unbound from an interface, or by entering a **clear access-list** command.

Enabling and disabling ACL accounting on XMR Series and MLX Series devices

ACL accounting is not automatically enabled on XMR Series and MLX Series devices. Before you can collect ACL accounting statistics, you must enter the following command.

```
device(config)# acl-policy
device(config-acl-policy)# enable-acl-counter
```

Syntax: **[no] enable-acl-counter**

> **NOTE**
> Enabling or disabling ACL accounting affects the gathering of statistics from all ACL types (Layer-2, IPv4 and IPv6).

ACL accounting on CES 2000 Series and CER 2000 Series devices

The following special considerations affect how ACL accounting is configured on CES 2000 Series and CER 2000 Series devices.

Enabling ACL accounting on CES 2000 Series and CER 2000 Series devices

On CES 2000 Series and CER 2000 Series devices you enable ACL accounting explicitly in each clause of an ACL for which you want to gather statistics. Enable ACL accounting in an individual filter by including the keyword **enable-accounting** immediately after the **permit** or **deny** keyword.

To create an ACL filter clause with ACL accounting enabled, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)# access-list 100 permit enable-accounting ip any any
```

The example above will add a permit clause to ACL 100 with accounting enabled.

Syntax: [no] access-list *num* | *name* **permit** | **deny enable-accounting**

> NOTE
> ACL accounting on CES 2000 Series and CER 2000 Series devices is applicable only on the outbound counter, not the inbound counter.

ACL rate-limiting and ACL accounting

CAM resources are shared on CES 2000 Series and CER 2000 Series devices between ACL accounting and ACL rate-limiting. This limits the number of ACL accounting instances available on the system.

To check the availability of ACL accounting and ACL rate-limiting resources, use the **show resource** command.

```
device# show resource
 . . .
[I cntr/mtrs(1)]  2048(size),  1982(free), 03.22%(used), 0(failed)
[O cntr/mtrs(1)]  2048(size),  1984(free), 03.12%(used), 0(failed)
 . . .
```

The above example shows only the output related to ACL rate-limiting and ACL accounting resources, and indicates that 3.22% of input resources and 3.12% of output resources have been used.

> NOTE
> On a CES 2000 Series or CER 2000 Series device, each outbound ACL clause has 2 clauses in the ternary content addressable memory (TCAM). The additional clause is for virtual ports that correspond to the physical ports. Accordingly any outbound ACL requests two separate TCAM indices. For a full TCAM, this results in 2 failure counts.

ACL deny logging and ACL accounting

On CES 2000 Series and CER 2000 Series devices, if ACL deny logging and ACL accounting are enabled on the same ACL clause deny logging takes precedence and ACL accounting statistics will not be available for that clause.

ACL Accounting interactions between L2 ACLs and IP ACLs

You can bind dual inbound ACLs (one L2 ACL and one IP ACL) to a single port on a CES 2000 Series and CER 2000 Series device. Extreme recommends enabling ACL accounting in only one of the ACLs bound to the same port. Including ACL-accounting-enabled clauses in both ACLs can result in anomalous reporting of filtering results.

Displaying accounting statistics for all ACLs

To display a summary of the number of hits in all ACLs on a Multi-Service device, enter the following command.

```
device (config)#show access-list accounting brief
Collecting ACL accounting summary for VE 1 ...  Completed successfully.
ACL Accounting Summary: (ac = accumulated since accounting started)
    Int    In ACL            Total In Hit   Out ACL            Total Out Hit
    VE 1    111               473963(1s)
                              25540391(1m)
                              87014178(5m)
                              112554569(ac)
```

The display shows the following information:

TABLE 15

This field...	Displays...
Collecting ACL accounting summary for *interface*	Shows for which interfaces the ACL accounting information was collected and whether or not the collection was successful.
Int	The ID of the interface for which the statistics are being reported.
In ACL	The ID of the ACL used to filter the incoming traffic on the interface.
Total In Hit*	The number of hits from incoming traffic processed by all ACL entries (filters) in the ACL. A number is shown for each counter.
Out ACL	ID of the ACL used to filter the outgoing traffic on the interface.
Total Out Hit*	The number of hits from incoming traffic processed by all ACL entries (filters) in the ACL. A number is shown for each counter.
* The Total In Hit and Total Out Hit displays the total number of hits for all the ACL entries (or filters) in an ACL. For example, if an ACL has five entries and each entry processed matching conditions three times during the last minute, then the total Hits for the 1m counter is 15.	

Syntax: show access-list accounting brief [l2 | policy-based-routing | rate-limit]

The **l2** parameter limits the display to Layer 2 ACL accounting information.

The **policy-based-routing** parameter limits the display to policy based routing accounting information.

The **rate-limit** parameter limits the display to rate limiting ACL accounting information.

IPv4 ACL accounting statistics are displayed if no option is specified.

Displaying statistics for an interface

To display statistics for an interface, enter commands such as the following.

```
device (config)#show access-list accounting ve 1 in
Collecting ACL accounting for VE 1  ...  Completed successfully.
ACL Accounting Information:
Inbound: ACL 111
    0: sequence 40 permit tcp any any
       Hit count: (1 sec)          555520   (1 min)          33488360
                  (5 min)               0   (accum)         133365917
    1: sequence 50 deny ip any any log
       Hit count: (1 sec)          281679   (1 min)          16976780
                  (5 min)               0   (accum)          67605380
    2: sequence 10 deny igmp any any
       Hit count: (1 sec)          102400   (1 min)           6173062
                  (5 min)               0   (accum)          24583600
    3: sequence 20 permit icmp any any
       Hit count: (1 sec)           76800   (1 min)           4629600
                  (5 min)               0   (accum)          18437700
    4: sequence 30 deny udp any any
       Hit count: (1 sec)           20480   (1 min)           1234560
                  (5 min)               0   (accum)           4916720
```

The display shows the following information:

TABLE 16

This field...	Displays...
The IP multicast traffic snooping state	The first line of the display indicates whether IP multicast traffic snooping is enabled or disabled. If enabled, it indicates if the feature is configured as passive or active.
Collecting ACL accounting summary for *interface*	Shows the interface included in the report and whether or not the collection was successful.

TABLE 16 (continued)

This field...	Displays...
Outbound or Inbound ACL ID	Shows the direction of the traffic on the interface and the ID of the ACL used.
#	Shows the index of the ACL entry, starting with 0, followed by the permit or deny condition defined for that ACL entry. (The first entry created for an ACL is assigned the index 0. The next one created is indexed as 1, and so on.) ACL entries are arranged beginning with the entry with the highest number of hits for IPv4 ACLs. For all other options, ACL entries are displayed in order of ascending ACL filter IDs.
Hit count	Shows the number of hits for each counter. NOTE On CES 2000 Series and CER 2000 Series devices this field will show "Accounting is not enabled" for those clauses which do not include the keyword **enable-accounting** and "Accounting is not available due to deny logging" for those clauses in which accounting and deny logging are both enabled.

Syntax: show access-list accounting ethernet [slot/port | ve ve-number] in | out [l2 | policy-based-routing | rate-limit]

Use **ethernet** *slot* or *port* to display a report for a physical interface.

Use **ve** *ve-number* to display a report for the ports that are included in a virtual routing interface. For example, if ports 1/2, 1/4, and 1/6 are all members of ve 2, the report includes information for all three ports.

Use the **in** parameter to display statistics for incoming traffic; **out** for outgoing traffic.

The **l2** parameter limits the display to Layer 2 ACL accounting information.

The **policy-based-routing** parameter limits the display to policy-based routing accounting information. This option is only available for incoming traffic.

The **rate-limit** parameter limits the display to rate limiting ACL accounting information.

> NOTE
> If the ACL definition is modified after it is applied to the interface, then discrepancy may exist in the ACL accounting information

Clearing the ACL statistics

Statistics on the ACL account report can be cleared:

- When a software reload occurs
- When the ACL is bound to or unbound from an interface
- When you enter the **clear access-list** command, as in the following example.

```
device(config)# clear access-list all
```

Syntax: clear access-list all | ethernet slot/port | ve ve-num

Enter **all** to clear all statistics for all ACLs.

Use **ethernet** *slot*/*port* to clear statistics for ACLs bound to a physical port.

Use **ve** *ve-number* to clear statistics for all ACLs bound to ports that are members of a virtual routing interface.

IPv6 ACLs

IPv6 ACL overview and guidelines

Extreme devices support IPv6 access control lists (ACLs), which you can use for traffic filtering.

For details on Layer 2 ACLs, refer to Layer 2 ACLs on page 83.

For details on IPv4 ACLs, refer to IPv4 ACLs on page 97.

An ACL contains one or more rules that permit or deny packets matching a specified source or destination prefix. For the maximum numbers of ACLs and rules supported, refer to ACL and rule limits on page 82. In ACLs with multiple statements, you can specify a sequence number for each statement. The sequence number determines the order in which the statements run. The last statement is an implicit deny statement for all packets that do not match the previous statements in the ACL.

You can configure an IPv6 ACL on a global basis, then apply it to the incoming or outgoing IPv6 packets on specified Extreme device interfaces. You can apply only one IPv6 ACL to incoming traffic for an interface and only one IPv6 ACL to outgoing traffic on an interface. When an interface sends or receives an IPv6 packet, it applies the statements within the ACL (in their order of occurrence in the ACL) to the packet. When a match occurs, the Extreme device takes the specified action (permits or denies the packet) and stops further comparison for that packet.

On XMR Series and MLX Series devices, you cannot bind Layer 2 ACLs and IPv4 or IPv6 ACLs for inbound traffic to the same port. However, you can perform the following configuration:

- On one port, bind a Layer 2 ACL.
- On a second port, bind one or both of the following:
  - An IPv4 ACL (standard or extended)
  - An IPv6 ACL (standard or extended)

CES 2000 Series and CER 2000 Series devices enable you to bind a Layer 2 ACL, an IPv4 ACL, and an IPv6 ACL for inbound traffic on the same port.

On all platforms, for outbound traffic, you can bind only one ACL—L2 or IPv4 or IPv6.

For dynamic LAG creation and deletion using IPv6 ACLs, before a LAG is formed, all ports which will be part of the LAG must have the same configuration. After the LAG is removed, all ACL bindings (if any) are propagated to all of the secondary ports.

IPv6 ACLs enable traffic filtering based on the following information:

- IPv6 protocol
- Source IPv6 address
- Destination IPv6 address
- ICMP message type (if the protocol is ICMP)
- Source TCP or UDP port (if the IPv6 protocol is TCP or UDP)
- Destination TCP or UDP port (if the IPv6 protocol is TCP or UDP)
- DSCP values

The IPv6 protocol can be one of the following well-known names, or any IPv6 protocol number from 0 - 255:

- Authentication Header (AHP)
- Encapsulating Security Payload (ESP)
- Internet Control Message Protocol (ICMP)
- Internet Protocol Version 6 (IPv6)

- Stream Control Transmission Protocol (SCTP)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)
- VLAN ID

For TCP and UDP, you also can specify a comparison operator and port name or number. For example, you can configure a policy to block Web access to a specific site by denying all TCP port 80 (HTTP) packets from a specified source IPv6 address to the IPv6 address of the site.

To disable IPv4 and IPv6 ACLs on the terminating node of a GRE tunnel, see the "Bypassing ACLs in a GRE tunnel" topic in the *Extreme NetIron Layer 3 Routing Configuration Guide*.

To disable IPv6 ACLs on the terminating node of an IPv6-over-IPv4 tunnel, see the "Bypassing ACLs in an IPv6-over-IPv4 tunnel" topic in the *Extreme NetIron Layer 3 Routing Configuration Guide*.

IPv6 ACL limitations

The following limitations apply to IPv6 ACLs:

- On both VE and physical interfaces, IPv6 ACLs are supported for routed traffic only, but not for switched traffic. This limitation also applies to mirror and deny-log actions.
- You cannot apply an IPv6 ACL to a subset of ports within a VE.
- You cannot apply an IPv6 ACL on a management interface.

Configuration considerations for dual inbound ACLs on CES 2000 Series and CER 2000 Series devices

CES 2000 Series and CER 2000 Series devices enable you to bind a Layer 2 ACL, an IPv4 ACL, and an IPv6 ACL to the same port, as follows:

- A Layer 2 ACL is bound to the port.
- An IPv4 and an IPv6 ACL, or one of the two, are bound to the same port.

The filtering algorithm is as follows:

1. An incoming packet is first examined by the IPv4 ACL or the IPv6 ACL, depending on the protocol.
2. If the packet is denied by the IPv4/IPv6 ACL, the packet is dropped, without being examined by the L2 ACL.
3. If the packet is permitted by the IPv4/IPv6 ACL, it is then examined by the L2 ACL.

However, there is an implicit "deny" at the end of any ACL. To enable the above algorithm, IPv4/IPv6 ACLs intended for this scenario must include a "permit any" filter as the last rule. Such a rule ensures that even packets not explicitly permitted by the IPv4/IPv6 ACL are passed to the L2 ACL.

Configuration considerations for IPv6 ACL and multicast traffic for 2X100GE modules installed on NetIron MLX and NetIron XMR devices

When applied to a 100GE interface, the following behavior will be applicable for IPv6 inbound ACLs:

1. You cannot match IPv6 multicast packets using filters with matching enabled on one or more of the following fields:
   a) TCP/UDP source port
   b) TCP/UDP destination port
   c) ICMP type/code

      The exception to this rule will be ICMP filters to match neighbor solicitation and router solicitation packets, such as "permit icmp any any nd-ns" and "permit icmp any any router-solicitation", that will be programmed to match all ICMP multicast packets irrespective of the ICMP type or code value.

2. Implicit "deny ipv6 any any" will not match multicast packets. However explicit "deny ipv6 any any" or any other filter with matching based on IPv6 header fields only will match multicast packets.

      NOTE
      The above rules are only applicable for IPv6 inbound ACLs. They are not applicable for IPv6 outbound ACLs.

Configuration considerations for IPv6 outbound ACLs on VPLS, VLL, and VLL-local endpoints

The following considerations apply to IPv6 outbound ACLs on VPLS, VLL, and VLL-local endpoints:

- Configure the port as a VPLS, VLL, or VLL-local endpoint and then bind the IPv6 outbound ACL to the port.
- Remove the IPv6 outbound ACL from a VPLS, VLL, or VLL-local endpoint before removing the port from the VPLS, VLL, or VLL-local instance or corresponding VLAN.
- Remove the IPv6 outbound ACL from a VPLS, VLL, or VLL-local endpoint before deleting the VPLS, VLL, or VLL-local instance or corresponding VLAN.
- If the VPLS, VLL, or VLL-local endpoint is a LAG port, you must first remove the IPv6 outbound ACL from the primary LAG port before deleting the LAG. This restriction is applicable even if you attempt to delete the lag using **force** keyword.
- If a VLL or VLL-local endpoint is a LAG port with an IPv6 outbound ACL, you must first remove the IPv6 outbound ACL from the primary LAG port before dynamically removing a port from the LAG.
- Ensure that no VPLS, VLL, or VLL-local endpoint exists with an IPv6 outbound ACL before entering the command: **no router mpls** .

Configuration considerations before upgrade or downgrade

Before performing a downgrade from NetIron 6.0.0, the following IPv6 ACL configuration considerations should be reviewed.

Before a downgrade from NetIron 6.0.0, the priority-mapping keyword in IPv6 ACL filters should be suppressed using the **suppress-ipv6-priority-mapping** command.

      NOTE
      The **suppress-ipv6-priority-mapping** command is intended for downgrade purposes only. This command will not be displayed in **show running-config** output, but will be displayed in **show acl-policy** output.

Suppress priority-mapping keyword in IPv6 ACL filters

```
device(acl-policy)#suppress-ipv6-priority-mapping
```

> **NOTE**
> If you need to downgrade to a version earlier than 5.6, refer to Upgrade and downgrade considerations (5.6.00) on page 197.

Configuration notes: Layer 4 filters in IPv6 ACLs

The following configuration considerations apply to IPv6 ACLs with Layer 4 filters:

- There are two lookups available for ingress direction. In ingress direction, you can bind an IPv6 layer 4 ACL with IPv4 layer 4 ACLs and layer 3 ACLs on the same port.
- XMR Series and MLX Series devices have one CAM lookup for outbound ACLs.
- Only one ingress L2 or IPv6 ACL is allowed per port. However, they cannot be applied simultaneously.
- There is only one lookup available for egress direction. When you bind outbound IPv6 L4 ACL to a port, the port does not allow L2, IPv4, or IPv6 ACL in that egress direction.
- Layer 4 ACLs filter incoming traffic based on IPv6 packet header fields. The following attributes can be added to the IPv6 packet header fields:
    - VLAN ID
    - Source IPv6 address (SIP) prefix
    - Destination IPv6 address (DIP) prefix
    - IP protocol (SPI matching is not supported for AHP or ESP
    - UDP, TCP, or SCTP source port
    - UDP, TCP, or SCTP destination port
    - TCP flags - established (RST or ACK)
    - TCP flags - SYN
    - ICMP type and code
    - DSCP value
    - IPv6 fragments
    - source routed packets
    - specific routing header type
- The following actions are available for the ingress ACL:
    - Permit
    - Deny
    - Copy-sflow
    - Drop-precedence
    - Drop-precedence-force
    - Priority-force
    - Mirror
    - Log

The following actions are available for the egress ACL:

- Permit
- Deny

Unsupported features for CES 2000 Series and CER 2000 Series devices

The following features are not supported on the CES 2000 Series and CER 2000 Series devices:

- The **acl-outbound exclude-switched-traffic** command to exclude switched traffic from outbound ACL filtering is not supported.
- The **acl-frag-conservative** command to change the operation of ACLs on fragmented packets is not supported.
- The **suppress-rpf-drop** command to suppress RPF packet drops for a specific set of packets using inbound ACLs is not supported.
- For all NetIron devices, if a port has an IPv4 or IPv6 ACL applied, you must remove the ACL bindings before adding that port to a VLAN that has a VE interface.
- The only ACL logging supported is for IPv4 deny rules.
- If you need to downgrade to a version earlier than 5.6, refer to

Using IPv6 ACLs as input to other features

You can use an IPv6 ACL to provide input to other features such as route maps and distribution lists. When you use an ACL this way, permit statements in the ACL specify traffic that you want to send to the other feature. If you use deny statements, the traffic specified by the deny statements is not supplied to the other feature.

Configuring an IPv6 ACL

To configure an IPv6 ACL, you must perform the following tasks:

- Create the ACL.
- Apply the ACL to an Extreme device interface.

The following configuration tasks are optional:

- Re-sequence the ACL table
- Control access to and from an Extreme device.

Example configurations

To configure an access list that blocks all Telnet traffic received on port 1/1 from IPv6 host 2000:2382:e0bb::2, enter the following commands.

```
device(config)# ipv6 access-list fdry
device(config-ipv6-access-list-fdry)# deny tcp host 2000:2382:e0bb::2 any eq telnet
device(config-ipv6-access-list-fdry)# permit ipv6 any any
device(config-ipv6-access-list-fdry)# exit
device(config)# int eth 1/1
device(config-if-1/1)# ipv6 traffic-filter fdry in
device(config)# write memory
```

Here is another example of how to configure an ACL and apply it to an interface.

```
device(config)# ipv6 access-list netw
device(config-ipv6-access-list-netw)# permit icmp 2000:2383:e0bb::/64 2001:3782::/64
device(config-ipv6-access-list-netw)# deny ipv6 host 2000:2383:e0ac::2 host 2000:2383:e0aa:0::24
device(config-ipv6-access-list-netw)# deny udp any any
device(config-ipv6-access-list-netw)# permit ipv6 any any
```

The first condition permits ICMP traffic from hosts in the 2000:2383:e0bb::x network to hosts in the 2001:3782::x network.

The second condition denies all IPv6 traffic from host 2000:2383:e0ac::2 to host 2000:2383:e0aa:0::24.

The third condition denies all UDP traffic.

The fourth condition permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL denies all incoming or outgoing IPv6 traffic on the ports to which the ACL is assigned.

The commands in the next example apply the ACL "netw" to the incoming and outgoing traffic on port 1/2 and to the incoming traffic on port 4/3.

```
device(config)# int eth 1/2
device(config-if-1/2)# ipv6 traffic-filter netw in
device(config-if-1/2)# ipv6 traffic-filter netw out
device(config-if-1/2)# exit
device(config)# int eth 4/3
device(config-if-4/3)# ipv6 traffic-filter netw in
device(config)# write memory
```

Here is another example of an ACL.

```
device(config)# ipv6 access-list rtr
device(config-ipv6-access-list rtr)# deny tcp 2001:1570:21::/24
2001:1570:22::/24
device(config-ipv6-access-list rtr)# deny udp any range 5 6 2001:1570:22::/24
device(config-ipv6-access-list rtr)# permit ipv6 any any
device(config-ipv6-access-list rtr)# write memory
```

The first condition in this ACL denies TCP traffic from the 2001:1570:21::x network to the 2001:1570:22::x network.

The second condition denies UDP packets from any source with source UDP ports in ranges 5 to 6 and with the 2001:1570:22::/24 network as a destination.

The third condition permits all packets containing source and destination addresses that are not explicitly denied by the first two conditions. Without this entry, the ACL would deny all incoming or outgoing IPv6 traffic on the ports to which you assign the ACL.

A **show running-config** command output resembles the following.

```
device(config)# show running-config
ipv6 access-list rtr
deny tcp 2001:1570:21::/24 2001:1570:22::/24
 deny udp any range 5 6 2001:1570:22::/24
 permit ipv6 any any
```

A **show ipv6 access-list** command output resembles the following.

```
device(config)# show ipv6 access-list rtr
ipv6 access-list rtr: 3 entries
 10: deny tcp 2001:1570:21::/24 2001:1570:22::/24
 20: deny udp any range 5 6 2001:1570:22::/24
 30: permit ipv6 any any
```

The following commands apply the ACL "rtr" to the incoming traffic on ports 2/1 and 2/2.

```
device(config)# int eth 2/1
device(config-if-2/1)# ipv6 traffic-filter rtr in
device(config-if-2/1)# exit
device(config)# int eth 2/2
device(config-if-2/2)# ipv6 traffic-filter rtr in
device(config)# write memory
```

The ACL functionality for filtering traffic is enhanced with sequence numbers that enable users to insert, modify or delete rules at any position, without having to remove and reapply the entire ACL. A sequence number is assigned to each ACL entry and ACL rules are applied in the order of lowest to highest sequence number. Therefore, you can insert a new filter rule at any position you want in the ACL table, by specifying the sequence number. If you do not specify a sequence number a default sequence number is applied to each ACL entry. The default value is 10+ the sequence number of the last ACL entry provisioned in the ACL table. Therefore, when you do not specify a sequence number, the rule is added to the end of the ACL table. The default value for the first entry in an IPv6 ACL table is "10".

To configure an ACL filter rule with the sequence number "23" for "ipv6_acl", enter the following commands:

```
device(config)# ipv6 access-list ipv6_acl
device(config-ipv6-access-list-ipv6_acl)# sequence 23 deny esp 2::/64 any
```

If the sequence number "23" is already used by another ACL filter rule, the following error message is displayed.

```
"Error: Entry with sequence 23 already exists!"
```

If you specify a sequence number which is greater than the limit (214748364) the following error message is displayed.

"Error: Valid range for sequence is 1 to 214748364"

Default and implicit IPv6 ACL action

The default action when no IPv6 ACLs are configured is to permit all IPv6 traffic. Once you configure an IPv6 ACL and apply it to an interface, the default action for that interface is to deny all IPv6 traffic that is not explicitly permitted on the interface. The following actions can be taken:

- To tightly control access, configure ACLs with permit entries for the access you want to permit. These ACLs implicitly deny all other access.
- To secure access in environments with many users, configure ACLs with explicit deny entries, then add an entry to permit all access to the end of each ACL. The Extreme device permits packets that are not denied by the deny entries.

   NOTE
   Refer to Configuration considerations for IPv6 ACL and multicast traffic for 2X100GE modules installed on NetIron MLX and NetIron XMR devices on page 155 regarding 2x100 GE IPv6 ACL rule exceptions for multicast traffic.

Every IPv6 ACL has the following implicit conditions as the last match condition.

1. **permit icmp any any nd-na** – Allows ICMP neighbor discovery acknowledgement.
2. **permit icmp any any nd-ns** – Allows ICMP neighbor discovery solicitation.

3. **deny ipv6 any any** - Denies IPv6 traffic. You must enter a **permit ipv6 any any** as the last statement in the ACL to permit IPv6 traffic that was not denied by the previous statements.

   The conditions are applied in the order shown above, with **deny ipv6 any any** as the last condition.

   For example, to deny ICMP neighbor discovery acknowledgement, then permit any remaining IPv6 traffic, enter commands such as the following.

   ```
   device(config)# ipv6 access-list netw
   device(config-ipv6-access-list-netw)# permit icmp 2000:2383:e0bb::/64 2001:3782::/64
   device(config-ipv6-access-list-netw)# deny icmp any any nd-na
   device(config-ipv6-access-list-netw)# permit ipv6 any any
   ```

   The first permit statement permits ICMP traffic from hosts in the 2000:2383:e0bb::x network to hosts in the 2001:3782::x network.

   The deny statement denies ICMP neighbor discovery acknowledgement.

   The last command permits all packets that are not explicitly denied by the other entries. Without this entry, the ACL denies all incoming or outgoing IPv6 traffic on the ports to which you assigned the ACL.

   Furthermore, if you add the statement **deny icmp any any** in the access list, then all neighbor discovery messages will be denied. You must explicitly enter the **permit icmp any any nd-na** and **permit icmp any any nd-ns** statements just before the **deny icmp** statement if you want the ACLs to permit neighbor discovery as in this example.

   ```
   device(config)# ipv6 access-list netw
   device(config-ipv6-access-list-netw)# permit icmp 2000:2383:e0bb::/64 2001:3782::/64
   device(config-ipv6-access-list-netw)# permit icmp any any nd-na
   device(config-ipv6-access-list-netw)# permit icmp any any nd-ns
   device(config-ipv6-access-list-netw)# deny icmp any any
   device(config-ipv6-access-list-netw)# permit ipv6 any any
   ```

Re-sequencing an IPv6 ACL table

To allow new ACL entries to be inserted between ACL entries that have consecutive sequence numbers, you can create space between sequence numbers of adjacent filters by regenerating the ACL table.

To re-sequence ACL table "ipv6_acl", use the following commands.

```
device(config)# ipv6 access-list ipv6_acl
device(config-ipv6-access-list-ipv6_acl)# regenerate-seq-num
```

This command regenerates the filter sequence numbers in steps of 10, assigning the default sequence number "10" to the first entry in the table. Any unused IPv6 remarks will be deleted after executing this command. For further information about remarks refer to Adding a comment to an IPv6 ACL entry on page 165.

> **NOTE**
> If sequence numbers generated by the **regenerate-seq-num** command cross the limit (214748364), then re-sequencing of ACL filters will not take place and the following error message is displayed."Error: Valid range for sequence is 1 to 214748364".

> **NOTE**
> The **regenerate-seq-num** command is not allowed while **tftp copy** in progress.

Deleting an IPv6 ACL entry

You can delete an ACL filter rule by providing the sequence number or without providing the sequence number. To delete an ACL filter rule without providing a sequence number you must specify the filter rule attributes. To delete an ACL filter rule providing a sequence number you can provide the sequence number alone or the sequence number and the other filter rule attributes.

To delete a filter rule with the sequence number "23" from access list "ipv6_acl" by specifying the sequence number alone, enter the following command.

```
device(config)# ipv6 access-list ipv6_acl
device(config-ipv6-access-list-ipv6_acl)# no sequence 23
```

You can also delete this entry by specifying both the entry sequence number and filter rule attributes. For example:

```
device(config)# ipv6 access-list ipv6_acl
device(config-ipv6-access-list-ipv6_acl)# no sequence 23 deny esp 2::/64 any
```

Alternatively, you can delete this rule by providing the filter rule attributes only. For example:

```
device(config)# ipv6 access-list ipv6_acl
device(config-ipv6-access-list-ipv6_acl)# no deny esp 2::/64 any
```

> **NOTE**
> If you try to delete an ACL filter rule using the sequence number, but the sequence number that you specify does not exist, the following error message will be displayed. "Error: Entry with sequence 23 does not exist!"

Filtering packets based on DSCP values

To filter packets based on DSCP values, enter commands such as the following.

```
device(config)# ipv6 access-list netw
device(config-ipv6-access-list netw) deny ipv6 any any dscp 3
```

Syntax: [no] ipv6 access-list *name* deny | permit *ipv6-source-prefix/prefix-length* | any *ipv6-destination-prefix/prefix-length* | any [sequence *number*] dscp *dscp-value*

Enter a value from 0 - 63 for the **dscp** *dscp-value* parameter to filter packets based on their DSCP value.

Marking the DSCP value in a packet

To specify the DSCP value to a packet, enter commands such as the following.

> **NOTE**
> **Dscp-marking** is not supported on outbound ACLs.

```
device(config)# ipv6 access-list dscp-markingv6
device(config-ipv6-access-list dscp-markingv6) permit ipv6 any any dscp 20 dscp-marking 10
device(config-ipv6-access-list dscp-markingv6) permit ipv6 any any
```

Syntax: [no] ipv6 access-list *name* deny | permit *ipv6-source-prefix/prefix-length* | any *ipv6-destination-prefix/prefix-length* | any [sequence *number*] dscp *dscp-value* | dscp marking *dscp-value*

Enter a value from 0 through 63 for the **dscp** marking *dscp-value* parameter to mark the DSCP value in the incoming packet with the value you specify.

Filtering packets based on routing header type

You can filter IPv6 packets based on their routing header type. This is of particular value when you want to filter IPv6 source-routed packets to prevent DoS attacks. These packets are type 0.

To filter IPv6 packets based on the routing header type, enter commands such as the following.

```
device(config)# ipv6 access-list drop-source-routed
device(config-ipv6-access-list drop-source-routed) deny ipv6 any any routing-header-type
```

Syntax: [no] ipv6 access-list *name* deny | permitrouting-header-type *type-value*

Enter a value from 0 - 255 for the **routing-header-type***type-value* parameter to filter packets based on their IPv6 header type value.

> **NOTE**
> The **routing-header-type** option is separate and independent of the **routing** option. The **routing-header-type** and **routing** options are mutually exclusive and cannot be used in the same filter.

> **NOTE**
> For more information on configuring the **acl-mirror-port** command, refer to *Extreme NetIron Layer 2 Switching Configuration Guide*.

Displaying IPv6 ACL definitions

To display the IPv6 ACLs configured on an Extreme device, use the **show ipv6 access-list** command.

To display the total number of IPv6 access lists and the number of filters configured for each list, use the **show ipv6 access-list count** command.

```
device(config)# show ipv6 access-list count
Total 4 IPv6 ACLs exist.
IPv6 ACL cust1, total 10 clauses
IPv6 ACL cust2, total 15 clauses
IPv6 ACL cust3, total 12 clauses
IPv6 ACL cust4, total 3 clauses
```

To display information about a specific IPv6 ACL table, you can enter a command such as the following.

```
device# show ipv6 access-list rtr
ipv6 access-list rtr: 3 entries
 10: permit ipv6 host 3000::2 any
 20: deny udp any any
 30: deny ipv6 any any
```

Syntax: show ipv6 access-list { count | access-list-name }

The **count** parameter specifies displaying the total number of IPv6 access lists and the number of filters configured for each list.

The *access-list-name* variable specifies displaying information for a specific IPv6 ACL.

CAM partitioning

CES 2000 Series and CER 2000 Series devices support CAM partitioning.

The size of the extended ingress IPv6 L4 key is 640 bits. The size of the standard ingress ACL key is 320 bits. In internal TCAM, different sized keys can reside next to each other in the same block. In external TCAM, blocks are allocated for ACLs, and different sized keys cannot reside in the same block. An ingress IPv6 L4 key cannot reside in the same block with other ingress ACLs.

You can configure CAM partition to have an ingress ACL into internal TCAM and an egress ACL into external TCAM. The ingress IPv6 L4 key can reside in the same TCAM with other ingress ACLs, but must reside in a different block in the external TCAM.

You can select one key per interface for the following packet types (port or VLAN).

- IPv6 packets
- IPv4 and ARP packets

- Non-IP packets

The following key types apply to layer 2 ACLs:

- Ingress L2 non-IP Key 0
- Egress L2+IPv4+L4 Key

The following keys apply to ether type IPv4, IPv6, or ARP:

- Ingress L2+IPv4/6 Key 1 -- ether type = IPv4 or IPv6
- Ingress IPv4+L4 Key 2 -- ether type = ARP
- Egress L2+IPv6 Key -- ether type = IPv6
- Egress L2+IPv4+L4 Key - ether type = ARP or IPv4

At ingress, each packet is subjected to two lookups. You can direct the system to use a different key for each lookup. Make sure that the source MAC, destination MAC, VLAN ID and ether type are the same for all layer 2 ACL fields. If layer 2 field locations are not same, you will have to create a separate TCAM entry for each layer 2 IPv6 ACL rule or packet type (IPv4, IPv6, and non-IP) combination, for the layer 2 IPv6 ACL to work on all packet types.

TCAM IFSR

The TCAM In-Field Soft Repair (IFSR) feature enhances the data integrity of the TCAM device used in knowledge based processors (KBP). The TCAM is configured to invalidate the entry for which a parity error has been detected. Using the IFSR feature, you can notify users about TCAM memory parity errors through a system log message. The feature supports in lower RMA requests in the field due to repairable errors in the KBP database.

The indices of the CAM entries having parity error are stored in a first in first out (FIFO) method in the KBP. The software reads the FIFO programming periodically and sends out the syslog message with the specific CAM index. The TCAM is configured to invalidate the entry for which a parity error has been detected. This ensures that there are no look-up issues on entry with an error.

To enable the IFSR feature, enter the following command.

```
device(config)# cam ifsr enable
```

To disable the IFSR feature, enter the following command.

```
device(config)# cam ifsr disable
```

cam ifsr enable | disable

Limitation

The KBP FIFO program holding the TCAM indices of the error entries is a limited resource. When the error rate is high, FIFO over flow might happen and some of the errors are lost and are not reported. However, the system informs the user about the IFSR FIFO overflow and sends out another syslog message.

Syslog Messages

The following syslog message is generated when the TCAM entry error is observed by the user:

SYSLOG: <14>Jul 23 11:02:41 sys-np-mac-224 IFSR: Soft Error at TCAM index 0x0002211a of PPCR 1

The following syslog message is generated when the FIFO overflow error is observed by the user:

SYSLOG: <14>Jul 23 11:02:41 sys-np-mac-224 IFSR: Error FIFO Overflow on PPCR 1

Applying an IPv6 ACL

To apply an IPv6 ACL, (for example "access1"), to an interface, enter commands such as the following.

```
device(config)# interface ethernet 3/1
device(config-if-e100-3/1)# ipv6 traffic-filter access1 in
```

This example applies the IPv6 ACL "access1" to incoming IPv6 packets on Ethernet interface 3/1. As a result, Ethernet interface 3/1 denies all incoming packets from the site-local prefix fec0:0:0:2::/64 and the global prefix 2001:100:1::/48 and permits all other incoming packets.

Syntax: [no] ipv6 traffic-filter *ipv6-acl-name* in | out

For the **ipv6-acl-name** parameter, specify the name of an IPv6 ACL created using the **ipv6 access-list** command.

The **in** keyword applies the specified IPv6 ACL to incoming IPv6 packets on the Extreme device interface.

The **out** keyword applies the specified IPv6 ACL to outgoing IPv6 packets on the Extreme device interface.

Reapplying modified IPv6 ACLs

If you make an IPv6 ACL configuration change, you must reapply the ACLs to their interfaces to place the change into effect.

An ACL configuration change includes any of the following:

- Adding, changing, or removing an ACL or an entry in an ACL
- Changing ToS-based QoS mappings

To reapply ACLs following an ACL configuration change, enter either of the following commands at the global CONFIG level of the CLI.

```
device(config)#
ipv6 rebind-acl
```

Syntax: [no] ip rebind-acl num | name

```
device(config)#
ipv6 rebind-all-acl
```

Syntax: [no] ip rebind-all-acl

Applying an IPv6 ACL to a VRF Interface

A VRF interface can be one physical port, a virtual interface, or a trunk port consisting of multiple physical ports. As with regular IPv6 ports, you can apply an inbound or outbound IPv6 ACL to a VRF interface to filter incoming and outgoing traffic respectively. This type of ACL is called a IPv6 VRF ACL.

Distinction between IPv6 ACLs applied to regular and VRF interfaces

IPv6 ACLs (both inbound and outbound) can only be applied at the IPv6 interface-level, which may be a physical or a virtual interface. If a physical port is a member of one or more virtual interfaces, the IPv6 ACL must be bound at the corresponding ve level (not at the physical port level). You cannot change the VLAN membership of a physical port with an IPv6 ACL.

When an IPv6 VRF is dynamically configured on an interface port, all IPv6 addresses on that interface are deleted. IPv6 ACL binding on the interface is not be cleared because IPv6 ACL programming is independent of the VRF membership of the interface.

To apply an IPv6 ACL, for example "access1", to a VRF interface, enter commands such as the following.

device (config)# vif 20 device (config-vif-20)#ipv6 traffic-filter access1 in

Syntax: [no] ipv6 traffic-filter ipv6-acl-name in | out

For the *ipv6-acl-name* parameter, specify the name of an IPv6 ACL created using the **ipv6 access-list** command.

The **in** keyword applies the specified IPv6 ACL to incoming IPv6 packets on the Extreme device interface.

The **out** keyword applies the specified IPv6 ACL to outgoing IPv6 packets on the Extreme device interface.

Controlling access to a device

You can use an IPv6 ACL to filter control incoming and outgoing connections to and from an Extreme device. To do so, you must create an ACL and then specify the sequence in which the ACL is applied to incoming or outgoing connections to the Extreme device.

For example, to permit incoming connections from remote hosts (2000:2383:e0bb::2/128 and 2000:2383:e0bb::3/128) to an Extreme device (30ff:3782::ff89/128), enter the following commands.

```
device(config)# ipv6 access-list remote-hosts permit 2000:2383:e0bb::2/128 30ff:3782::ff89/128 sequence 10
device(config)# ipv6 access-list remote-hosts permit 2000:2383:e0bb::3/128 30ff:3782::ff89/128 sequence 20
device(config)# ipv6 access-class remote-hosts in
```

Because of the implicit deny command at the end of each IPv6 ACL, the Extreme device denies incoming connections from all other IPv6 hosts.

> **NOTE**
> The **ipv6 access-class** command is applicable only to traffic coming in or going out the management
> port.

Syntax: [no] ipv6 access-list name deny | permit ipv6-source-prefix/prefix-length | any ipv6-destination-prefix/prefix-length | any [sequence number]

The **sequence***number* parameter specifies the order in which a statement appears in an IPv6 ACL and is therefore applied to a request. You can specify a value from 0 - 4294967295.

Adding a comment to an IPv6 ACL entry

You can optionally add a comment to describe entries in an IPv6 ACL. The comment appears in the output of **show** commands that display ACL information.

You can add a comment by entering the **remark** command immediately preceding an ACL entry, or specify the ACL entry to which the comment applies.

For example, to enter comments for preceding an ACL entry, enter commands such as the following.

```
device(config)#ipv6 access-list rtr
device(config-ipv6-access-list rtr)# remark This entry permits ipv6 packets from 3002::2 to any destination
device(config-ipv6-access-list rtr)# permit ipv6 host 3000::2 any
device(config-ipv6-access-list rtr)# remark This entry denies udp packets from any source to any destination
device(config-ipv6-access-list rtr)# deny udp any any
device(config-ipv6-access-list rtr)# remark This entry denies IPv6 packets from any source to any
destination
device(config-ipv6-access-list rtr)# deny ipv6 any any
device(config-ipv6-access-list rtr)# write memory
```

In the following example, remarks are entered immediately preceding ACL entries that specify sequence numbers.

```
device(config)# ipv6 access-list ipv6_acl
device(config-ipv6-access-list-ipv6_acl)# remark test-entry
device(config-ipv6-access-list-ipv6_acl)# deny sctp any any sequence 1
device(config-ipv6-access-list-ipv6_acl)# remark-entry sequence 5 test_acl
device(config-ipv6-access-list-ipv6_acl)# permit esp 2::/64 any sequence 5
device(config-ipv6-access-list-ipv6_acl)# remark test_remark
```

device(config-ipv6-access-list-ipv6_acl)# deny ipv6 any any sequence 23

Syntax: [no] remark comment-text

The *comment-text* can be up to 256 characters in length.

The **remark** command provisions a default comment. Only one default comment is maintained; it is overwritten by any subsequent **remark** command. The default remark is associated with the next provisioned filter as follows:

- If the immediately following filter is provisioned without a sequence number, the system assigns a default sequence number:
  - And a remark for this system-assigned sequence number already exists, then the filter gets associated with that remark and default remark remains unused.
  - And a remark for this system-assigned sequence number does not exist, then the default remark gets associated with the filter.
- If the immediately following filter is provisioned with a sequence number:
  - And a remark for this sequence number already exists, then the filter gets associated with that remark and default remark remains unused.
  - And a remark for this sequence number does not exist, then the default remark gets associated with the filter.
- Once the default remark gets associated with a filter:
  - It gets the same sequence number as the filter.
  - You can provision another default remark which may be used by another filter.

To apply a comment to a specific ACL entry, specify the ACL's entry number with the **remark-entry sequence** command. Use the **show ipv6 access-list** command to list ACL entry number. Enter commands such as the following.

```
device(config)# ipv6 access-list netw
device(config-ipv6-access-list netw) remark-entry sequence 10 This entry permits ipv6 packets from 3000::2
to any destination
device(config-ipv6-access-list netw)# remark-entry sequence 20 This entry denies UDP packets from any
source to any destination
device(config-ipv6-access-list netw)# remark-entry sequence 30 This entry denies IPv6 packets from any
source to any destination
```

Syntax: [no] remark-entry sequence sequence number comment-text

The *sequence number* is the line number assigned to the ACL entry. For a list of ACL entry numbers, use the **show ipv6 access-list** command.

The *comment-text* can be up to 256 characters in length. The comment must be entered separately from the actual ACL entry; that is, you cannot enter the ACL entry and the ACL comment with the same command.

You can use the **show running-config** or **show ipv6 access-list** commands to display IPv6 ACLs and comments.

The following shows the comment text for the ACL named "rtr" in a **show running-config** display.

```
device# show running-config
ipv6 access-list rtr
 remark This entry permits ipv6 packets from 3002::2 to any destination
 permit ipv6 host 3000::2 any
 remark This entry denies udp packets from any source to any destination
 deny udp any any
 remark This entry denies IPv6 packets from any source to any destination
 deny ipv6 any any
```

Syntax: show running-config

> **NOTE**
> If "suppress-acl-seq" is ON; All unused "remark-entry" statements will be hidden while the **running-config** is displayed or stored.If "suppress-acl-seq" is ON; All used "remark-entry" statements will be displayed as "remark" statements while the **running-config** is displayed or stored.

The following example shows the comment text for the ACL named "rtr" in a **show ipv6 access-list** display.

```
device# show ipv6 access-list rtr
ipv6 access-list rtr: 3 entries
 10: remark This entry permits ipv6 packets from 3002::2 to any destination
 10: permit ipv6 host 3000::2 any
 20: remark This entry denies udp packets from any source to any destination
 20: deny udp any any
 30: remark This entry denies IPv6 packets from any source to any destination
 30: deny ipv6 any any
```

The following example shows the comment text for the ACL named "ipv6_acl".

```
device(config)# sh ipv6 access-list ipv6_acl
ipv6 access-list ipv6_acl: 3 entries
 1: remark test-entry
 1: deny sctp any any sequence 1
 5: remark-entry sequence 5 test_acl
 5: permit esp 2::/64 any sequence 5
 23:remark test_remark
```

23: deny ipv6 any any sequence 23

Syntax: show ipv6 access-list [**access-list-name**]

For the *access-list-name* parameter, specify the name of an IPv6 ACL created using the **ipv6 access-list** command.

ACL CAM sharing for inbound IPv6 ACLs

ACL CAM sharing allows you to conserve CAM by sharing it between ports that are supported by the same packet processor (PPCR). If this feature is enabled globally, you can share CAM space that is allocated for inbound ACLs between instances on ports that share the same packet processor (PPCR). For example, if you have bound- inbound ACL 101 to ports 1/1 and 1/5, the ACL is stored in a single location in CAM and used by both ports. Table 10 describes which ports share PPCRs and can participate in ACL CAM sharing.

Module type	PPCR number	Ports supported by PPCR
20 x 1G	PPCR 1	1 - 20
4 x 10G	PPCR 1	1 - 2
	PPCR 2	3 - 4
2 x 10G	PPCR 1	1 - 2

Considerations when implementing this feature

The following consideration apply when implementing this feature:

- If you enable ACL CAM sharing, ACL statistics will be generated per-PPCR instead of per-port. If you require the statistics per-port granularity for your application, you cannot use this feature.
- This feature cannot be applied to a virtual interface.
- CAM entry matching within this feature is based on the ACL group ID.
- This feature cannot co-exist with IP Multicast Routing or IP Multicast Traffic Reduction.

Configuring ACL CAM sharing for IPv6 ACLs

When enabled, ACL CAM sharing for IPv6 inbound ACLs is applied across all ports in a system. To apply ACL CAM sharing for IPv6 ACLs on an Extreme device, use the following command.

```
device(config)# acl-policy
device(config-acl-policy)# ipv6 enable-acl-cam-sharing
```

Syntax: ipv6 enable-acl-cam-sharing

Filtering and priority manipulation based on 802.1p priority

Filtering and priority manipulation based on a packet's 801.1p priority is supported in the Extreme devices through the following QoS options:

- **priority-force** – Assigns packets of outgoing traffic that match the ACL to a specific hardware forwarding queue, even though the incoming packet may be assigned to another queue. Specify one of the following QoS queues:

  - 0 - qosp0
  - 1 - qosp1
  - 2 - qosp2
  - 3 - qosp3
  - 4 - qosp4
  - 5 - qosp5
  - 6 - qosp6
  - 7 - qosp7

If a packet's 802.1p value is forced to another value by its assignment to a lower value queue, it will retain that value when it is sent out through the outbound port.

The default behavior on previous revisions of this feature was to send the packet out with the higher of two possible values: the initial 802.1p value that the packet arrived with or the new (higher) priority that the packet has been "forced" to.

- **priority-mapping** – Matches on the packet's 802.1p value. This option does not change the packet's forwarding priority through the device or mark the packet.

Example using the priority force option

In the following IPv6 ACL example, access list acl1 assigns tcp packets with the source address specified and any destination address to the internal priority 7.

```
device(config)# ipv6 access-list acl1
device(config-ipv6-access-list acl1)# permit tcp 4000:1::/64 any priority-force 7
```

The **priority-force** parameter specifies one of the 8 internal priorities of the Extreme device. Possible values are between 0 and 7.

ACL accounting

Multi-Service devices monitor the number of times an ACL is used to filter incoming or outgoing traffic on an interface. The **show ipv6 access-list accounting** command displays the number of "hits" or how many times ACL filters permitted or denied packets that matched the conditions of the filters.

> **NOTE**
> ACL accounting does not tabulate nor display the number of implicit denials by an ACL.

Counters, stored in hardware, keep track of the number of times an ACL filter is used.

The counters that are displayed on the ACL accounting report are:

- **1s** - Number of hits during the last second. This counter is updated every second.
- **1m** - Number of hits during the last minute. This counter is updated every one minute.
- **5m** - Number of hits during the last five minutes. This counter is updated every five minutes.
- **ac** - Accumulated total number of hits. This counter begins when an ACL is bound to an interface and is updated every one minute. This total is updated until it is cleared.

The accumulated total is updated every minute. For example, a minute after an ACL is bound to a port, it receives 10 hits per second and continues to receive 10 hits per second. After one minute, the accumulated total hits is 600. After 10 minutes, there will be 6000 hits.

The counters can be cleared when the device is rebooted, when an ACL is bound to or unbound from an interface, or by entering a **clear ipv6 access-list** command.

Enabling and Disabling ACL accounting on XMR Series and MLX Series devices

ACL accounting is disabled by default on XMR Series and MLX Series devices. To enable ACL accounting, enter the following command:

```
device(config)# acl-policy
device(config-acl-policy)# enable-acl-counter
```

Syntax: **[no] enable-acl-counter**

> NOTE
> Enabling or disabling ACL accounting affects the gathering of statistics from all ACL types (Layer-2, IPv4 and IPv6).

> NOTE
> The **enable-acl-counter** command is not supported on CES 2000 Series and CER 2000 Series devices.

ACL accounting on CES 2000 Series and CER 2000 Series devices

The following special considerations affect how IPv6 Layer 4 ACL accounting is configured on the CES 2000 Series and CER 2000 Series devices:

- You can enable ACL accounting at the filter level by adding an **enable-accounting** keyword in each clause of an IPv6 ACL for which you want to gather statistics.
- IPv6 ACL rate limiting and IPv6 deny logging are not supported.
- CAM resources are shared on the devices between Layer 2, IPv4, and IPv6 ACL accounting. This limits the number of ACL accounting instances available on the system.
- For inbound ACL accounting, you can bind a Layer 2, IPv4, and IPv6 ACL accounting to the same port. Refer to "Configuration considerations for dual inbound ACLS on CES 2000 Series and CER 2000 Series devices" and "ACL Accounting interactions between L2 ACLs and IP ACLs" for further information.
- For outbound ACL accounting, you can bind an IPv4 and IPv6 ACL accounting to the same port. However, Layer 2 ACL accounting does not coexist with either IPv4 or IPv6 ACL accounting on the same port.
- The port-level configuration to enable or disable the counters is not applicable.

For detailed information about ACL accounting considerations for CES 2000 Series and CER 2000 Series devices, refer to "ACL accounting".

Enabling and disabling IPv6 ACL accounting on CES 2000 Series 2000 and CER 2000 Series 2000 devices

By default, the ACL accounting is disabled on the CES 2000 Series 2000 and CER 2000 Series 2000 devices. You can enable the ACL accounting explicitly in each clause of an IPv6 ACL for which you want to gather statistics by including the keyword **enable-accounting** immediately after the **permit** or **deny** keyword. To enable ACL accounting, enter the following command:

```
device(config)# ipv6 access-list netw permit enable-accounting ip any any
```

Syntax: **[no] ipv6 access-list** *aclname* **permit | deny enable-accounting**

The *acl name* variable defines the name of the IPv6 ACL. The *acl name* can contain up to 199 characters and numbers, but cannot begin with a number and cannot include any spaces or quotation marks.

The **permit** keyword indicates that enabling IPv6 ACL accounting will be permitted for the clauses that match a policy in the access list.

The **deny** keyword indicates that enabling IPv6 ACL accounting will be denied for the clauses that match a policy in the access list.

The **enable-accounting** keyword enables the IPv6 ACL accounting.

The **no** option is used to turn off the previously enabled IPv6 ACL accounting.

> **NOTE**
> The rules of action merging and counter precedence must be considered to determine which action to take and which accounting to count while binding multiple ACL accountings to the same port.

Displaying statistics for IPv6 ACL accounting

To display statistics for IPv6 accounting, enter commands such as the following.

```
device# show ipv6 access-list accounting brief
Collecting IPv6 ACL accounting summary for 1/26 ...  Completed successfully.
Collecting IPv6 ACL accounting summary for 1/25 ...  Completed successfully.
Int          In ACL                    Total   In Hit                Out ACL              Total Out
Hit
1/26          ipv6-port2                   450375(1s)
                                                  27009259(1m)
                                                  135046361(5m)
                                                  464353218(ac)
1/25                                              ipv6-port            2234540(1s)

          11321736(1m)

          0(5m)

          11321736(ac)
```

The table below describes the output parameters of the **show ipv6 access-list accounting brief** command.

TABLE 17

Field	Description
Collecting IPv6 ACL accounting summary for *interface*	Shows the interface for which the ACL accounting summary is collected and specifies whether or not the collection is successful.
Int	Shows the ID of the interface for which the statistics are being reported.
In ACL	Shows the ID of the ACL used to filter the incoming traffic on the interface.
Total In Hit	Shows the number of hits from the incoming traffic processed by all the ACL entries (filters) in the ACL.
Out ACL	Shows the ID of the ACL used to filter the outgoing traffic on the interface.

TABLE 17 (continued)

Field	Description
Total Out Hit1	Shows the number of hits from the outgoing traffic processed by all the ACL entries (filters) in the ACL.

Syntax: show ipv6 access-list accounting brief

Displaying IPv6 accounting statistics for an interface

To display statistics for an interface, enter commands such as the following.

```
device# show ipv6 access-list accounting ethernet 1/24 out
IPv6 ACL ipv6-protocol-ahp-deny-test-66
Collecting IPv6 ACL accounting for 1/24  ...  Completed successfully.
   10:  deny enable-accounting ahp any host 3002::10
      Hit count: (1 sec)                    2   (1
min)                                                                              15
                  (5 min)                  92
   (accum)                                                                          107
```

The table below describes the output parameters of the **show ipv6 access-list accounting ethernet** command.

TABLE 18

Field	Description
IPv6 ACL	Shows the name of the IPv6 traffic filter for the collected statistics.
Collecting IPv6 ACL accounting for *interface*	Shows the interface for which the ACL accounting information is collected and specifies whether or not the collection is successful.
#	Shows the index of the IPv6 ACL entry, starting with 0, followed by the permit or deny condition defined for that ACL entry. (The first entry created for an ACL is assigned the index 0. The index of the subsequent entries created are incremented by 1.)
deny enable-accounting ahp	Shows the name of the matching clause in the ACL.
Hit count	Shows the number of matching frames for each sample interval and the accumulated value since the last clear or rebind action.

Syntax: show ipv6 access-list accounting ethernet [slot/port | ve ve-number] in | out [policy-based-routing]

Use **ethernet***slot*/*port* to display a report for a physical interface.

Use **ve** *ve-number* to display a report for the ports that are included in a virtual routing interface. For example, if ports 1/2, 1/4, and 1/6 are all members of ve 2, the report includes information for all three ports.

Use the **in** parameter to display statistics for incoming traffic; **out** for outgoing traffic.

The **policy-based-routing** parameter limits the display to policy based routing accounting information.

Clearing the ACL statistics

Statistics on the ACL account report can be cleared:

- When a software reload occurs

- When the ACL is bound to or unbound from an interface

- When you enter the **clear ipv6 access-list** command, as in the following example.

```
device(config)# clear ipv6 access-list all
```

Syntax: clear ipv6 access-list all | ethernet slot/port | ve ve-num

Enter **all** to clear all statistics for all ACLs.

Use **ethernet**slot/port to clear statistics for ACLs a physical port.

Use **ve**ve-number to clear statistics for all ACLs bound to ports that are members of a virtual routing interface.

> **NOTE**
> Because IPv6 rate limiting is not supported on the CES 2000 Series 2000 and CER 2000 Series 2000 devices, the counts displayed in the accounting mode represent the number of packets that matched the IPv6 ACL.

IPv6 receive ACLs

This section discusses the following topics:

IPv6 receive ACLs overview

The IPv6 receive access-control list feature (rACL) provides hardware-based filtering capability for IPv6 traffic, destined for the CPU in the default VRF, such as management traffic. Its purpose is to protect the management module's CPU from overloading due to large amounts of traffic sent to one of the Extreme device's IP interfaces. The rACL feature applies the specified ACL to every interface on the Extreme device. This eliminates the need to add an ACL to each interface on an Extreme device.

The rACL feature is configured by creating an ACL to filter traffic and then specifying that ACL in the **ipv6 receive access-list** command. This applies the ACL to all interfaces on the device. The destination IP address in an ACL specified by the rACL command is interpreted to apply to all interfaces in the default VRF of the device. This is implemented by programming an ACL entry in CAM that applies the ACL clause for each interface.

You can bind multiple IPv6 rACLs, up to a maximum of 50.

CAM entries are programmed differently for Gen-1 and Gen-2 interface modules. For Gen-1 interface modules, each rule is programmed for every IPv6 address interface so that the number of CAM entries for each rule is equivalent to the number of interface addresses. For example, if "M" IPv6 address interfaces are configured and there are "N" rACL rules, then there will be "M x N" CAM entries in the IPv6 rACL CAM partition.

> **NOTE**
> The rACL feature does not program CAM entries on Gen-1 interface modules when an IPv6 interface is in the down state.

For Gen-2 interface modules, one rule is programmed for all local IPv6 address interfaces. Hence, if there are "N" IPv6 rACL rules to program, there will be "N" CAM entries in IPv6 rACL CAM partition.

The IPv6 rACL feature supports mirroring and sflow for traffic filtered by IPv6 rACL.

Configuration guidelines for IPv6 receive ACLs considerations

Use the following considerations when configuring IPv6 receive ACLs (rACLs):

- IPv6 rACLs support was introduced in *Multi-Service IronWare R05.6.00*. For backward compatibility, the IPv6 rACL sub-partition is set to "0" by default, so that other sub-partitions of the IPv6 CAM partition are not affected by this feature when the firmware is upgraded.

- After an upgrade to *Multi-Service IronWare R05.6.00* or later, the sub-partition size for IPv6 rACL will be "0". Refer to Specifiying the maximum number of rACLs supported in CAM on page 174 for more information about changing the default value.

- After a downgrade to a previous release, all configured IPv6 rACLs will be lost.

- An explicit **deny ip any any** filter does not match any multicast traffic. Since the destination field in the ACL contains "any", rACLs program interface IP addresses in the CAM instead of the destination's IP address. To match the multicast traffic, you should specify the multicast group address in the destination field in the ACL.

- (MLX Series and XMR Series devices only) For traffic destined to VRRP/E router instances, enable the **racl-vrrp-vrip-filter ipv6-packets** command. For details, refer to Summary of ACL system and policy parameters on page 181.

Interfaces on which a UDA or Layer 2 ACL is also applied

M-CPU-bound Layer 3 Broadcast, Unicast, or Multicast packets permitted by the UDA or Layer 2 ACL are subject to the following conditions:

- If the packet matches an rACL permit rule, the packet is forwarded to the M-CPU.
- If the packet matches an rACL deny rule, the packet is dropped.
- The default implicit rACL rule is deny. Packets that do not match any configured permit statement are dropped.

Interfaces on which no UDA or Layer 2 ACL is applied

If neither an ingress Layer 2 ACL nor a user-defined ACL (UDA) is applied on an interface, by default, rACLs are not supported for filtering CPU-bound traffic through that interface. To support rACLs on such interfaces, enable the **racl-cpu-filtering ipv6-packets** command at device level:

```
device(config)# acl-policy
device(config-acl-policy)# racl-cpu-filtering ipv6-packets
```

If **racl-cpu-filtering ip-packets** is enabled and rACLs are applied on a device, guidelines for Layer 3 Unicast packets are as follows:

- If the destination address in the incoming packet is a connected IPv4/IPv6 address (physical/VE Interface or loopback), the decision to permit or deny the packet is taken in hardware—under CAM programming—subject to the following conditions:
  - Packets matching a configured rACL permit rule are forwarded to the M-CPU for processing.
  - Packets that do not match any rACL permit rules are dropped in hardware, by a default deny rule applicable if the destination is a connected IPv4/IPv6 address.
- If the destination is not a connected IPv4/IPv6 address, the packet is not destined for M-CPU, and is not subject to rACL CPU filtering.

If **racl-cpu-filtering ipv6-packets** is enabled and rACLs are applied on a device, guidelines for Layer 3 Broadcast or Multicast packets are as follows:

- Packets that match an rACL permit rule are forwarded to the M-CPU.
- Packets that match an rACL deny rule are dropped.

- The default implicit rACL rule is permit; packets that do not match any rACL deny rule are permitted. (A default implicit permit enables protocol packets to be processed by the M-CPU. Multicast-based protocols such as OSPF, LDP, RIP, VRRP, and PIM work seamlessly, even when rACL CPU filtering is configured.)
- If you add an explicit "deny ip/ipv6 any any" rule to an rACL, make sure to add permit rules for the relevant multicast-protocol destination addresses.

IPv6 rACL example for multicast protocol

In the following example, racl-cpu-filtering is enabled. If you conclude the rACL with a "deny any any" rule (in order for IPv6 OSPFv2 peering to work seamlessly) you also need to include the following permit rules:

- Unicast peer IPv6 address
- Destination multicast addresses: FF02::5 and FF02::6
- Solicited Node Multicast address: FF02::1:FFXX:XXXX
- Link Local IPv6 address: FE80::/64

IPv6 receive ACL prerequisites

Specifiying the maximum number of rACLs supported in CAM

By default, the IPv6 rACL sub-partition in an IPv6 session CAM partition is set to "0". This must be resized before using the IPv6 receive ACL feature.

An IPv6 session CAM partition has sub-partitions for:

- IPv6 Multicast
- Receive ACL
- Rule-based ACL

The IPv6 Multicast sub-partition is configured using the **system-max ipv6-mcast-cam** command. The Receive ACL sub-partition is configured using the **system-max ipv6-receive-cam** command. The number of IPv6 Rule-based ACL entries is normalized after allocating space for IPv6 Multicast and IPv6 Receive ACL entries i.e. IPv6 Rule-based ACL entries take the remaining space after the allocation of IPv6 Multicast and IPv6 rACL entries. However, the system ensures that there are a minimum of 128 IPv6 Rule-based ACL entries.

When you set the size of the Receive ACL sub-partition using the **system-max ipv6-receive-cam** command, the size of the Rule-based ACL sub-partition is decreased. The following example shows how configure the Receive ACL sub-partition assuming that the IPv6 session CAM partition is initially configured as follows:

```
[IPV6 Session]    16384(size),   16384(free), 000.00%(used)
  :IPv6 Multicast:  1024(size),    1024(free), 000.00%(used)
  :Receive ACL:        0(size),       0(free), 000.00%(used)
  :Rule ACL:       15360(size),   15360(free), 000.00%(used)
```

Use the following command to set the maximum IPv6 Receive ACL entries to 2048.

```
device(config)# system-max ipv6-receive-cam 2048
```

Syntax: **[no] system-max ipv6-receive-cam** *num*

The *num* variable specifies the maximum number of IPv6 Receive ACL entries allowed in CAM. Acceptable values are powers of 2 in the range from 0 through 8192. The default value is 0. If you enter a value that is not a power of 2, the system rounds off the entry to a

number less than the input value. For example, if you enter 2044, which is not a power of 2, the system rounds it down to 1024 and shows an appropriate warning.

```
device(config)# system-max ipv6-receive-cam 2044
Warning - IPv6 Receive ACL CAM size requires power of 2, round down to 1024
Reload required. Please write memory and then reload or power cycle the system.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

The **no** form of the **system-max ipv6-receive-cam** command restores the default value.

> **NOTE**
> You must write this command to memory and perform a system reload for this command to take effect.

Setting IPv6 Receive ACL to 2048 decreases the size of the Rule ACL sub-partition so that the IPv6 session CAM partition profile is now:

```
[IPV6 Session]    16384(size),   16384(free), 000.00%(used)
   :IPv6 Multicast: 1024(size),    1024(free), 000.00%(used)
   :Receive ACL:    2048(size),    2048(free), 000.00%(used)
   :Rule ACL:          13312(size),   13312(free), 000.00%(used)
```

The Rule-based sub-partition size can be increased again by reducing the size of the IPv6 Multicast sub-partition. Use the following command to reduce the size of the IP Multicast sub-partition to 0.

```
device(config)# system-max ipv6-mcast-cam 0
Reload required. Please write memory and then reload or power cycle the system.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

Syntax: [no] system-max ipv6-mcast-cam *num*

The *num* variable specifies the maximum number of IPv6 Multicast entries allowed in CAM. The range is from 0 through 8192 . The default value is 1024.

> **NOTE**
> You must write this command to memory and perform a system reload for this command to take effect.

Changing the size of IPv6 Multicast from 1024 to 0 increases the maximum IPv6 Rule-based ACL entries by 1024 so that the CAM partition profile is:

```
[IPV6 Session]    16384(size),   16384(free), 000.00%(used)
   :IPv6 Multicast:    0(size),       0(free), 000.00%(used)
   :Receive ACL:    2048(size),    2048(free), 000.00%(used)
   :Rule ACL:          14336(size),   14336(free), 000.00%(used)
```

Maximum supported size of IPv6 Receive ACLs in CAM profiles

IPv6 rACLs are supported in the following CAM profiles which have space allocated for an IPv6 partition:

- Default
- IPv6
- IPv6 + IPv4
- IPv6 + IPv4-2
- Multi-Service
- Multi-Service-2

- Multi-Service-3
- Multi-Service-4

The table below shows the maximum supported size of IPv6 rACLs in supported CAM profiles when the size of the IPv6 Multicast sub-partition is 0. The maximum number of IPv6 rACL entries will vary when the number of IPv6 Multicast entries is not 0.

TABLE 19

Profile	Supported	IPv6 rACL size	Rule ACL Size	IPv6 Multicast
Default	Y	1024	3072	0
IPv4 Optimized	N			
IPv6 Optimized	Y	8192	16384	0
MPLS VPN Optimized	N			
MPLS VPLS Optimized	N			
L2 Metro Optimized	N			
L2 Metro Optimized #2	N			
MPLS VPN Optimized #2	N			
MPLS VPLS Optimized #2	N			
Multi-Service	Y	2048	6144	0
MPLS VPN+VPLS	N			
IPv4 + VPN	N			
IPv6 + IPv4	Y	8192	16384	0
IPv4 + VPLS	N			
IPv4 + Ipv6 2	Y	2048	6144	0
Multi-service 2	Y	1024	3072	0
Multi-service 3	Y	2048	6144	0
Multi-service 4	Y	2048	6144	0

NOTE
The table above shows the maximum supported IPv6 rACL entries for all supported XMR cards including legacy Gen-1 cards, 8x10G-M/X/D, 2x100G and 24x10G.

Checking for available space when configuring the IPv6 rACL sub-partition

Before allocating more space to the Receive ACL sub-partition, the system will check if there is enough space in the IPv6 ACL CAM partition. If there is not enough space, an error message is displayed.

```
device(config)# system-max ipv6-receive-cam 2048
Error - IPV6 Receive ACL CAM (2048) exceeding available CAM resources
Total IPv6 ACL CAM:            16384(Raw Size)
    Reserved IPv6 Rule ACL CAM:  1024(Raw Size)
    IPv6 Multicast CAM:             0(Raw Size)
    Available IPv6 Receive ACL CAM: 15360(Raw Size) 1920(User Size)
```

This error message shows that there are 1920 available user entries for IPv6 rACL CAM. This example uses the default CAM profile which supports a maximum of 1024 IPv6 rACL CAM entries. Refer to the table above for information on the maximum supported size of IPv6 Receive ACLs in different CAM profiles.

The **system-max ipv6-receive-cam** command executes successfully to set the maximum number of IPv6 Receive ACLs to 1024 for the default CAM profile. For example:

```
device(config)# system-max ipv6-receive-cam 1024
Reload required. Please write memory and then reload or power cycle the system.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

Checking for available space when changing the CAM profile

The system will check if there is enough space for the IPv6 Receive ACL sub-partition before changing the CAM profile. If there is not enough space, an error message is displayed.

```
device(config)# system-max ipv6-receive-cam 2048
device(config)# cam-partition profile multi-service-2
Error - IPV6 Receive ACL CAM (2048) exceeding available CAM resources
Total IPv6 ACL CAM:            16384(Raw Size)
    Reserved IPv6 Rule ACL CAM:   1024(Raw Size)
    IPv6 Multicast CAM:      0(Raw Size)
    Available IPv6 Receive ACL CAM: 15360(Raw Size) 1920(User Size)
```

This error message shows that there are 1920 available user entries for IPv6 rACL CAM. In this case, the CAM profile is "multi-service 2" which supports 1024 IPv6 rACL CAM entries. Refer to the table above for on the maximum supported size of IPv6 Receive ACLs in different CAM profiles. Use the following command to change the CAM profile to the "multi-service 3" which supports 2048 IPv6 rACLs.

```
device(config)# cam-partition profile multi-service-3
Reload required. Please write memory and then reload or power cycle the system.
```

IPv6 receive ACL: basic configuration

Configuring and applying an IPv6 rACL

Configuring IPv6 rACLs requires the following steps:

1. Configuring an IPv6 rACL sub-partition in the CAM partition on page 177
2. Creating an IPv6 access-list on page 177
3. Creating a policy-map on page 178 (if you want to rate limit traffic)
4. Applying an IPv6 rACL on page 178

Configuring an IPv6 rACL sub-partition in the CAM partition

To create an IPv6 rACL sub-partition and set the maximum number of IPv6 rACL entries at 1024, use the following commands.

```
device(config)# system-max ipv6-receive-cam 1024
device(config)# write memory
device(config)# reload
```

Creating an IPv6 access-list

To create an IPv6 access-list named "b1":

```
device(config)# ipv6 access-list b1
device(config-ipv6-access-list b1)# permit ipv6 any any
device(config-ipv6-access-list b1)# exit
```

Creating a policy-map

To create a policy map "m1" to rate-limit traffic:

```
device(config)# policy-map m1
device(config-policymap m1)# cir 1000000 cbs 2000000
device(config-policymap m1)# exit
```

Applying an IPv6 rACL

To configure IPv6 rACL to apply IPv6 access-list "b1" with a sequence number "15" to all interfaces within the default VRF for all CPU-bound traffic, enter the following command:

```
device(config)# ipv6 receive access-list b1 sequence 15
```

To configure IPv6 rACL to apply IPv6 access-list "b1" with a sequence number "15" with a policy-map "m1", enter the following command:

```
device(config)# ipv6 receive access-list b1 sequence 15 policy-map m1
```

To configure IPv6 rACL to apply IPv6 access-list "b1" with a sequence number "15" and a policy-map "m1" with strict -acl, enter the following command:

```
device(config)# ipv6 receive access-list b1 sequence 15 policy-map m1 strict-acl
```

Syntax: [no] ipv6 receive access-list *acl-name* sequence seq-num [policy-map policy-map-name [strict-acl]]

The *acl-name* variable specifies the name of the access-control list to apply to all interfaces within the default VRF, for all CPU-bound traffic. The maximum length is 256 characters.

The **sequence** *seq-num* variable specifies the sequence number of the access-control list being applied as a rACL. IPv6 rACL commands are applied in the order of the lowest to the highest sequence numbers. The range of values is from 1 through 50.

The **policy-map** *policy-map-name* variable specifies the name of a policy map. When the **policy-map** option is specified, traffic matching the "permit "clause of the specified IPv6 ACL is rate-limited as defined in the policy map and IPv6 traffic matching the "deny" clause in the IPv6 ACL is permitted without any rate limiting.

The **strict-acl** parameter specifies that traffic matching the "permit" clause of the specified IPv6 ACL is rate-limited as defined in the policy map and IPv6 traffic matching the "deny" clause in the IPv6 ACL is dropped in the hardware.

Rebinding an IPv6 rACL definition or policy-map

When access list rules are modified or a policy map associated with a rACL is changed, an explicit rebind must be performed to propogate the changes to the interfaces. To rebind an IPv6 access-control list, enter the following command:

```
device(config)# ipv6 receive rebind-acl-all
```

Syntax: ipv6 receive rebind-acl-all

Displaying access-list binding information

To display all IPv6 access-lists (both rule-based and rACL) that are bound to different interfaces, enter the following command:

```
device(config)# show ipv6 access-list bindings
!
ipv6 receive access-list b1 sequence 11
ipv6 receive access-list b2 sequence 12
!
```

Syntax: show ipv6 access-list bindings

Deactivating the IPv6 rACL configuration

To deactivate the IPv6 rACL configuration and remove all the rules from CAM, enter the following command.

```
device(config)# ipv6 receive deactivate-acl-all
```

Syntax: [no] ipv6 receive deactivate-acl-all

The **no** form of this command reactivates the IPv6 rACL configuration.

> **NOTE**
> To make this change permanent and prevent ACL binding to CAM after reload, you must save the configuration.

Deleting the IPv6 rACL configuration

To delete the IPv6 rACL configuration and remove all IPv6 rACL rules from the system, use the following command.

```
device(config)# ipv6 receive delete-acl-all
This command deletes all IP Receive ACLs from system.
Are you sure? (enter 'y' or 'n'): y
```

Syntax: ipv6 receive delete-acl-all

IPv6 receive ACL: additional configuration

> **NOTE**
> For details of rACL logging, refer to Enabling L3 rACL logging on page 186.

Configuring IPv6 rACL with acl-mirror-port

You can mirror traffic coming on to an interface, to any other interface. When specifying a destination port for IPv6 rACLs, you must configure the acl-mirror-port command on all ports supported by the same packet processor (PPCR).

Configuring IPv6 rACL with acl-mirror-port requires the following steps:

1. Creating an IPv6 ACL with a mirroring clause on page 179
2. Specifying the destination mirror port for physical ports on page 179
3. Applying the IPv6 rACL on page 180

Creating an IPv6 ACL with a mirroring clause

To create a named ACL "b1" with a mirroring clause, enter the following commands:

```
device(config)# ipv6 access-list b1
device(config-ipv6-access-list b1)# permit ipv6 any any mirror
device(config-ipv6-access-list b1)# permit ipv6 any any
device(config-ipv6-access-list b1)# exit
```

Specifying the destination mirror port for physical ports

In the following example, ports "ethernet 3/1" and "ethernet 3/2" belong to the same PPCR. To specify "ethernet 5/1" as the destination mirror port for these ports, use the following commands:

```
device(config)# interface ethernet 3/1
device(config-if-e1000-3/1)# acl-mirror-port ethernet 5/1
```

```
device(config-if-e1000-3/1)#interface ethernet 3/2
device(config-if-e1000-3/2)# acl-mirror-port ethernet 5/1
```

Applying the IPv6 rACL

To apply the IPv6 rACL, enter the following command:

```
device(config)#ipv6 receive access-list b1 sequence 15
```

Configuring IPv6 rACL with copy-sflow

You can direct data matching an IPv6 ACL permit clause, to the sFlow collector by specifying "copy sflow" when creating the IPv6 ACL. Configuring IPv6 rACL with "copy-sflow" requires the following steps:

Creating an IPv6 ACL that directs traffic to the sFlow collector

To create a named ACL "b1" that directs traffic to the sFlow collector, enter the following commands:

```
device(config)#ipv6 access-list b1
device(config-ipv6-access-list b1)# permit ipv6 any any copy-sflow
device(config-ipv6-access-list b1)# exit
```

Applying the IPv6 rACL

To apply the IPv6 rACL, enter the following command:

```
device(config)#ipv6 receive access-list b1 sequence 15
```

Displaying accounting information for IPv6 rACLs

Before you can collect ACL accounting statistics,you must enable accounting for IPv6 rACL. Refer to ACL accounting on page 168 for further information on how to enable accounting using the **enable-acl-counter** command.

To display rACL accounting information for the named ACL "b1", enter the following command

```
device(config)# show ipv6 access-list receive accounting name b1
IPv6 Receive ACL Accounting Information:
IPv6 Receive ACL b1
ACL hit count for software processing (accum)                   0
HW counters:
   0: permit tcp any host 2000::2
      Hit count: (1 sec)                 0  (1 min)                0
                 (5 min)                 0  (accum)                0
   1: permit udp any host 1000::1
      Hit count: (1 sec)                 0  (1 min)                0
                 (5 min)                 0  (accum)                0
```

Syntax: show ipv6 access-list receive accounting { brief | name *acl-name* **}**

The **brief parameter** displays IPv6 rACL accounting information in brief. The **name** *acl-name* variable specifies the name of a receive access-control list. The maximum name length is 256 characters.

To clear IPv6 rACL accounting information for an ACL named "acl_ext1", use the following command.

```
device(config)# clear ipv6 access-list receive name acl_ext1
```

To clear accounting statistics for all configured IPv6 rACLs, enter the following command.

```
device(config)# clear ipv6 access-list receive all
```

Syntax: clear ipv6 access-list receive (all | name *acl-name* }

The **all** parameter specifies clearing accounting statistics for all configured IPv6 rACLs.

The **name** *acl-name* variable specifies clearing accounting statistics for the named rACL.

General ACL topics

This section contains topics that apply to more than one ACL type.

Summary of ACL system and policy parameters

There is a full range of system-max and acl-policy parameters that affect ACLs.

For more information about the commands listed in the following tables, refer to the following resources:

- Other ACL topics in the current publication
- *Extreme NetIron Command Reference*
- *Extreme NetIron Layer 2 Switching Configuration Guide*

You enter the commands in the following table from global configuration mode:

TABLE 20 ACL system-max parameters

Command	Description
system-max config-file-size	Max config file size. Large ACL configurations may exceed the default system-max config-file-size.
system-max ip-filter-sys	Max IPv4 ACL filter rules. This is a shared resource for all IPv6 ACL filter definitions.
system-max ipv6-receive-cam	Max IPv6 Receive ACL SW CAM rules
system-max l2-acl-table-entries	Max L2 ACL rules per ACL table
system-max mgmt-port-acl-size	Max size for management port ACL
system-max session-limit	Max ACL active ACL session resource
system-max subnet-broadcast-acl-cam	Max IP Broadcast ACL SW CAM rules
system-max uda-acl-table-entries	Max UDA ACL rules per ACL table

You enter the commands in the following table from ACL policy configuration mode, which you access from global configuration mode:

```
device(config)# acl-policy
device(config-acl-policy)#
```

TABLE 21 ACL policy parameters

Command	Description
accounting-no-sort	Display IPv4 ACL accounting in configuration order and not sorted order.
acl-duplication-check	Perform ACL filter duplication check.
acl-frag-conservative	Conservative mode for ACL filtering of fragmented packets.
acl-skip-boot-checks	Disable ACL conflict and duplication checks during boot and tftp copy.
clear	Clear table/statistics/keys.

TABLE 21 ACL policy parameters (continued)

Command	Description
cls	Clear screen.
disable-acl-for-6to4	Disable 6to4 packet processing for l4.
disable-acl-for-gre	Disable gre packet processing for l4.
display-config-format	Display in ACL configured format.
display-def-acl-seq	Display in ACL configured format with sequence information.
display-pkt-bit-rate	Display Packet Rate and Bit Rates, PBR Only.
enable-acl-cam-sharing	Enable ACL CAM sharing.
enable-acl-counter	Enable access list counters.
force-delete-bound-acl	Enable forced deletion of ACLs bound to an interface.
max-uda-offset	Enable 124 bytes as the UDA offset values (rather than the default 116 bytes).
racl-cpu-filtering	Enable rACL filtering for CPU-bound traffic—even if neither Layer 2 ACLs nor User-defined ACLs (UDAs) are configured on the incoming interface.
racl-vrrp-vrip-filter	Enable rACL filtering for VRRP/E router instances.
statistics-load-interval	Configure Statistics load interval, PBR Only.
suppress-acl-seq	Hide Sequence information from ACL running config output.
suppress-ipv6-priority-mapping	Hide priority-mapping keyword from ACL running config output.

Layer 3 ACL logging

ACL logs can provide insight into permitted and denied network traffic.

This section covers logging for Layer 3 ACLs. For Layer 2 ACLs, refer to Configuring ACL Deny Logging for Layer-2 inbound ACLs on page 94.

ACL logging platform-support matrix

The following table indicates on which platforms ACL logging is supported for the various ACL and filter types.

TABLE 22 ACL logging platform-support matrix

ACL type	Rule type	CER/CES	MLX/XMR
L2	permit	N	N
L2	deny	Y	Y
IPv4	permit	N	Y
IPv4	deny	Y	Y
IPv6	permit	N	Y
IPv6	deny	N	Y
IPv4 rACL	permit	N	Y
IPv4 rACL	deny	N	Y
IPv6 rACL	permit	N	Y
IPv6 rACL	deny	N	Y

L3 ACL-logging configuration notes

Before configuring L3 ACL permit or deny logging on your device, consider these configuration notes.

General L3 ACL-logging configuration notes

- ACL logging is a CPU-intensive feature. To maintain maximum performance, log only specific filters.
- ACL logging generates Syslog entries only. No SNMP traps are issued.
- The ACL logging feature is supported for inbound ACLs only.
- By default, ACL logs (Syslogs) are generated every five minutes.
- ACL logging is not supported for:
  - ACL-based Rate Limiting
  - Policy Based Routing
- You can configure the maximum number of ACL session entries using the **system-max session-limit** command. Only the 2/3rd of the number of sessions specified using **system-max session-limit** command are available for ACL or uRPF logging.
- The ACL or RPF logging mechanism on the Interface modules log a maximum of 256 messages per minute, and send these messages to the Management module. A rate-limiting mechanism has been added to rate-limit the number of log messages from the Interface module CPU to the Management module CPU to 5 messages per second. Because this delays the delivery of messages to the Management module, in the worst case scenario with all 256 packets arriving at the same time on the Interface module, the time values stamped by the Management module on the messages will vary by as much as 60 seconds.

Deny-filter configuration notes

- The ACL deny logging feature may be enabled with the **ip access-group redirect-deny-to-interf** command. However, if the **ip access-group enable-deny-logging** command and the **ip access-group redirect-deny-to-interf** command are configured on the same interface, a syslog entry is created for packets matching the deny action filter containing the **log** keyword, and the packet is dropped. Packets matching a **log** enabled filter are not redirected to the specified interface.The **ip access-group redirect-deny-to-interf** command applies only to inbound ACLs.
- The **ip access-group redirect-deny-to-interf** command cannot be applied on VPLS, VLL, or VLL-local endpoints and vice versa. Please refer to Configuration considerations for IPv4 outbound ACLs on VPLS, VLL, and VLL-Local endpoints on page 100.

  NOTE
  Redirect-deny packets do not apply to outbound
  traffic.

- On CES 2000 Series and CER 2000 Series devices, ACL deny-logging takes precedence over ACL Accounting. If the **ip access-group enable-deny-logging** command is configured on the interface, and both keywords (**enable-accounting** and **log**) are present in the same ACL clause, the statistics for that specific ACL clause are not collected. Both keywords will appear in the output of the **show access-list accounting** command indicating that logging is enabled, and the statistics for that specific ACL clause are not available. In the example output below, the **deny enable-accounting** and **log** keywords are applied to ip host 10.1.2.104/16.

```
0: deny enable-accounting ip host 10.1.2.104 10.19.0.0 10.0.255.255 log
      Hit count:  Accounting is not available due to deny logging
```

Permit-filter configuration notes

- (Not supported on CER 2000 Series or CES 2000 Series devices) ACL permit logging is supported on user and management interfaces.

- To reduce CPU overhead, you can add **selective** to any of the permit-logging enablement commands. However, **selective** does not affect traffic destined for the CPU.
    - **ip access-group enable-permit-logging selective**
    - **ipv6 traffic-filter enable-permit-logging selective**
    - **ip receive access-list enable-permit-logging selective**
    - **ipv6 receive access-list enable-permit-logging selective**

L3 ACL logging enablement

Enablement is one of the conditions required for ACL logging. With the exception of rACLs, this enablement is at interface level.

Enabling L3 ACL logging at interface level

There are specific interface-level commands for enabling IPv4/IPv6 deny/permit ACL logging.

> NOTE
> Using these commands, ACL logging can be enabled and disabled dynamically and does not require you to rebind the ACLs using the **ip rebind-acl** and **ipv6 rebind-acl** commands.

Enabling IPv4 ACL deny logging

Perform this task on each interface for which you need to enable IPv4 ACL deny logging.

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

   ```
   device# configure terminal
   ```

2. Enter the **interface** command to access the relevant interface.

   ```
   device(config)# interface ethernet 5/1
   ```

3. To enable unrestricted logging for IPv4 ACL **deny ... log** rules on that interface, enter the **ip access-group enable-deny-logging** command.

   ```
   device(config-if-e1000-5/1)# ip access-group enable-deny-logging
   ```

4. To enable partial logging for IPv4 ACL **deny ... log** rules on that interface, enter the **ip access-group enable-deny-logging hw-drop** command.

   ```
   device(config-if-e1000-5/1)# ip access-group enable-deny-logging hw-drop
   ```

   The **hw-drop** keyword reduces CPU load by logging only the first denied packet in each time cycle. To modify the default value of five minutes, refer to .

Enabling IPv4 ACL permit logging

Perform this task on each interface for which you need to enable IPv4 ACL permit logging.

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

   ```
   device# configure terminal
   ```

2. Enter the **interface** command to access the relevant interface.

   ```
   device(config)# interface ethernet 5/1
   ```

3. To enable unrestricted logging for IPv4 ACL **permit … log** rules on that interface, enter the **ip access-group enable-permit-logging** command.

```
device(config-if-e1000-5/1)# ip access-group enable-permit-logging
```

4. To enable partial logging for IPv4 ACL **permit … log** rules on that interface, enter the **ip access-group enable-permit-logging selective** command.

```
device(config-if-e1000-5/1)# ip access-group enable-permit-logging selective
```

The **selective** keyword reduces CPU load by logging only the first permitted packet in each time cycle. To modify the default value of five minutes, refer to Configuring log timers on page 189.

Enabling IPv6 ACL deny logging

Perform this task on each interface for which you need to enable IPv6 ACL deny logging.

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

```
device# configure terminal
```

2. Enter the **interface** command to access the relevant interface.

```
device(config)# interface ethernet 5/1
```

3. To enable unrestricted logging for IPv6 ACL **deny … log** rules, enter the **ipv6 traffic-filter enable-deny-logging** command.

```
device(config)# ipv6 traffic-filter enable-deny-logging
```

4. To enable partial logging for IPv6 ACL **deny … log** rules, enter the **ipv6 traffic-filter enable-deny-logging hw-drop** command.

```
device(config-if-e1000-5/1)# ipv6 traffic-filter enable-deny-logging hw-drop
```

The **hw-drop** keyword reduces CPU load by logging only the first denied packet in each time cycle. To modify the default value of five minutes, refer to Configuring log timers on page 189.

Enabling IPv6 ACL permit logging

Perform this task on each interface for which you need to enable IPv6 ACL permit logging.

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

```
device# configure terminal
```

2. Enter the **interface** command to access the relevant interface.

```
device(config)# interface ethernet 5/1
```

3. To enable unrestricted logging for IPv6 ACL **permit … log** rules on that interface, enter the **ip access-group enable-deny-logging** command.

```
device(config-if-e1000-5/1)# ipv6 traffic-filter enable-permit-logging
```

4. To enable partial logging for IPv6 ACL **permit ... log** rules on that interface, enter the **ipv6 traffic-filter enable-permit-logging selective** command.

    ```
    device(config-if-e1000-5/1)# ipv6 traffic-filter enable-permit-logging selective
    ```

    The **selective** keyword reduces CPU load by logging only the first denied packet in each time cycle. To modify the default value of five minutes, refer to Configuring log timers on page 189.

Enabling L3 rACL logging

There are several commands for globally enabling IPv4/IPv6 deny/permit rACL logging.

> **NOTE**
> Using these command, ACL logging can be enabled and disabled dynamically and does not require you to rebind the ACLs using the **ip rebind-receive-acl** command.

Enabling IPv4 rACL deny logging

Perform this task if you need to enable IPv4 rACL deny logging, which you do globally.

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

    ```
    device# configure terminal
    ```

2. To enable unrestricted logging for IPv4 rACL **deny ... log** rules, enter the **ip receive access-list enable-deny-logging** command.

    ```
    device(config)# ip receive access-list enable-deny-logging
    ```

3. To enable partial logging for IPv4 rACL **deny ... log** rules, enter the **ip receive access-list enable-deny-logging hw-drop** command.

    ```
    device(config-if-e1000-5/1)# ip receive access-list enable-deny-logging hw-drop
    ```

    The **hw-drop** keyword reduces CPU load by logging only the first denied packet in each time cycle. To modify the default value of five minutes, refer to Configuring log timers on page 189.

Enabling IPv4 rACL permit logging

Perform this task if you need to enable IPv4 rACL permit logging, which you do globally.

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

    ```
    device# configure terminal
    ```

2. To enable unrestricted logging for IPv4 rACL **permit ... log** rules, enter the **ip receive access-list enable-permit-logging** command.

    ```
    device(config)# ip receive access-list enable-permit-logging
    ```

3. To enable partial logging for IPv4 rACL **permit ... log** rules, enter the **ip receive access-list enable-permit-logging selective** command.

    ```
    device(config)# ip receive access-list enable-permit-logging selective
    ```

    The **selective** keyword reduces CPU load by logging only the first permitted packet in each time cycle. To modify the default value of five minutes, refer to Configuring log timers on page 189.

Enabling IPv6 rACL deny logging

Perform this task if you need to enable IPv6 rACL deny logging, which you do globally.

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

    ```
    device# configure terminal
    ```

2. To enable unrestricted logging for IPv6 rACL **deny** ... **log** rules, enter the **ipv6 receive access-list enable-deny-logging** command.

    ```
    device(config)# ipv6 receive access-list enable-deny-logging
    ```

3. To enable partial logging for IPv6 rACL **deny** ... **log** rules, enter the **ipv6 receive access-list enable-deny-logging hw-drop** command.

    ```
    device(config)# ipv6 receive access-list enable-deny-logging hw-drop
    ```

    The **hw-drop** keyword reduces CPU load by logging only the first denied packet in each time cycle. To modify the default value of five minutes, refer to

Enabling IPv6 rACL permit logging

Perform this task if you need to enable IPv4 rACL permit logging, which you do globally.

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

    ```
    device# configure terminal
    ```

2. To enable unrestricted logging for IPv6 rACL **permit** ... **log** rules, enter the **ipv6 receive access-list enable-permit-logging** command.

    ```
    device(config)# ipv6 receive access-list enable-permit-logging
    ```

3. To enable partial logging for IPv6 rACL **permit** ... **log** rules, enter the **ipv6 receive access-list enable-permit-logging selective** command.

    ```
    device(config)# ipv6 receive access-list enable-permit-logging selective
    ```

    The **selective** keyword reduces CPU load by logging only the first permitted packet in each time cycle. To modify the default value of five minutes, refer to

L3 ACL log-enabled rules

When you create ACL rules for which you want to enable logging, you must include the log parameter.

IPv4 ACL log-enabled rules

The following examples contain IPv4 ACL rules enabled for logging.

The following example is an IPv4 numbered extended ACL with a log-enabled deny rule.

```
device# configure terminal
device(config)#access-list 101 deny ip any any log
```

The following example is an IPv4 named standard ACL with a log-enabled deny rule.

```
device# configure terminal
device(config)# ip access-list standard vfour2
device(config-std-nacl-vfour2)# deny any any log
```

The following example is an IPv4 named extended ACL with a log-enabled permit rule.

```
device# configure terminal
device(config)# ip access-list extended vfour1
device(config-ext-nacl-vfour1)# permit ip any any log
```

The following example displays typical log entries (including a Layer 2 ACL).

```
[IPv4 Inbound ACL]
Dec 16 12:12:29:I:list 102 denied tcp 10.10.10.1(1024)(Ethernet 3/1 0000.0000.0010) - 10.20.20.1(1025),
27298224 event(s)

[L2 MAC ACL]
Dec 16 12:12:29:I: MAC ACL 400 denied 1 packets on port 3/16 [SA:0000.0000.0020, DA:0000.0000.0010,
Type:IPV4-L5, VLAN:1]
```

IPv6 ACL log-enabled rules

The following example contains IPv6 ACL rules enabled for logging.

The following example is an IPv6 ACL with a log-enabled permit rule.

```
device# configure terminal
device(config)# ipv6 access-list vsix1
device(config-ipv6-access-list vsix1)# permit tcp 2001:1570:21::/24 2001:1570:22::/24 log
```

Syslog messages for IPv6 rACLs

The following Syslog messages will be logged corresponding to the commands and conditions indicated.

1. **ipv6 receive rebind-acl-all**

```
SYSLOG: <14>Jun  6 10:37:54 FWD14 IPv6-rACL: rebinded by operator from console session.
```

2. **ipv6 receive deactivate-acl-all**

```
SYSLOG: <14>Jun  6 10:38:14 FWD14 IPv6-rACL: deactivated by operator from console session.
```

3. **no ipv6 receive deactivate-acl-all**

```
SYSLOG: <14>Jun  6 10:38:14 FWD14 IPv6-rACL: Activated by operator from console session.
```

4. **ipv6 receive delete-acl-all**

```
SYSLOG: <13>Jun 6 10:39:45 FWD14 IPv6-rACL: Deleting IPv6 Receive ACLs.
```

5. **IPv6 rACL CAM partition exhaust error (this unbinds rACL from the interface module)**

```
SYSLOG: <9>May 10 10:17:48 IxANVL-14 CAM IPv6 Receive ACL partition warning: total 1024 (reserved 0), free
0, slot 3, ppcr 0
SYSLOG: <13>Jun  1 07:15:10 IxANVL-1 IPv6-RACL: Port 3/1, IPv6 Receive ACL b1 exceed configured CAM size,
larger partition size required.
SYSLOG: <13>Jun  1 07:15:10 IxANVL-1 IPv6-RACL: Unbinding IPv6 Receive ACL b1
```

6. **Policy-map limit exhaust error (this unbinds rACL from the interface module).**

```
SYSLOG: <13>Jun  1 07:15:10 IxANVL-1 IPv6-RACL: Port 3/1, IPv6 Receive ACL b1 exceed configured RL class
limit.
SYSLOG: <13>Jun  1 07:15:10 IxANVL-1 IPv6-RACL: Unbinding IPv6 Receive ACL b1
```

7. **CAM malloc error (this unbinds rACL from the interface module).**

```
SYSLOG: <13>Jun  1 07:15:10 IxANVL-1 IPv6-RACL: Port 3/1, IPv6 Receive ACL b1 CAM malloc error.
SYSLOG: <13>Jun  1 07:15:10 IxANVL-1 IPv6-RACL: Unbinding IPv6 Receive ACL b1
```

Configuring log timers

You can specify how long the system waits before it sends a message in the Syslog.

The timer default value is 5 minutes, and values range from 1 through 10:

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

   ```
   device# configure terminal
   ```

2. To change the IPv4 timer value, enter the **ip access-list logging-age** command, specifying the new value.

   ```
   device(config)# ip access-list logging-age 8
   ```

3. To change the IPv6 timer value, enter the **ipv6 session-logging-age** command, specifying the new value.

   ```
   device(config)# ipv6 session-logging-age 3
   ```

Support for ACL CAM sharing

For ports sharing a PPCR to which the same ACLs are bound, ACL CAM sharing only applies if all or none of the ports have ACL Deny Logging configured.

In the following example, ports 4/1 and 4/2 in same packet processor (PPCR) are bound with inbound ACL 101 but only port 4/2 has the **ip access-group enable-deny-logging** command configured.

```
device(config)# acl-policy
device(config-acl-policy)# enable-acl-cam-sharing
device(config)# interface ethernet 4/1
device(config-if-e1000-4/1)# ip access group 101 in
device(config)# interface ethernet 4/2
device(config-if-e1000-4/2)# ip access group 101 in
device(config-if-e1000-4/2)# ip access-group enable-deny-logging
```

Because they do not have the same ACL Deny Logging configuration, a separate set of ACL CAM entries are programmed for ports 4/1 and 4/2.

UDAs for security and for PBR

User-defined Access Control Lists (UDAs) examine packet fields at specified offsets, applying permit and deny rules.

> NOTE
> Policy-based routing (PBR) also supports UDAs.

Up to four user defined data and masks can be specified for the ACL lookup. UDA field offsets are configurable for each physical port level. The ACL action processing of the packet is same as L2 ACL processing.

The UDA ACL supports 1000 numbered ACLs and 500 Named ACLs. The numbered UDA ACL starts from 2000 to 2999. The named UDA starts from 4000 to 4499.

The UDA ACL classifies the packets based on VLAN ID, 802.1p priority, and four user defined values (32 bit) and masks at the user defined offsets. These offsets are defined for each UDA ACL table. The UDA offset defined for ACL table are applied to the physical when the ACL is bound to any physical port.

The offsets are arrived based on the normalized packet format. In the normalized format, the VLAN headers in the packets are stripped. The offset specified in the UDA is based on the normalized packet format. For example, considering the packet format Ethernet/VLAN/IP/UDP, the normalized form is Ethernet/IP/UDP, the offsets will be calculated in the normalized packet format.

The normalization is applied when the Tag Protocol Identifier (TPID) of the packet is matching the interface configuration.

Interactions with other features

Interactions and limitations with the user defined ACL feature.

- When a physical port is bound to an UDA ACL (ingress), it cannot be bound to L2 ingress ACLs. Similarly, when a port is bound to L2 ingress ACL, UDA ingress ACL cannot be bound to the port.
- L3 PBR and UDA PBR can be configured on the same interface. But L2 PBR and UDA PBR cannot be configured on the same interface. When L3 and UDA PBR configured on the same interface, UDA PBR is applied only on non-IP packets.
- When configuring UDA rate limiting in the physical port, you cannot apply UDA ACL on that port.
- UDA inbound ACLs and UDA inbound ACL-based rate limiting are not supported on Layer-3 VPNs.
- IPv4 and IPv6 ACL-based rate limiting and UDA ACL-based rate limiting cannot be configured on the same port.
- Multiple rate limiting policies can be bound to a single port. However, once finding a matching ACL clause for a packet, the device does not evaluate subsequent clauses in that rate limiting ACL and subsequent rate limiting ACLs. UDA can not mixed with other ACLs (like L2).
- The Statistics display of the UDA ACL/PBR are subject to acl-policy configurations such as statistics-load-interval, display-pkt-bit-rate, and display-config-format.

Pre-requisites and limitations

- UDA ACL is for inbound operation only. UDA ACL is not supported for egress.
- UDA ACLs can be bound to physical interfaces only; it cannot be bound to virtual interfaces. In case of lag, it must be bound to the primary port of the lag.
- UDA ACL-based rate limiting can be applied on a physical port but cannot be applied on a virtual interface.
- UDA ACL having mirroring and log must not be used in UDA PBR. If used, mirror and log is ignored.
- UDA PBR cannot be applied globally.
- UDA PBR cannot be applied on interfaces where ACL based rate limiting is already applied.
- When both UDA PBR and IPv4 PBR are applied on an interface, the IPv6 packets do not UDA ACL because the UDA ACL applies only for non-IPv4 or non-IPv6 packets. This holds true for IPv6 + UDA, and IPv4 traffic. The IPv4 traffic is ignored.
- Open flow cannot be enabled on a port in which UDA ACL is configured.
- UDA cannot be applied as rACL.
- UDA ACLs are not supported on MPLS interfaces.

Security UDA configuration

This section contains tasks for configuring security user-defined ACLs (UDAs).

Configuring offsets for UDAs (fixed offsets)

To define UDA offset values, use the **uda-offsets** { *offset0* | **ignore** } { *offset1* **ignore** } { *offset2* **ignore** } { *offset3* **ignore** } command. This command configures in the physical interface.

This CLI command defines offsets with the following requirements:

- This CLI command defines offsets at which the user defined field starts.
- The width of each user defined fields is 32 bits.
- The offset must be on a four-byte boundary.
- The offset specified is the offset from the beginning of the normalized packet.

- By default, the maximum value of the offset is 116. To increase the maximum value to 124, use the **max-uda-offset** command.

You can define one or more valid user offsets based on the requirements, and other offsets can be specified as "ignore".

- If the offset is not in the four-byte boundary or greater than 116, an error message displays.
- The UDA offsets can be modified when the UDA bounds to the physical port. The UDA rules dynamically update to mask the "ignored" UDA fields.
- Deleting UDA offsets when a UDA is bound to the physical port is not allowed.

To define two offsets, use the following command:

```
device# configure terminal
device(config)# interface ethernet 1/1
device(config-intf-e1000-1/1)# uda-offsets  0  4 ignore ignore
```

To define up to four offsets, use the following command:

```
device# configure terminal
device(config)# interface ethernet 1/1
device(config-intf-e1000-1/1)# uda-offsets  0  4  8 12
```

To delete the UDA offset configuration, use the following command:

```
device# configure terminal
device(config)# interface ethernet 1/1
device(config-intf-e1000-1/1)# no uda-offsets
```

Configuring offsets for UDAs (flexible offsets)

Flexible offsets enable you to specify *offset2* and *offset3* values that are not multiples of 4 bytes. The width of all offset fields is 4 bytes. Support for flexible offsets is as follows:

- The XMR Series supports only fixed UDA offsets.
- The MLX Series BR-MLX-40Gx4-X line card supports flexible UDA offsets.
- The MLX Series BR-MLX-10Gx20 and BR-MLX-100Gx2-CFP2 line-cards support flexible UDA offsets only in Network Packet Broker (NPB) mode.
- All other MLX Series line cards support only fixed UDA offsets.
- All CER 2000 Series and CES 2000 Series configurations support only fixed UDA offsets.

The following example utilizes the "flexible" option to specify *offset2* and *ofset3* that are not multiples of 4.

```
device# configure terminal
device(config)# uda access-list uda_flex
device(config-uda-acl-uda_flex)# permit any any 4d455353 ffffffff 41474500 ffffff00
device(config-uda-acl-uda_flex)# exit
device(config)# interface ethernet 1/2
device(config-if-e1000-1/2)# ignore ignore 98 102
device(config-if-e1000-1/2)# uda access-group uda_flex in
```

Configuring numbered UDAs.

To define a User Defined ACL table, use the following command. This command creates a User defined ACL Table with a list of user defined values.

[no] access-list *num* [**sequence** *num*] **permit** | **deny** { *vlan_id* | **any** } { *uda_val0 uda_mask0* | **any**} { *uda_val1 uda_mask1* | **any**} { *uda_val2 uda_mask2* | **any**} { *uda_val3 uda_mask3* | **any**} [**priority_mapping 802.1p-value**] [**priority 802.1p-value** | **priority-force 802.1p-value** | **drop-precedence** *value*] [**drop-precedence-force** *value*] [**mirror**] [**log**]

On the CER 2000 Series and CES 2000 Series series, the **exclude-pcp-marking** option is available if you include the **priority** or **priority-force** keywords in ACL permit or deny statements. The **exclude-pcp-marking** option prevents PCP modification both for regular Layer 2/Layer 3 traffic and for VPLS packets.

The uda_val0, uda_val1, uda_val2, uda_val3 values are matched against the packet offsets defined for the inteface to which the ACL is applied.

If any of the UDA offsets are not configured while defining UDA offsets, the UDA rule internally masks that particular UDA field.

The behavior of parameters such as priority, priority-mapping, priority-force, drop-precedence, drop-precedence-force, mirror, and log is the same as the regular L2 ACL.

The configuration option **arp-guard** is supported in the L2 ACL but not supported in the UDA.

Use the optional **sequence** keyword to specify the sequence number for each ACL rule. It is helpful to insert a new rule between the existing ACL rules. Without this option, the sequence number assigns internally starting from 10 and increments by 10 for each rule. The behavior of this option is same as L2 ACLs. Review the sample configuration below:

```
device# configure terminal
device(config)# access-list 2000 sequence 100 permit 100 aabbccdd ffffffff eeff0000 ffff0000  08000000
ffff0000 0a0a ffff
```

Deleting a numbered UDA rule

The UDA configuration can be deleted with the **no** option. The UDA rule can be deleted with the sequence number, the rule parameters, or with the sequence number and rule parameters.

To delete the UDA rule with the sequence number, use the following command:

```
device# configure terminal
device(config)# no access-list 2000 sequence 100
```

To delete the UDA rule with the sequence number and parameters, use the following command:

```
device# configure terminal
device(config)# no access-list 2000 sequence 100 permit 100  aabbccdd ffffffff eeff0000 ffff0000  08000000
ffff0000 0a0a ffff
```

To delete the UDA rule with parameters, use the following command:

```
device# configure terminal
device(config)# no access-list 2000 permit 100 aabbccdd ffffffff eeff0000 ffff0000  08000000  ffff0000 0a0a
ffff
```

Binding a numbered UDA to a physical port

To bind a numbered UDA to a physical port, use the following command:
uda access-group *uda-acl-num* **in**

- The user-defined ACL created passes to this CLI command.
- The user-defined ACLs is only supported on the ingress side. The UDA offsets must be defined for the access list before binding the ACL to any physical port. If not, an error message displays **"UDA offsets are not defined for this port"** and the binding fails.
- All of the UDA rules defined in the UDA are programmed into the hardware. The UDA offsets configured as "ignore" are masked in the ACL rule while programming in the hardware.
- If the empty UDA is bound to a physical port, the UDA lookup does not start until rules are added.

Displaying numbered UDAs

To display the list of all user-defined ACLs (UDAs) configured in the system or specified numbered UDA, use the following commands:
show access-list *acl-num* **show access-list uda**

Below is a sample output.

```
device# show access-list 2000
UDA Access List 2000:
10: access-list 2000 permit 100 any any 00001122 0000ffff 00003344 0000ffff
20: access-list 2000 permit any any any any any
```

Sample numbered UDA configurations

The following contains a list of CLI commands for creating a UDA for filtering packets based on VLAN, MAC DA, Ethernet type, and DIP (MSB 16 bits).

- To define an Offset for MAC DA and Ethernet type. The offsets of the MAC DA is 0 and 4 (DA is split into two words). The offset of the Ethernet type is 12. The offset of DIP is 30 and after converting to the four-byte boundary, the value is 28. The VLAN Header is not considered for offset calculation. See the sample command configuration below.

```
device# configure terminal
device(config)# interface ethernet 1/1
device(config-intf-e1000-1/1)# uda-offsets  0  4  12 28
```

- To create a UDA for VLAN, MAC DA, Ethernet type and DIP. Assume the MAC DA is aabb:ccdd:eeff and IP Address is 10.10.*.*. See the sample command configuration below.

```
device# configure terminal
device(config)# access-list 2000 permit 100 aabbccdd ffffffff eeff0000 ffff0000  08000000  ffff0000
0a0a ffff
```

- To bind the UDA to port. See the sample command configuration below.

```
device# configure terminal
device(config)# interface ethernet 1/1
device(config-intf-e1000-1/1)# uda access-group 2000 in
```

The following contains a list of CLI commands for creating a UDA for filtering packets based on VLAN, MAC DA and Ethernet type.

- To define Offset for MAC DA and Ethernet type. The offsets of the MAC DA is 0 and 4 (DA is split into two words). The offset of the Ethernet type is 12. The VLAN Header is not considered for offset calculation. See the sample command configuration below.

```
device# configure terminal
device(config)# interface ethernet 1/1
device(config-intf-e1000-1/1)# uda-offsets  0  4  12 ignore
```

- To create a UDA for VLAN, MAC DA, and Ethernet type. The UDA field 3 can be specified as "any" as the uda-offset is not defined. Even when the UDA field 3 is configured any other value, this field is still masked in the hardware table because the uda-offset is not configured. Assume the MAC DA is aabb:ccdd:eeff and Ethernet type is 0x800. See the sample command configuration below.

```
device# configure terminal
device(config)# access-list 2000 permit 100 aabbccdd ffffffff eeff0000 ffff0000  08000000  ffff0000
any
```

- To bind the UDA to port. See the sample command configuration below.

```
device# configure terminal
device(config)# interface ethernet 1/1
device(config-intf-e1000-1/1)# uda access-group 2000 in
```

Configuring named UDAs

Use the **uda access-list** command to create a named user-defined ACL (UDA).

You add rules under the named UDA, using the **permit** and **deny** commands.

If the UDA offsets are configured as "ignore" while defining UDA offsets, the UDA rule internally masks that particular UDA field.

Deleting a named UDA rule

UDA rules can be deleted with the **no** option. Rules can be deleted with the sequence number, the rule parameters, or with the sequence number and rule parameters.

- The following example deletes with the sequence number.

```
device# configure terminal
device(config)# uda access-list udaAcl
device(config-uda-acl-udaAcl)# no sequence 100
```

- The following example deletes with the parameters.

```
device# configure terminal
device(config)# uda access-list udaAcl
device(config-uda-acl-udaAcl)# no permit 100 aabbccdd ffffffff eeff0000 ffff0000  08000000  ffff0000
0a0a ffff
```

- The following example deletes with the sequence number and the parameters.

```
device# configure terminal
device(config)# uda access-list udaAcl
device(config-uda-acl-udaAcl)# no sequence 100 permit 100 aabbccdd ffffffff eeff0000 ffff0000
08000000  ffff0000 0a0a ffff
```

Binding a named UDA to a physical port

Use the following command to bind a user-defined ACL (UDA) to a physical port.

```
device# configure terminal
device(config)# interface ethernet 1/7
device(config-if-e1000-1/7)# uda access-group uda_01 in
```

If the UDA rules are not previously created, all the traffic on the port drops until the UDA is defined.

The user-defined ACLs are only supported on the ingress side. The UDA offsets must be defined for the interface before binding the ACL to that interface. If not, an error message **"UDA offsets are not defined for this port"** displays and binding fails.

All of the UDA rules defined in the UDA are programmed into the hardware. The UDA offsets not configured are masked in the ACL rule while programming in the hardware.

Displaying named UDAs

The **show access-list uda** [*uda-name*] command displays the details of a named UDA.

The following example is a sample configuration.

```
device# show access-list uda TestUdaAcl
UDA Access List TestUdaAcl:
access-list 2000 uda-offsets 12     20     36      72
10: access-list 2000 permit 100 any any 00001122 0000ffff 00003344 0000ffff
20: access-list 2000 permit any any any any any
```

Sample configurations for named UDAs

The following examples create named UDAs for filtering packets based on VLAN, MAC DA, Ethernet type, and DIP (16 MSB Bits).

- The following example defines offsets for MAC DA and Ethernet type. The offsets of the MAC DA is 0 and 4 (DA is split into two words). The offset of the Ethernet type is 12. The offset of DIP is 26 and after converting it to 32-bit boundary, the value is 24. The VLAN Header is not considered for the offset calculation.

    ```
    device# configure terminal
    device(config)# interface ethernet 1/1
    device(config-if-e1000-1/1) uda-offsets 0 4 12 24
    ```

- The following example creates a UDA for VLAN, MAC DA, Ethernet type, and DIP.

    ```
    device# configure terminal
    device(config)# uda access-list testudaAcl
    device(config-uda-acl-testudaAcl)# permit 100 aabbccdd ffffffff eeff0000 ffff0000  08000000
    ffff0000 0a0a ffff
    ```

- The following example binds the UDA to a port.

    ```
    device# configure terminal
    device(config)# interface ethernet 1/1
    device(config-if-e1000-1/1)# uda access-group testUdaAcl in
    ```

Configuring the maximum number of UDA rules

The **system-max uda-acl-entries** command configures the system max value for numbered-UDA rules. The default number of rules per UDA is 64. The maximum number of entries ranges from 64 through 256.

This configuration requires a reload of the system to take effect.

UDA access group enable deny logging

The **uda access-group enable-deny-logging** command enables deny logging for the UDA access group.

A typical log is as follows.

```
UDA ACL 2001 denied 1 packets on port 3/4 [SA:0024.389e.2b03, DA:0180.c200.0002, Type:UNKNOWN, VLAN:1]
```

UDA configuration for PBR route maps

This section contains tasks for configuring user-defined ACLs (UDAs) for policy-based routing (PBR).

The UDA can be defined along with IPv4, IPv6, and Layer 2 ACL route map entries. The user can also define the interface or VLAN using the **set interface** or **set next-hop-flood-vlan** commands. As like IPv4, IPv6, and Layer 2 ACL policies, up to five UDA policies can be configured under the route map.

See the sample command configuration below.

```
device(config)# route-map testRouteMap permit 1
device(config-routemap testRouteMap permit 1)# match ip address 102
```

```
device(config-routemap testRouteMap permit 1)# match ipv6 address v6Filter
device(config-routemap testRouteMap permit 1)# match l2acl 400 401
device(config-routemap testRouteMap permit 1)# match uda 2000 2001 2002 2003 2004
device(config-routemap testRouteMap permit 1)# set ip next-hop 100.10.10.1
device(config-routemap testRouteMap permit 1)# set ipv6 next-hop 100:10::10.1
device(config-routemap testRouteMap permit 1)# set interface ethernet 1/1
```

Binding the UDA route map

The route map can be bound to a physical interface as a UDA policy. When bound as a UDA policy, only the UDA rules are applied on that interface. The UDA policy and L3 policies can be configured on the same interface. When the L3 policy is configured, the UDA is applied only for non IP packets. If the L3 policy is not applied, the UDA is applied to all the packets on that interface.

The UDA offset must be defined before binding the UDA PBR to the interface. The UDA offsets cannot be modified while any UDA PBR policy is present in the interface.

When the UDA policy is bound to an interface, the L2 policy cannot be applied on that interface and vice versa.
uda policy route-map *route_map_name*

See the sample command configuration below.

```
device# configure terminal
device(config)# interface eth 1/7
device(config-if-e1000-1/7)# ip policy route-map testRouteMap
device(config-if-e1000-1/7)# ipv6 policy route-map testRouteMap
device(config-if-e1000-1/7)# uda policy route-map testRouteMap
device(config-if-e1000-1/7)# allow-all-vlan pbr
```

Configuring rate limit

Use this CLI configuration for configuring the rate limiting for the UDAs. This is similar to the L2 ACL and support is also extended to UDAs. The UDA offsets must be defined before binding the UDA to rate limiting configuration.
rate-limit input access-group name [*ipv4_name* | *lpv6_name* | *mac_name* | *uda_name*]

Similarly, the Numbered ACL support in the rate limiting configuration is extended to UDAs. The ACL ID also supports the UDA numbered ACL range.
rate-limit input access-group *num*

Displaying the UDA PBR statistics

Use this command to display the UDA PBR statistics on the specified interface. Use the clear command to clear the PBR statistics on the specified interface.
show access-list accounting ethernet *slpt/port* **in uda policy-based-routing clear access-list ethernet** *slot/port* **uda policy-based-routing**

See the sample command configuration below.

```
device# configure terminal
device(config)# show access-list accounting ethernet 3/1 in uda policy-based-routing
Policy based Routing Accounting Information:
Routemap route1
ACL ACLNameTest112345679-023456789-0123456789
   0: sequence 10 permit 100 any any 1234 ffff any
        Hit count: (1 sec) 0 (1 min) 0
(5 min) 0 (accum) 0
```

Displaying the UDA PBR brief statistics

Use this CLI command to display the brief version of the configured UDA PBR.

show access-list accounting brief uda policy-based-routing

See the sample command configuration below.

```
device# show access-list accounting brief uda  policy-based-routing
 3/1                        2000              0(1s)
                                                  0(1m)
                                             0(5m)
                                             0(ac)
```

Upgrade and downgrade considerations (5.6.00)

Multi-Service IronWare versions later than R05.6.00 support ACL entry sequence numbers for Layer-2, IPv4 and IPv6 ACLs. Where ACL filters have been configured on R05.6.00 and you want to downgrade a device to an earlier version of software, you should enable **suppress-acl-seq** prior to the downgrade.

> **NOTE**
> If **suppress-acl-seq** is not enabled before downgrade to versions earlier than Multi-Service IronWare R05.6.00, ACL configurations created with the **sequence** parameter on R05.6.00 will not be allowed on older releases and will result in an error.

By default, the **suppress-acl-seq** switch is OFF. When it is turned ON, the system:

- Hides or suppresses sequence numbers for ACL filters while:
    - Executing **show access-list** commands
    - Displaying the **running-config**
    - Saving the running-config using **write memory**
    - Copying the **running-config** to a tftp server
- Hides all unused IPv6 **remark-entry** configuration statements when **running-config** is displayed or stored.
- Shows all used IPv6 **remark-entry** configuration statements as **remark** statements when **running-config** is displayed or stored.

The following example displays **show access-list** command output for IPv6 ACL "ip6_"when **suppress-acl-seq** is OFF.

```
device(config)# show access-list ip6_
ipv6 access-list ip6_: 11 entries
 0: remark unused default comment
 1: remark-entry sequence 1 unused comment
 5: remark allowonly  udp traffic from 1::5
 5: permit udp host 1::5 any sequence 5
 7: remark-entry sequence 7 permit all ipv6 traffic for 1::3
 9: remark-entry sequence 9 deny udp traffic for 1::2
 9: deny udp host 1::2 any sequence 9
 10: remark-entry sequence 10 permit all ipv6 traffic for 1::1
 10: permit ipv6 host 1::1 any
 12: remark allow only sctp traffic for 1::10
 12: permit sctp host 1::10 any sequence 12
 15: remark-entry sequence 15 deny all tcp traffic for 1::9
 17: remark-entry sequence 17 deny tcp traffic for 1::2
 17: deny tcp host 1::2 any sequence 17
 23: remark-entry sequence 23 allow rest of the ipv6 traffic for 1::2
 23: permit ipv6 host 1::2 any sequence 23
 28: remark-entry sequence 28 permit all ipv6 traffic for 1::9
```

To turn **suppress-acl-seq** ON and display the **show access-list** command output again, enter the following commands.

```
device(config)# acl-policy
device(config-acl-policy)# suppress-acl-seq
device(config-acl-policy)# exit
device(config)# show access-list ip6_
ipv6 access-list ip6_: 11 entries
 0: remark unused default comment
 1: remark-entry sequence 1 unused comment
```

```
    5: remark allowonly  udp traffic from 1::5
    5: permit udp host 1::5 any
    7: remark-entry sequence 7 permit all ipv6 traffic for 1::3
    9: remark-entry sequence 9 deny udp traffic for 1::2
    9: deny udp host 1::2 any
    10: remark-entry sequence 10 permit all ipv6 traffic for 1::1
    10: permit ipv6 host 1::1 any
    12: remark allow only sctp traffic for 1::10
    12: permit sctp host 1::10 any
    15: remark-entry sequence 15 deny all tcp traffic for 1::9
    17: remark-entry sequence 17 deny tcp traffic for 1::2
    17: deny tcp host 1::2 any
    23: remark-entry sequence 23 allow rest of the ipv6 traffic for 1::2
    23: permit ipv6 host 1::2 any
    28: remark-entry sequence 28 permit all ipv6 traffic for 1::9
```

Because **suppress-acl-seq** is ON, the system hides the user-configured sequence numbers for ACL filters.

The following examples show how the **suppress-acl-seq** state affects the display of **remark-entry** configuration statements. When **suppress-acl-seq** is OFF, the **running-config** for IPv6 ACL "ip6_" is:

```
ipv6 access-list ip6_
 remark-entry sequence  1 unused comment
 remark allow only  udp traffic from 1::5
 permit udp host 1::5 any sequence 5
 remark-entry sequence  7 permit all ipv6 traffic for 1::3
 remark-entry
 sequence 9 deny udp traffic for 1::2
 deny udp host 1::2 any sequence 9
 remark-entry
 s
equence 10 permit all ipv6 traffic for 1::1
 permit ipv6 host 1::1 any
 remark allow only sctp traffic for 1::10
 permit sctp host 1::10 any sequence 12
 remark-entry sequence  15 deny all tcp traffic for 1::9
 remark-entry
 sequence 17 deny tcp traffic for 1::2
 deny tcp host 1::2 any s
equence 17
 remark-entry
 sequence 23 allow rest of the ipv6 traffic for 1::2
 permit ipv6 host 1::2 any s
equence 23
 remark-entry sequence 28 permit all ipv6 traffic for 1::9

 remark unused default comment
```

When **suppress-acl-seq** is turned ON, the **running-config** display for IPv6 ACL "ip6_" is:

```
ipv6 access-list ip6_
 remark allow only  udp traffic from 1::5

permit udp host 1::5 any
 remark
 deny udp traffic for 1::2
 deny udp host 1::2 any
 remark
 permit all ipv6 traffic for 1::1
 permit ipv6 host 1::1 any
 remark allow only sctp traffic for 1::10
 permit sctp host 1::10 any
 remark
 deny tcp traffic for 1::2
 deny tcp host 1::2 any
 remark
 allow rest of the ipv6 traffic for 1::2
 permit ipv6 host 1::2 any
 remark unused default comment
```

When suppress-acl-seq is ON, the system hides unused **remark-entry** statements and displays used **remark-entry** statements as **remark** statements.

Syntax: [no] suppress-acl-seq

The **no** version of this command turns **suppress-acl-seq** OFF.

HTTP and HTTPS

Configuring SSL security for the Web Management Interface

When enabled, the SSL protocol uses digital certificates and public-private key pairs to establish a secure connection to the Extreme device. Digital certificates serve to prove the identity of a connecting client, and public-private key pairs provide a means to encrypt data sent between the device and the client.

> **NOTE**
> The Web Management Interface is only supported on the XMR Series and MLX Series devices

Configuring SSL for the Web Management Interface consists of the following tasks:

- Enabling the SSL server on the Extreme device
- Importing an RSA certificate and private key file from a client (optional)
- Generating a certificate

Enabling the SSL server on a device

To enable the SSL server on an Extreme device, enter the following command.

```
device(config)# web-management https
```

Syntax: [no] web-management http | https

You can enable either the HTTP or HTTPs servers with this command. You can disable both the HTTP and HTTPs servers by entering the following command.

```
device(config)# no web-management
```

Syntax: [no] web-management

Specifying a port for SSL communication

By default, SSL protocol exchanges occur on TCP port 443. You can optionally change the port number used for SSL communication.

For example, the following command causes the device to use TCP port 334 for SSL communication.

```
device(config)# ip ssl port 334
```

Syntax: [no] ip ssl port port-number

The default port for SSL communication is 443.

Importing digital certificates and RSA private key files

To allow a client to communicate with the other Extreme device using an SSL connection, you configure a set of digital certificates and RSA public-private key pairs on the device. A digital certificate is used for identifying the server to the connecting client. It contains information about the issuing Certificate Authority, as well as a public key. You can either import digital certificates and private keys from a server, or you can allow the Extreme device to create them.

If you want to allow the Extreme device to create the digital certificates, refer to the next section, Generating an SSL certificate on page 202. If you choose to import an RSA certificate and private key file from a client, you can use TFTP to transfer the files.

For example, to import a digital certificate using TFTP, enter a command such as the following.

```
device# copy tftp flash 10.168.9.210 certfile server-certificate
```

Syntax: copy tftp flash ip-address file-name server-certificate

> **NOTE**
> If you import a digital certificate from a client, it can be no larger than 2048
> bytes.

To import an RSA private key from a client using TFTP, enter a command such as the following.

```
device# copy tftp flash 10.168.9.210 keyfile server-private-key
```

Syntax: copy tftp flash ip-address file-name server-private-key

The *ip-addr* is the IP address of a TFTP server that contains the digital certificate or private key.

Generating an SSL certificate

If you did not already import a digital certificate from a client, the device can create a default certificate. To do this, enter the following command.

```
device(config)# crypto-ssl certificate generate
```

Syntax: [no] crypto-ssl certificate generate

Deleting the SSL certificate

To delete the SSL certificate, enter the following command.

```
device(config)# crypto-ssl certificate zeroize
```

Syntax: [no] crypto-ssl certificate zeroize

IPsec

IPsec overview

Internet Protocol security (IPsec) is a suite of protocols that provide secure communication between devices at the network layer (Layer 3) across public and private networks.

IPsec provides end-to-end security for data traffic by using encryption and authentication techniques that ensure data privacy. Encrypted packets are routed in the same way as ordinary IP packets.

IPsec components include:

- Authentication Header (AH)
- Encapsulating Security Payload (ESP)
- Internet Key Exchange (IKE)

Authentication Header (AH) is a security protocol that provides source authentication and data integrity.

Encapsulating Security Payload (ESP) is a security protocol that provides data confidentiality in addition to the source authentication and data integrity that is also provided by AH. The Extreme implementation of IPsec uses ESP.

ESP supports two modes of use: transport mode and tunnel mode. In transport mode, an IPsec header is inserted into the IP packet and the packet payload is encrypted. In tunnel mode, the original IP packet is encrypted as an inner IP payload and an IPsec header and outer IP header are added so that the IPsec header and encrypted IP packet become the data component of a new and larger IP packet, as shown in the following figure.

Extreme devices support ESP in tunnel mode.

FIGURE 2 IPsec tunnel mode versus transport mode

Internet Key Exchange (IKE) is used to establish an IPsec tunnel. IKE performs mutual authentication of peer devices and establishes and maintains a secure channel for communication between the devices.

Extreme devices support Internet Key Exchange version 2 (IKEv2).

A security association (SA) is another important concept in IPsec. The SA is a logically secure relationship between peer devices. Both IKEv2 and IPsec SAs are used to establish an IPsec tunnel.

FIGURE 3 Basic IPsec functionality

The preceding figure shows the secure transfer of IP data between two routers, R1 and R3, over an insecure or public network by using IPsec. First, the tunnel parameters, transform set, and crypto algorithms to be used for encryption and authentication along with associated policy filters are configured on both R1 and R3. IKE negotiations are used to establish the tunnel.

Once the tunnel is up, all packets going out over the tunnel are encrypted and packets received on the tunnel interface are decrypted.

Acronyms

The following acronyms may be used in this document in the description of IP security.

AES-128-GCM	Advanced Encryption Standard with 128-bit key size in Galois Counter Mode
AES-256-GCM	Advanced Encryption Standard with 256-bit key size in Galois Counter Mode
AES-CBC-128	Advanced Encryption Standard with 128-bit key size in Cipher Block Chaining Mode
AES-CBC-256	Advanced Encryption Standard with 256-bit key size in Cipher Block Chaining Mode
AH	Authentication Header protocol
CA	Certificate Authority
CDP	CRL Distribution Point
Cert	Certificate
CRL	Certificate Revocation List
DH	Diffie-Hellman
DN	Distinguished Name
DPD	Dead Peer Detection
ECDSA	Elliptical Curve Digital Signature Algorithm
ESP	Encapsulating Security Payload protocol
GCM	Galois/Counter Mode
HMAC	Hash-based Message Authentication Code
IKEv2	Internet Key Exchange protocol, version 2
IPsec	IP security
IPsec LP	IPsec Line Card
ISAKMP	Internet Security Association and Key Management Protocol
OCSP	Online Certificate Status Protocol, RFC 6960

PKI	Public Key Infrastructure
RA	Registration Authority
SA	Security Association
SAD	Security Association Database
SCEP	Simple Certificate Enrollment Protocol, draft RFC
SHA	Secure Hash Algorithm
SPD	Security Policy Database
SPI	Security Parameters Index (used to identify a SA)
TC	Traffic Class field in IPv6 header
TOS	Type of Service field in IPv4 header
VTI	Virtual Tunnel Interface

IPsec standards

Various RFCs define IPsec standards.

For further information, refer to the RFCs listed in the following table.

TABLE 23 RFCs that define IPsec standards

RFC	Title
RFC 2408	Internet Security Association and Key Management Protocol (ISAKMP)
RFC 2560	X.509 Internet Public Key Infrastructure Online Certificate Status Protocol – OCSP
RFC 4106	The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)
RFC 4301	Security Architecture for the Internet Protocol
RFC 4302	IP Authentication Header
RFC 4303	IP Encapsulating Security Payload (ESP)
RFC 4835	Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)
RFC 4868	Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec
RFC 6379	Suite B Cryptographic Suites for IPsec
RFC 7296	Internet Key Exchange Protocol Version 2 (IKEv2)

Supported hardware

The use of IPv4 and IPv6 over IPsec is supported on the following hardware.

TABLE 24 Supported hardware for IPsec

Hardware	Description
MLXe device	Router
BR-MLX-10GX4-M-IPSEC	IPsec LP line card (module)

IPsec tunnel setup process

There are a number of factors that impact the setup process for an IPsec tunnel.

The factors that can affect the setup procedure include:

- The use of PKI—You must complete the PKI configuration before you set up the IPsec tunnel.
- The use of non-default settings for mandatory tunnel elements .

  IPsec tunnels can be set up by using default settings for the mandatory tunnel elements. When the default settings are not acceptable, you can configure the mandatory tunnel elements. Using the default settings simplifies the configuration and reduces the number of steps in the setup process.

- The use of other options that are not mandatory, such as authentication algorithms.

  When enabling an option that is disabled by default, or configuring an option that is not required to set up an IPsec tunnel (for example, generating key pairs), make sure you are familiar with the option before you begin the tunnel set up procedure. The use of non-mandatory tunnel options adds steps to the tunnel setup process and depending on the option, can involve dependencies or require certain hardware.

  Some of the options you can choose to enable or configure as part of the tunnel set up process include:

  – AES-GCM-128 algorithm
  – NAT traversal
  – Generating key-pairs using ECDSA P-256 or ECDSA P-384 for signature generation and verification.

- The version of IP.

  The process you use to set up IPv4 and IPv6 tunnels is very similar. The differences are:

  – Tunnel mode—Use the **tunnel mode ipsec** command to set the tunnel mode to IPv4 or IPv6.
  – Tunnel source—The options for specifying the tunnel source are different.
  – Tunnel destination—For IPv6, set the tunnel destination to the IPv6 address of the remote endpoint of the tunnel.

  > **NOTE**
  > IKE cryptographic algorithms and key sizes that are configured for an IPsec tunnel must not be weaker than the IPsec cryptographic algorithms and key sizes that are configured for the tunnel.

PKI configuration

Public Key Infrastructure (PKI) provides certificate management to support secured communication for security protocols such as IP security (IPsec). When Public Key Infrastructure (PKI) is used in the configuration of an IPsec tunnel, you must complete the PKI configuration before you begin the IPsec tunnel configuration.

The PKI options such as PKI entity, PKI enrollment profile and PKI trustpoint must be configured before you can select them during the configuration of the IPsec tunnel.

PKI configuration includes:

- Generating a key pair for the local certificate
- Configuring an entity Distinguished Name
- Creating a trustpoint
- Configuring Certificate Authority (CA) authentication
- Creating a PKI enrollment profile
- Generating a certificate request
- Installing identity certificates

Setup considerations for IPsec tunnels

Some considerations apply to the planning and setup of IPsec tunnels.

- The simultaneous use of IPsec IPv4 and IPsec IPv6 tunnels on the same Extreme router is supported. When mixing tunnels on a device, the maximum number of tunnels for a system is 2048, and the maximum number of tunnels for a IPsec line card is 256.
- Although you can configure IPsec IPv4 and IPsec IPv6 tunnels on the same Extreme router, you cannot:
  - Use IPsec IPv4 tunnels to encrypt or decrypt IPv6 packets.
  - Use IPsec IPv6 tunnels to encrypt or decrypt IPv4 packets.
- Security options are configured independently for IPv6 IPsec and IPv4 IPsec tunnels. Different levels of security can be defined for each type of tunnel.

You secure the IPv6 payload by configuring the security of the IPv6 IPsec tunnel used to deliver the payload. You use IPsec parameters to configure the security options of the tunnel.

Establishment of an IPsec tunnel

IKEv2 negotiations are used to establish an IPsec tunnel.

There are two phases in the IKEv2 negotiation process.

In Phase 1, the tunnel endpoints exchange proposals for mutual authentication and for securing the communication channel. The IKEv2 protocol is used to dynamically negotiate and authenticate keying material and other security parameters that are required to establish secure communications. A secret shared key for encrypting and decrypting the IKEv2 packets themselves is derived in this phase. When the Phase 1 negotiations are successful, an IKEv2 SA, which contains the negotiated security encryption and keying material, is established. The IKEv2 SA is a secure "control channel" where keys and other information for protecting Phase 2 IKEv2 negotiations are maintained. The IKEv2 SA established in Phase 1 is bidirectional.

In Phase 2, the tunnel endpoints exchange proposals for securing the data that is to be sent over the tunnel. Phase 2 communication is secure. When Phase 2 negotiations are successful, a pair of IPsec SAs is established. An IPsec SA is unidirectional; one IPsec SA is needed for inbound traffic and one for outbound traffic. SAs that are negotiated by using the IKE SA (such as IPsec SAs) are also known as child SAs. The negotiated data path keys are then programmed into the specialized hardware crypto engine, and the tunnel state is set to up. At this point, the user data can be exchanged through the encrypted tunnel.

The set of IPsec parameters that describe an IPsec tunnel connection is known as an IPsec security policy. The IPsec security policy describes how both endpoints use the security services, such as encryption and hash algorithms for secure communication.

Configuration of an IPsec tunnel

Configuration of an IPsec tunnel includes the configuration of virtual tunnel interfaces (VTIs) at the tunnel endpoints and the configuration of both the IKEv2 and IPsec parameters that are used to establish the tunnel and secure the tunnel traffic.

To configure an IPsec tunnel, you must complete the following tasks at the tunnel endpoints:

- Configure a virtual tunnel interface and set the mode of the tunnel to ipsec.
- Configure the following values (when the default values are not acceptable):
  - Global parameters for IKEv2
  - An IKEv2 proposal
  - An IKEv2 policy
  - An IKEv2 authentication proposal
  - An IKEv2 profile
  - An IPsec proposal

- – An IPsec profile
- Bind the IPsec profile to the VTI by using the **tunnel protection ipsec profile** command.

Configuration of traffic to route over an IPsec tunnel

Traffic is routed over an IPsec tunnel by using static route configuration, a dynamic routing protocol, or by using policy-based routing (PBR).

MLX Series supports dynamic routing of traffic over an IPsec tunnel by using routes learned by way of Routing Information Protocol (RIP), Open Shortest Path First (OSPF), or Border Gateway Protocol version 4 (BGP4).

Because IPsec tunnels are established by using virtual tunnel interfaces (VTIs), they can plug into the routing protocol infrastructure of a router. IPsec VTIs impact a change in the packet path based on routing metrics or by toggling the link state of the tunnel. VTIs provide termination points for site-to-site IPsec VPN tunnels and allow them to behave like routable interfaces. VTIs simplify the IPsec configuration and enable common routing capabilities to be used because the endpoint is associated with an actual interface.

Using VTIs offers the following advantages:

- IPsec configuration does not require a static mapping of IPsec sessions to a physical interface.
- IPsec VTIs simplify encapsulation and do not require the use of crypto maps for IPsec.
- Common interface capabilities can be applied to the IPsec tunnel because there is a routable interface at the tunnel endpoint.
- IPsec VTIs support flexibility for the sending and receiving of unicast encrypted traffic on the IPsec module.

Recommendation for using LAG on the same IPsec line card

There are recommended settings for using LAG on physical ports of the same IPsec line card.

When using LAG on physical ports of the same IPsec line card, make sure that replay-check is **disabled**. If replay-check is enabled, packet loss can occur due to packet re-ordering and a small replay window setting (for example, 64 packets).

If you encounter packet re-ordering, make sure that LAG is formed using one of the following port groups:

- **Port group 1**: 1, 2, 5 and 7
- **Port group 2**: 3, 4, 5, 6, and 8

Using one of these port groups can improve performance because ingress IPsec processing for IPv4 and IPv6 is done using two separates engines. One engine handles the packets received on **Port group 1** (ports 1, 2, 5, and 7), and the second engine handles the packets received on **Port group 2** (ports). Using the same engine for a given LAG can minimize IPsec packet re-ordering.

Supported algorithms

A number of algorithms are supported for IKEv2 negotiations and data path encryption.

TABLE 25 Algorithms supported for IKEv2 negotiations and data path encryption

	Diffie-Hellman group	Encryption	Integrity	Pseudorandom function
IKEv2 algorithms	14	AES-CBC-128	SHA-256	SHA-256
	19	AES-CBC-256	SHA-384	SHA-384
	20			
Data path algorithms		AES-GCM-128	none	none
		AES-GCM-256	none	none

> **NOTE**
> On MLX Series routers any combination of Diffie-Hellman group, encryption algorithm, and integrity algorithm is allowed. For example, you can use DH group 19, AES-CBC-128 as the encryption algorithm, and SHA-384 as the integrity algorithm. MLX Series routers (with the BR-MLX-10GX4-M-IPSEC module installed) interoperate with devices from other manufacturers that support the algorithms listed in the preceding table.

Support for PSK for IKEv2 SAs

Pre-shared key (PSK) authentication is supported for IKEv2 SAs.

The pre-shared key authentication method is set when configuring an IKEv2 authentication proposal.

Both text-based PSK and hexadecimal-based PSK are supported.

Hexadecimal-based PSK is binary bit-based PSK in which the binary bits are configured as hexadecimal digits.

> **NOTE**
> For Network Device Protection Profile (NDPP) with VPN gateway, the PSK can be either text-based or bit-based (entered as hex digits).

PSK values are encrypted using simple base64 encryption. The encrypted form of PSKs is used in both saved and run-time configurations. PSKs are displayed in encrypted format in **show** command output.

Self MAC addresses and IPsec tunnels

When encrypted or decrypted IP packets are looped back to the device for an additional level of encryption or decryption, the packets are sent to the CPU for SA learning.

Since, by default SA learning is enabled and self MAC address learning is disabled, all looped packets are continuously sent to CPU resulting in high utilization. However, the device may be configured to learn self MAC addresses for all configured IPsec IPv4 and IPv6 tunnels by using the **ipsec self-sa-learning-enable** command.

Unicast IPv4 and IPv6 over IPsec Tunnels

Both IPv4 and IPv6 are supported for point-to-point (unicast) communications that are secured by using IPsec.

> **NOTE**
> IPv4 multicast over IPv4 IPsec tunnels is supported. IPv6 multicast over IPv6 IPsec tunnels is not supported.

Using Unicast IPSec IPv4 with Network Address Translation (NAT)

NAT Traversal (NAT-T) enables the set up and management of unicast IPSec IPv4 tunnels across gateways that employ NAT.

Because NAT modifies the values of IP packet headers (including source and destination IP addresses), a typical implementation of IPSec with NAT results in significant loss of data and data integrity (for example, the discarding of IP packets at the tunnel end nodes and invalid checksums).

There are some operational issues when using IPsec with NAT. IPsec ensures the integrity of IP packets by discarding packets that violate integrity checks. When IP packets pass through a NAT device, the IP/TCP header is automatically modified, which causes a violation of integrity, and the packets are discarded by the IPsec tunnel end nodes.

NAT-T supports the features described in the following table.

TABLE 26 NAT-T feature support

Feature	Supported
IP version	IPv4 only
MLX Series	IPsec module
Operating mode	Tunnel mode
High availability	Supported during switchover and Hitless Operating System Switchover (HLOS)
Keepalive messages	Configurable NAT keepalive message time interval

> **NOTE**
> NAT-T supports only static mappings on the NAT device.

The Extreme NAT-T feature is designed for use in topologies where there is either a NAT device between the IKEv2 peers or where one of IKEv2 peers supports NAT functionality.

> **NOTE**
> When a network configuration does not meet at least one of these requirements, the NAT detection phase of the IKEv2 negotiation process fails, resulting in a discontinuation of further probing for NAT-T.

You enable the feature using the **ikev2 nat-enable** command. You must configure the command on both IKEv2 peers to ensure that IKEv2 negotiation is successful. When you enter the command, the negotiation of NAT-T between the IKEv2 peers is initiated. If the negotiation is successful and the tunnel is up, all ESP packets transmitted over the tunnel are encapsulated in the UDP header.

When you disable the feature by using the **no** version of the **ikev2 nat-enable**, any existing IPsec IPv4 tunnels on which NAT-T was enabled are automatically brought down (but not deleted) by the system. IKEv2 negotiations on the tunnel are restarted with NAT-T disabled, and ESP packets sent across the tunnel are no longer encapsulated in the UDP header.

You use the **ikev2 nat keepalive** command to ensure that the NAT mappings are retained (kept alive). You can modify the NAT keep-alive time interval as needed based on network bandwidth.

If you disable the NAT **keepalive** option by using the **no ikev2 nat keepalive** command, the NAT keep-alive value is reset to the default value.

> **NOTE**
> There is no impact for upgrades. There is potential impact for downgrades. If the NAT-T feature is enabled and the configuration is saved, a downgrade to a version that does not support the feature results in a system error, indicating the system detects an unrecognized configuration.
> To prevent this error, remove any configuration for this feature before the downgrade. You can restore the configuration for the feature if after the downgrade, you upgrade again to a version that supports this feature.

IKEv2 configuration for using NAT-T

Internet Key Exchange version 2 (IKEv2) configuration includes settings related to the use of the NAT-T feature.

The following examples describe global IKEv2 configuration commands to enable the NAT-T feature and specify the keep alive time for the NAT mappings. In addition, various IKEv2 **show** commands indicate whether the feature is enabled, the keep alive time for NAT mappings, and so on.

The following example shows the use of the **ikev2 nat-enable** command to enable the NAT-T feature. You must be in global configuration mode to enter the command.

NOTE
You must configure the command on both IKE peers to ensure that IKEv2 negotiation is successful.

```
device(config)# ikev2 nat-enable
```

The following example shows the use of the **ikev2 nat keepalive** command to specify the keep alive time for the NAT mappings. You must be in global configuration mode to enter the command.

In this example, the specified keep alive time is 9 seconds.

```
device(config)# ikev2 nat keepalive 9
```

The following example shows the use of the **show ikev2** command to show the current IKEv2 configuration. The NAT-T feature settings are shown at the bottom of the output.

```
device(config)# show ikev2
IKEv2 Global data:
Retry Count          : 5          Max Exchange Time      : 60
Retransmit Interval  : 5          Cookie Challenge Number : 0
Max Sa In Nego per LP: 256        Max Sa per LP          : 256
Allow Duplicate      : False      Http Cert Enable       : False
Total Peers          : 1          Total Ipsec Intf       : 1
Total Ike Sa         : 1          Total Ipsec Sa         : 2
Sync In Progress     : 0          Time to syn peer to LP : 0
Pending Job          : 0          Sw pending time        : 0
NAT-T enabled        : True          NAT-T keep-alive time   : 5 sec
```

The following example shows the use of the **show ikev2 session** command to show the details for the most recent IKEv2 session. The output from this command indicates whether a NAT device was discovered during the IKEv2 session.

```
device(config)# show ikev2 session
IKE count:1, Child Sa Count:2
tnl-id                local                 remote    status vrf(i)        vrf(f)
-------------------------------------------------------------------------------
tnl 257     151.152.2.51/4500    151.152.2.52/4500 active default-vrf  default-vrf
-------------------------------------------------------------------------------
    Encr: aes-cbc-256, Hash: sha384, DH Grp:384_ECP/Group 20, Auth: pre_shared
    Local spi: 0xa9a3494d5ebb00e9     Remote SPI: 0x3f0cb6041144d971
    Life/Active Time: 0/0 sec
    NAT Discovered: True
    Child Sa:
     id 1
        Local selector  0.0.0.0/0 - 255.255.255.255/65535
        Remote selector 0.0.0.0/0 - 255.255.255.255/65535
        ESP SPI IN/OUT: 0x00019a48/0x000327f7
        Encryption: aes-gcm-256, ICV Size: 16 octects, Esp_hmac: null
        Authetication: null  DH Group:none , Mode: tunnel
```

The following example shows the use of the **show ikev2 sa detail** command to show the details for the current security associations on the IPsec line card. The output from this command indicates whether NAT-T is detected (enabled), and whether the local and remote end nodes are behind the NAT device.

```
LP# show ikev2 sa detail
-------------------------------------------------------------------------------------------------------
--------
tnl 23      100.23.23.2/4500                           100.23.23.1/4500        active vrf3
default-vrf
-------------------------------------------------------------------------------------------------------
---------
Role                 : Responder
Local SPI            : 0x2708deebeec67f20     Remote SPI: 0x84628d4dea442668
Ike Profile          : 23
Ike Policy           : 251
Auth Proposal        : 23
NAT-T is detected
```

```
Local node behind NAT: True
Remote node behind NAT: False
```

Multicast IPv4 over IPsec Tunnels

Multicast IPv4 over IPsec tunnels is used to transport data and control multicast traffic securely by encrypting the multicast IPv4 packets sent over external networks.

> **NOTE**
> The Extreme implementation of this feature currently supports only the tunnel mode.

To transport multicast data packets securely, an IPsec tunnel interface known as the Virtual Tunnel Interface (VTI) is configured on both of the tunnel end points. The VTI connects the two tunnel endpoints and is similar to a GRE tunnel interface. Once the VTI is configured, you can configure PIM on the VTI. This enables the tunnel endpoints to exchange PIM control packets to set up the multicast distribution tree and the multicast forwarding entries on both tunnel end points.

On the VTI ingress node, the native IP multicast data packet, including the IP header (inner), is encrypted using Encapsulating Security Payload (ESP) and encapsulated by another IP header (outer) corresponding to the tunnel end points. The encrypted and encapsulated unicast packet is routed in the external network based on the outer IP destination address. At the VTI egress node, the received IP packet is decapsulated and decrypted to obtain the native IP multicast data packet, and sent downstream of the egress node. The packet continues as a native multicast data packet from the point it egresses the VTI tunnel until it reaches the end receivers.

Multicast control packets sent over VTI interfaces are also encrypted. The encryption and decryption and encapsulation and decapsulation of control packets and data packets takes place at the hardware level. This tunnel interface at the ingress node will be one of the outgoing interfaces (OIFs) for the multicast forwarding cache entry and at the tunnel egress node it is the incoming interface (IIF). PIM control messages sent over the VTI interfaces are also encrypted and encapsulated by the outer IP header. This takes place at the hardware level.

FIGURE 4 Multicast data packet ingressing and egressing the VTI.

In the figure, R3 has two VTI tunnels (blue and green) configured (one to R7 and another to R6). PIM is configured on the tunnel interfaces on R3, R6 and R7, as well as on the non-tunnel interface connecting R3 and R8. The receivers are attached to R8, R7, and

R6. R3 is the VTI ingress for both blue and green tunnels and R7 and R6 are VTI egress nodes for the blue and green tunnels, respectively. As the packet reaches the VTI ingress, the packet needs to be replicated on the two tunnel interfaces (blue and green) and on the interface connecting R8.

The following events occur at R3:

- The multicast packet is encrypted and encapsulated by a unicast IPsec header, which includes an IP Encapsulating Security Payload (ESP) header corresponding to the encryption on the green tunnel and gets forwarded out of the interface which is the next hop interface for the tunnel destination R6.
- Similarly, the packet gets encrypted and encapsulated and sent out of the blue tunnel.
- On the interface connecting R3 and R8, the packet is forwarded as a native multicast packet.

At R6 and R7, the received packet is decapsulated and then decrypted, and the native multicast packet is forwarded down to the interfaces connecting the receivers as a native multicast packet.

At R8, the incoming packet is a native multicast packet and gets forwarded as a native multicast packet. The PIM control packets exchanged between R3 and R6 and R3 and R7 are also encrypted and encapsulated.

Supported features

The following features are supported as part of IP multicast over IPsec tunnels:

- PIM-SM and SSM, PIM-DM, and static IGMP groups on the VTIs
- PIM over multi-VRF (different outer and inner VRF) IPsec tunnel VTIs

Unsupported features

The following features are not supported as part of IP multicast over IPsec tunnels:

- IPv6 multicast over IPv6 or IPv4 VTI
- IGMP over IPv4 VTI (Static IGMP groups configuration continues to be supported)
- Multicast security associations (Only point-to-point unicast VTI is supported)
- Multicast traffic ingress on the IPsec tunnel on the BR-MLX-10Gx4-M-IPSEC module and egress on GRE tunnel on the BR-MLX-10Gx24-DM 24-port module.

Configuration considerations for multicast over IPsec tunnels

The following configuration considerations apply to the multicast over IPsec tunnels functionality:

- The **ip pim tunnel rpf-strict** command is not supported on PIM-enabled IPsec VTIs.
- The PIM fast hello timer is not supported with IPsec tunnels.
- The multicast source and rendezvous point must be reachable through the tunnel for multi-VRF instances, and the source and receivers should be reachable on the same VRF. A source and receivers residing on different VRFs are not supported over IPsec tunnels and by native multicast.
- The maximum number of PIM IPsec tunnel interfaces is 128 at the system level.
- IPsec tunnels currently support 2000 multicast cache entries.
- Multicast over IPsec tunnels depends on the unicast infrastructure for the IPsec tunnels. Thus, the unicast limitations for IPsec tunnels are also applicable to this functionality. For example, IPsec tunnels over multi-slot LAGs are not supported, so the same restriction applies to multicast.

Multi-VRF support

For IPsec VTIs, the source and destination addresses of the outer IPv4 or IPv6 header of the tunneled packet can reside on a different VRF from the VRF for which the packet is received (including the default global VRF) or on the same VRF for which the packet is received.

PIM is supported on these VTIs.

IPsec Scalability Limits

There are two scalability limits that can affect your use of IPsec.

One is the maximum number of IPsec tunnel elements (such as IKE sessions, SAs, and policies and proposals), and the second is the maximum number of tunnels. The system and the IPsec line cards you are using have both IPsec tunnel element thresholds, and number of total tunnel thresholds.

If you exceed the IPsec tunnel element thresholds for an IPsec line card or for the system, one or more of the IPsec tunnels you are trying to set up cannot be used to transmit protected IP packets.

> **NOTE**
> The maximum number of IPSec tunnels for the system is 2048. The maximum number of IPsec tunnels for an IPsec line card is 256.

IPsec tunnel elements are the various components that make up an IPsec tunnel (for example, IKE sessions, SAs, and policies and proposals). The following tables list the IPsec tunnel element thresholds for IPsec (applies to IPv4 and IPv6).

The system thresholds are based on a system with the maximum number of IPsec line cards (8 line cards). If you have fewer than 8 IPsec line cards, the system thresholds will be lower than the listed values.

TABLE 27 IPsec tunnel element thresholds

Tunnel Element	System Threshold	Line Card Threshold
IKE Peer/ IKE Session/IKESA	2048 (this includes any IPsec tunnels that are used to carry L2 traffic)	256
IPsec SA	4096 (this includes any IPsec tunnels that are used to carry L2 traffic)	512
IKE Proposal	512	512
IKE Policy	512	512
IKE Profile	512	512
IPsec Proposal	512	512
IPsec Profile	512	512

An error message displays when you attempt to configure an IPsec tunnel that exceeds the system tunnel limit.

For example, for a system with the maximum number of IPsec line cards (8), the maximum number of IPsec tunnels you can have is 2048. If you try to configure more than 2048 tunnels, the system displays the following error message:

```
Error - Max number (2048) of ipsec tunnels already configured.
```

If this occurs, you must remove an existing IPsec tunnel in order to configure another tunnel.

Supported Features and Functionality

The features and functionality needed to set up IPv4 and IPv6 IPsec tunnels and transmit protected IP packets across the tunnels are provided. Depending on the version of IP you use (IPv4 or IPv6), there are small differences in the features and functionality.

The following table lists the supported features and functionality for using IPv4 and IPv6 over IPsec.

Feature	IPv4	IPv6
Mixing of IPv4 IPsec and IPv6 tunnels on the device and on the IPsec line card (BR-MLX-10GX4-IPSEC-M). The maximum number of IPsec tunnels is 2000 per system and 256 for an IPsec line card.	Yes	Yes
Independent security configuration for IPv4 IPsec and IPv6 IPsec tunnels. This enables you to have different levels of security on IPv4 IPsec tunnels than you do on IPv6 IPsec tunnels.	Yes	Yes
Static point-to-point tunnel setup between two IP endpoints using IKEv2	Yes	Yes
Single slot LAG (on IPsec line card)	Yes	Yes
Multiple LAGs on different IPsec line cards when: • Each LAG has all physical ports on same IPsec line card.	Yes	Yes
Tunnel destination is reachable using VE interface and the VE interface has physical ports located on multiple IPsec line cards.	Yes	Yes
Dead Peer Detection (DPD) using IKEv2 Keep Alive	Yes	Yes
Configurable options for tunnel elements, such as IKE SA Lifetime, IKEv2 Keep Alive	Yes	Yes
ESP replay window (used to check received packets when replay is enabled)	Yes	Yes
Replay protection for each SA (by default anti-replay check is disabled)	Yes	Yes
64 bit ESN (extended sequence number) for ESP when replay is enabled	Yes	Yes
Different Quality of Service settings on IPv6 IPsec tunnels	Yes	Yes
VRF forwarding		
Source and destination addresses of the outer header of the tunneled packet can be: • In a different VRF from the VRF for which the packet is received (including the default global VRF). • In same VRF that receives the packet.	Yes	Yes
Configurable VRF (outer header of the packet)	Yes	Yes
Multi-VRF forwarding to same remote end point (Multiple IPsec tunnels are set up on the same remote endpoint, one for each inner VRF. Also, a separate IKE session is set up for each IPsec tunnel.)	Yes	Yes
Address translation		
Network Address Translation (NAT) NOTE Requires that you enable NAT-T feature when setting up IPsec IPv4 tunnels.	Yes (unicast traffic only)	No
Protocols		
Encapsulation Security Protocol (ESP) in tunnel mode	Yes	Yes
IKE (for tunnel setup and key management)	IKEv2 only	IKEv2 only
BGP4	BGP4	BGP+
OSPF	OSPFv2 only	OSPFv3 only
RIP	RIPv1/v2 only	RIPng
Cryptography	**IPv4**	**IPv6**

Suite B cryptography to provide Top Secret, 192 bits strength security	Yes	Yes
AES-128-GCM and AES-256-GCM For ESP combined mode authentication, encryption, decryption of data and control packets	Yes	Yes
AES-CBC-256, AES-CBC-128 (confidentiality)	Yes	Yes
Diffie Hellman groups (key exchange). You can accept the default (group 20), or select one or multiple groups, including group 14 and 19.	Yes	Yes
Elliptical Curve Digital Signature Algorithm (ECDSA) P-384 and P-256 For Digital Signature generation and verification	Yes	Yes
SHA-384 and SHA256 (IKEv2)	Yes	Yes
FPGA encryption and decryption (no encryption or decryption is done by software)	Yes	Yes
Line rate support (encryption and decryption at 44G line rate)	Yes	Yes
Interface to interface traffic forwarding (IP packets)	**IPv4**	**IPv6**
• Link local and /127 addresses	N/A	Link Local IPv6 address and /127 IPv6 addresses
• Jumbo Frame support	Yes	Yes
• Non-IPsec module to IPsec module packet forwarding, **and** • IPsec module to non-IPsec module packet forwarding	Yes	Yes
• Physical interface to IPsec tunnel interface	Yes	Yes
• VE interface to IPsec tunnel interface	Yes	Yes
• IPsec IPv4 tunnel interface to IPv4 IPsec tunnel interface	Yes	N/A
• IPsec IPv6 tunnel interface to IPv6 IPsec tunnel interface	N/A	Yes
• Manual IPv6 tunnel interface to IPv6 IPsec tunnel interface	N/A	Yes
• 6to4 tunnel interface to IPv6 IPsec tunnel interface	N/A	Yes
Authentication	**IPv4**	**IPv6**
IP peer authentication using PKI	Yes	Yes
Endpoint authentication using pre-shared keys and digital certificates (ECDSA)	Yes	Yes
IPsec statistics	**IPv4**	**IPv6**
Packet counts and byte counts, including: • Transmit and Receive packet counts for each tunnel. • Byte counts for each tunnel.	Yes	Yes
IKEv2 packet counters, including IKEv2 Keep Alive packets.	Yes	Yes
Traps and syslogs	**IPv4**	**IPv6**
UP/DOWN traps and syslogs.	Yes	Yes
Errors counters and syslog for specific ESP errors.	Yes	Yes

Unsupported features

Some features are not supported for unicast IP communications over IPsec.

TABLE 28 Unsupported features for unicast communications over IPsec

Feature	Description
Tunnel setup (IPv4 and IPv6)	Because of a hardware limitation and the fact that IPsec tunnels can only carry IP inner packets, IPv4 or IPv6 IPsec tunnels cannot be set up if the remote endpoint can only be reached by: • GRE tunnel

TABLE 28 Unsupported features for unicast communications over IPsec (continued)

Feature	Description
	• VEoVPLS uplink • MPLS tunnel • IGP shortcut • Manual IPv6 tunnels • 6to4 tunnel • IPsec tunnel (IPv4 or IPv6)
Multicast IPv6 over IPv6 IPsec tunnels	Multicast IPv6 traffic over IPv6 IPsec tunnels is not currently supported. NOTE Multicast IPv4 traffic over IPv4 IPsec tunnels **is** currently supported.
ESP in transport mode	ESP in transport mode is not currently supported.
Cryptography (other than Suite B)	Cryptography algorithms other than Suite B cryptography are not currently supported.

Limitations

There are some limitations that impact the use of IPsec for creating secure tunnels.

> NOTE
> The limitations apply to IPv4 and IPv6, unless indicated otherwise.

- Encryption and decryption of packets sent or received on an IPsec tunnel is not supported by the CPU of the LP or MP module.
- Encryption and decryption of IP packets and the tunnel setup using IKEv2 can be done only on the BR-MLX-10GX4-M-IPSEC line card.
- Global and link local IPv6 addresses can be configured only on IPv6 IPsec tunnel logical interfaces. IPv4 address configuration is not supported on IPv6 IPsec tunnel logical interfaces.
- Incoming and outgoing paths to tunnel endpoints must be on same IPsec module interface card.
- Fragmentation is not supported when traffic is routed over an IPSec tunnel; a fragmented IPsec packet received on an IPv4 IPsec tunnel is dropped because IPsec packets are not re-assembled before decryption.
- GRE and IPsec encapsulation are not performed together for the same flow in the same device.
- For LAG, all physical ports must be on same IPsec module (line card).
- LAG across IPsec and non-IPsec line cards is not supported.
- Multi-slot LAG using two or more IPsec LP modules is not supported.
- IPsec supports LAG only if the anti-replay check is disabled. If anti-replay check is enabled, there is a possibility of packet drop due to packet reordering and the small window size of 64 packets. To minimize the packet reordering issue, it is recommended to for the LAG using on of the following groups of ports: ports 1, 2, 5, and 7 or port 3, 4, 7, and 8.
- If the anti-replay check is enabled, an anti-replay check or integrity check failure can occur in certain topologies where packet drop or reordering of packets takes place, and the packet reordering is beyond the 64-packet replay window. In this case, the anti-replay check, and the extended sequence number can be disabled. In non-LAG cases, the IPsec tunnel TOS can be used instead of using the inner packet TOS.
- The anti-replay check and ESN settings must be identical at both endpoints for a given IPsec tunnel. If they are not identical, the IPsec tunnel will not operate correctly. In addition, the anti-replay check and ESN settings should match each other's

configuration (both should be disabled or enabled). Because the settings are independent, you must configure them so they are identical.

- When multiple IPSec tunnels are configured on the same device, each IPsec tunnel must have a unique tunnel source and tunnel destination combination.
- For each protected inner VRF (IVRF), there is one IPsec tunnel and at least one IKE session. These are needed to ensure binding of the child SA to the VRF.
- IP packet forwarding for IPsec tunnels in a non-default VRF is not supported during Hitless Operating System Switchover (HLOS).
- Equal-cost multi-path routing (ECMP) is not supported on single (individual) IPsec tunnels. Tunnels are setup using the first path available to reach the remote endpoint of the tunnel. All traffic transported on the tunnel use the same path as the path used to setup the tunnel.

The use of ACLs to control IPsec tunnel traffic

Access control lists (ACLs) can be applied to incoming or outgoing traffic on specific ports.

ACLs cannot be applied to IPsec tunnel interfaces. To control the traffic entering an IPsec tunnel, you can apply ACLs on the interfaces where the traffic (that is destined for the IPsec tunnel) is received.

Similarly, you can apply ACLs to underlying interfaces that are used to filter traffic received on an IPsec tunnel. These ACLs are applied to the inner IP packet (after decryption and de-encapsulation of the packet received on the IPsec tunnel).

You configure ACLs on a global basis and apply them to the incoming or outgoing traffic on specific ports. You can apply only one IPv4 ACL to a port's inbound traffic and similarly, only one IPv4 ACL to a port's outbound traffic. The software applies the entries within an ACL in the order they appear in the ACL's configuration. As soon as a match is found, the software takes the action specified in the ACL entry (permit or deny the packet) and stops further comparison for that packet. For further information on how to configure ACLs refer to *ACLs*.

> **NOTE**
> To configure the Target of Evaluation (TOE) to meet the Security Requirements for Network Devices (NDPP) with VPN gateway requirements, the last rule configured for every interface must be the one which denies all traffic that is not already explicitly permitted by a previous rule.

SPD rules

Each security policy database (SPD) rule requires two ACL rules, one defining the traffic flows on the tunnel, and one defining the traffic flows within the tunnel. The protocol to which the SPD rule applies must be identical for both ACL rules, and the sequence number used with the ACL rules, should be consecutive.

SPD rule	Action	Address	Protocol	Sequence number
BYPASS	Permit	Public address	Applicable protocol	N
	Deny	Tunnel internal address	Applicable protocol	N + 1
PROTECT	Deny	Public address	Applicable protocol	N
	Permit	Tunnel internal address	Applicable protocol	N + 1
DISCARD	Deny	Public address	Applicable protocol	N
	Deny	Tunnel internal address	Applicable protocol	N + 1

Impact of upgrades or downgrades of NetIron

There is an important interoperability issue that can affect IPsec tunnels and may result in tunnel shutdown and traffic loss. There are steps you can take to avoid this issue as well as steps you can take to restore the tunnels to full functionality.

The potential issue occurs in situations in which all of the following conditions exist:

- One tunnel node is running version 5.9 and the other node is running version 5.8 or 5.8x.
- Both tunnel nodes are using the default IKEv2 authentication proposal.

Because a single tunnel node can be upgraded independently of the other node of the tunnel, you can end up with a configuration in which one tunnel node is running version 5.9 and the other node if running version 5.8 or 5.8x.

> **NOTE**
> This issue occurs even if you have enabled the NetIron compatibility option on the node running version 5.8 or 5.8x. (This option is designed to enable version 5.8 or 5.8x to interoperate with version 5.9.)

The cause of the issue is that the default IKEv2 authentication proposals are different in version 5.9 and version 5.8 (or 5.8x). In 5.9, the default method for the IKEv2 authentication proposal is pre-shared key (the value is "def-ike-pre-shared-password").

There are two options you have for resolving this issue. They are:

- Upgrade all nodes (boxes) simultaneously to version 5.9.
- Make sure that both nodes have the same IKEv2 authentication proposal configuration.

Support for PKI

Public Key Infrastructure (PKI) provides certificate management to support secured communication for security protocols such as IP security (IPsec).

A PKI is composed of the following entities:

- Peers communicating on a secure network.
- At least one public-private key pair for the local certificate
- In the case of dynamic PKI, at least one Certificate Authority (CA) that grants and maintains certificates.
- Digital certificates, which contain information such as the certificate validity period, peer identity information, encryption keys that are used for secure communications, and the signature of the issuing CA.
- An optional registration authority (RA) to offload the CA by processing enrollment requests.
- A distribution mechanism such as Lightweight Directory Access Protocol (LDAP) for certificate revocation lists (CRLs).

PKI provides customers with a scalable, secure mechanism for distributing, managing, and revoking encryption and identity information in a secured data network. Every entity participating in the secured communications is enrolled in the PKI, a process where the device generates a key pair (one private key and one public key) using an asymmetric encryption algorithm and has their identity validated by a trusted entity (also known as a CA or trust point).

After each entity enrolls in a PKI, every peer in a PKI is granted a digital certificate that has been issued by a CA. When peers must negotiate a secured communication session, they exchange digital certificates. Based on the information in the certificate, a peer can validate the identity of another peer and establish an encrypted session with the public keys contained in the certificate.

PKI terminology

Various terms are used to describe elements of a Public Key Infrastructure (PKI).

TABLE 29 PKI terminology

Term	Description
Certificate	A public key certificate (also known as a digital certificate or identity certificate) is an electronic document used to prove ownership of a public key. The certificate includes information about the key, information about its owner's identity, and the digital signature of an entity that has verified the certificate's contents are correct. If the signature is valid, and the person examining the certificate trusts the signer, then they know they can use that key to communicate with its owner. Certificates have a finite lifetime, defined by the start and end time within the certificate. The certificate is invalid if it is outside of that lifetime. • Certificate authority (CA) certificates are further classified into three classes: cross-certificates, self-issued certificates, and self-signed certificates. Cross-certificates are CA certificates in which the issuer and subject are different entities. Cross-certificates describe a trust relationship between the two CAs. Self-issued certificates are CA certificates in which the issuer and subject are the same entity. Self-issued certificates are generated to support changes in policy or operations. Self-signed certificates are self-issued certificates where the digital signature may be verified by the public key bound into the certificate. Self-signed certificates are used to convey a public key for use to begin certification paths. • End entity certificates are issued to subjects that are not authorized to issue certificates.
Certificate authority	A certificate authority (CA) is an authority in a network that issues and manages security credentials and public keys for message encryption. As part of a PKI, a CA checks with a registration authority (RA) to verify information provided by the requestor of a digital certificate. When the RA verifies the requestor's information, the CA can then issue a certificate. The CA also specifies the validity periods of certificates, and revokes certificates as needed by publishing certificate revocation lists (CRLs).
Certificate revocation list	A Certificate revocation list (CRL) is a list of signed certificates (signed by the CA) that are now prematurely invalid.
CRL distribution point	A CRL distribution point (CDP) is used to retrieve a CA's latest CRL. A CDP is usually an LDAP server or HTTP (web) server. The CDP is normally expressed as an ldap://host/dir or http://host/path URL.
Distinguished name	Distinguished name (DN) is a set of fields and values that uniquely define a certificate and VPN gateway or RAS VPN client identity. This is sometimes called the "Subject" of the certificate. The DN identity can be used as the Internet Key Exchange (IKE) ID.
Entity	An entity is an end user of PKI products or services, such as a person, an organization, a device such as a switch or router, or a process running on a computer.
Lightweight Directory Access Protocol	Lightweight Directory Access Protocol (LDAP) is used for accessing and managing PKI information. An LDAP server stores user information and digital certificates from the RA server and provides a directory navigation service. From an LDAP server, an entity can retrieve local and CA certificates of its own as well as certificates of other entities.
PKI repository	A PKI repository can be a Lightweight Directory Access Protocol (LDAP) server or a common database. It stores and manages information such as certificate requests, certificates, keys, CRLs, and logs while providing a simple query function.
Registration authority	A registration authority (RA) is an authority in a network that verifies user requests for a digital certificate and tells the Certificate Authority (CA) to issue it. An RA can implement functions including identity authentication, CRL management, key pair generation and key pair backup. The PKI standard recommends that an independent RA be used for registration management to achieve higher security of application systems.
Requester	A requester is the client of Simple Certificate Enrollment Protocol (SCEP) exchanges, such as exchanges involved in requesting and issuing certificates. The requester typically submits SCEP messages for itself (it can also submit SCEP messages for peers).

Certificate enrollment using SCEP

Extreme provide functionality for certificate enrollment through the use of the Simple Certificate Enrollment Protocol (SCEP). Certificate enrollment is an essential PKI operation in which a system requester successfully receives a certificate from the Certificate Authority (CA).

SCEP is designed to supports the following PKI operations:

•   Certificate Authority (CA) public key distribution

- Registration Authority (RA) public key distribution
- Certificate enrollment
- Certificate query
- Certificate Revocation List (CRL) query

Types of enrollment

There are two basic types (modes) of enrollment. They are:

- Manual enrollment (manual mode)

    In manual mode, the requester's messages are placed in the PENDING state until the CA operator (person) authorizes or rejects them. Manual authorization is used when the client has only a self-signed certificate, or a challenge password is not available.

- Automatic enrollment (auto-enrollment)

    In automatic mode, all messages between the requester and the CA are managed by the network resources. This mode is used to efficiently enroll many system requesters, and the requester systems usually have a common, templated configuration.

Requirements for requesting a certificate

There are a few requirements that must be satisfied before a requestor can request a certificate.

The requirements are:

- The requester must have at least one appropriate key pair (for example, an EC key pair).
- The following information must be configured locally on the requester (client).
    - The CA IP address, or fully qualified domain name (FQDN).
    - The CA HTTP Computer Gateway Interface (CGI) script path.
    - The identifying information used to authenticate the CA. This can be obtained from the user or provided (presented) to the end user for manual authorization during the exchange.

        > NOTE
        > Multiple independent configurations that contain this information (items 1, 2, and 3) can be maintained by the requester if needed to enable interactions with multiple CAs.

Communications between requesters and the CA

Communications between requesters and the CA are made secure using SCEP Secure Message Objects, which specifies how the PKCS #7 cryptographic message syntax is used to encrypt and sign the data. To ensure that the signing operation can be completed, the client uses an appropriate (valid) local certificate.

The following rules apply to the communications between requesters and the CA.

- If the requester already has a certificate issued by the SCEP server and the server supports certificate renewal, the certificate that was already issued should be used (renewed).
- If the requester does not have a certificate issued by the new CA, but does have credentials from an alternate CA, the certificate issued by the alternate CA may be used.

    > NOTE
    > In this case, policy settings on the new CA determine whether or not the request can be accepted. It can be beneficial to use a certificate from the old domain as credentials when enrolling with a new administrative domain.

- If the requester does not have an appropriate existing certificate, then a locally generated, self-signed certificate must be used. The self-signed certificate must use the same subject name used in the PKCS #10 request.

Offline certificate enrollment and import of certificates and CRLs

Public Key Infrastructure (PKI) certificate enrollment can be done offline. Certificates and certificate revocation lists (CRLs) can also be manually imported.

When a certification authority (CA) server is offline or when Simple Certificate Enrollment Protocol (SCEP) is disabled or not supported, certificates can be manually imported. When a certificate revocation list (CRL) distribution point is offline or when Online Certificate Status Protocol (OSCP) is disabled or not supported, CRLs can be manually imported .

When a CA or trustpoint server is online, an MLXe device uses the SCEP protocol to generate a certificate signing request (CSR) in Public Key Cryptography Standards (PKCS) #7 format and send it to the CA server over HTTP. The server then processes the CSR and returns the SCEP response which contains the client certificate. When the CA server is offline, this process can be done manually by generating the CSR on the MLXe device in PKCS #10 format and manually transporting it to the CA server to obtain the client certificate. The X.509 client certificate is then taken back to the MLXe device and manually imported into the device either by loading from flash memory or by pasting onto the device terminal. When using the option to import a certificate by pasting onto the device terminal you must first enable enrollment terminal for the trustpoint by issuing the **enrollment terminal** command.

When a CRL distribution point (CDP), which is typically the CA server also, is online and an MLXe device needs to check the revocation status of a certificate for a peer device, it dynamically downloads the CRL from the distribution point over HTTP. When the CRL distribution point is offline, a CRL can be loaded manually by copying the CRL from the CRL distribution point server, transporting it to the MLXe device and importing it into the device either by loading from flash memory or pasting onto the device terminal.

> NOTE
> When validating a certificate chain for a peer device any CRL for an issuer in the certificate chain that is not present in the trustpoint, is dynamically downloaded when the relevant CDP is online.

In network environments where file copy protocols such as Secure Copy Protocol (SCP) or Trivial File Transfer Protocol (TFTP) are not available, the option to manually import certificates and CRLs by pasting onto the device terminal can be used.

Support for Certificate Path Construction and Validation

You can ensure that your system's certificate management complies with the requirement for Network Devices profile protection (NDPP) VPN Gateway Certification.

The essential requirement for VPN Gateway Certification includes two main items, which are:

- A chain of certificates, or a certification path between the certificate and an established point of trust (typically a Certificate Authority), must be established.
- Every certificate within the certificate path must be validated.

The process used to satisfy the requirements for VPN Gateway Certification is referred to as **certificate processing**. The Extreme certificate processing process involves two phases, which correspond to the two items of the VPN Gateway Certification requirement. The two phases of the Extreme certificate processing process are:

- Path construction
- Path validation.

Path construction phase

During this phase, one or more paths, called candidate certification paths, are built that are used to forward certificates to and from the entities involved in the processing of certificates. The candidate paths are alternative paths that can be used to process the certificates.

In some cases, the system attempts to identify an acceptable certification path as the path is created. This helps to maximize the probability that the candidate path can be validated as well as reduce the time required to create an acceptable path. The process used to identify an acceptable path during path construction is referred to as certification path processing.

> **NOTE**
> Although the certificates may be linked or chained properly, the candidate certification paths may not be valid in terms of the path length, name, or certificate policy constraints or restrictions.

Path validation phase

Once the candidate certification path (or paths) are built, the system initiates the path validation phase.

During this phase, the system checks each certificate in the path to make sure the certificates are valid. Among other checks, each certificate is checked to make sure that:

- The validity period for the certificate has not expired
- The certificate has not been revoked
- The certificate's integrity is intact
- Constraints levied on part or all of the certification path have not been violated. Constraints can include path length constraints, name constraints, and policy constraints.

Operations that occur during the path validation

A candidate certification path must perform an operation called **name chaining**. Name chaining occurs between the recognized trust point (for example, a CA), and the target certificate (the end-entity certificate).

The purpose of the name chaining operation is to ensure that the certificates in the path are ordered or sequenced correctly. To do this, the system uses the Subject Name and Issuer Name in certificates to chain or link the certificates in the path.

The requirement for valid name chaining is that the Subject Name in one certificate must be the Issuer Name in the next certificate in the path. The system checks the name chaining of the certificates in the path beginning at the trust anchor, and moving from there toward the target certificate.

The following figure represents a valid certificate path.
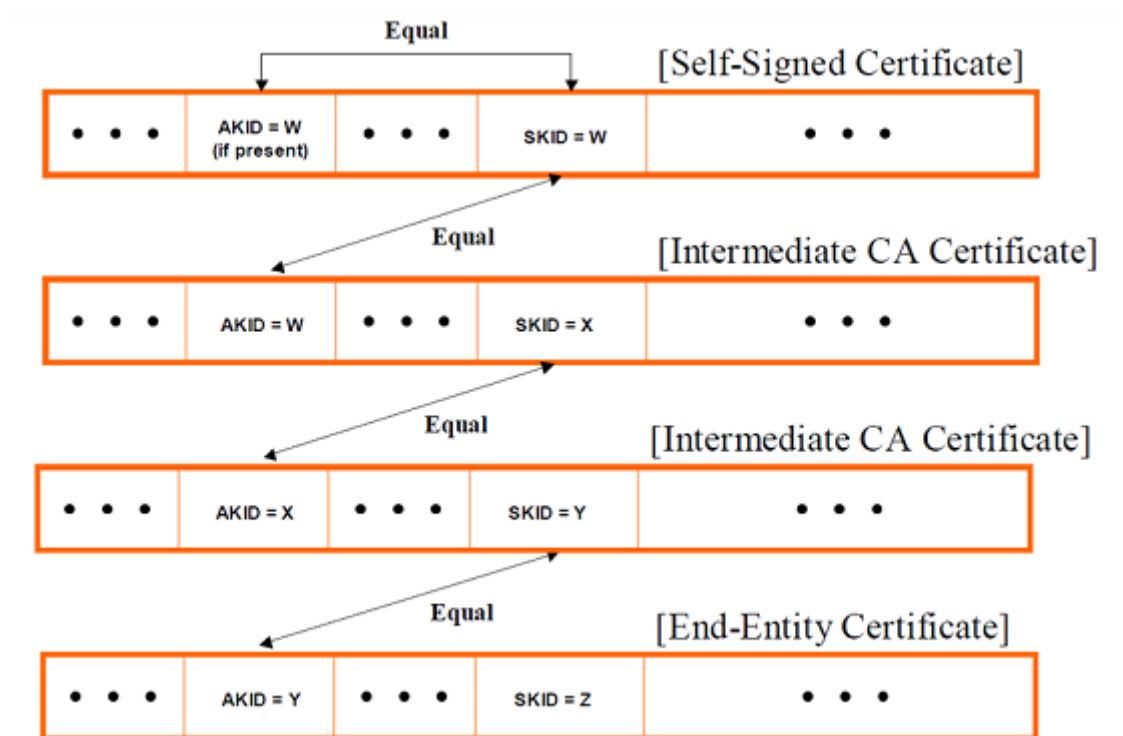
In this example:

- The path begins with a self-signed certificate that contains the public key of the trust anchor.
- The path ends with the end-entity certificate.
- The intermediate CA certificates in the path have correct name chaining (the Subject Name in one certificate is the Issuer Name in the next certificate). The certificates in the path between the two certificates at the beginning and end of the path are referred to as intermediate CA certificates.

Support for hash and URL encoding of X.509 certificates

To avoid fragmentation of User Datagram Protocol (UDP) datagrams, Internet Key Exchange version 2 (IKEv2) supports hash and URL encoding of X.509 certificates.

When hash and URL encoding of X.509 certificates is enabled by using the **ikev2 http-url-cert** command, the IKEv2 authentication process transmits a hash and URL of each certificate (in place of a DER-encoded certificate).

An X.509 certificate is prepared for hash and URL encoding by obtaining a SHA1 hash of the DER-encoded certificate file data, renaming the certificate file to the hash value, and uploading it to a web server for use in the IKEv2 negotiations. The hash and URL must resolve to the DER-encoded certificate.

IKEv2 supports both sending and receiving certificates that are encoded in this manner.

Support for Logging IKE and PKI Transaction Details

Logging of IKE and PKI transaction details is supported. The log files are automatically generated syslog messages that contain the transaction details.

There are two types or levels of logging. Standard (default) logging is enabled by default. The second type of logging is called extended logging, which you must enable using commands. This type of logging allows you to log additional IKE or PKI transaction details.

The hardware requirements are identical for default logging and extended logging. The following table lists the required hardware.

TABLE 30 Hardware requirements for logging IKE and PKI transaction details

Hardware	Requirement
Device	MLXe router
Line card	IPsec module (BR-MLX-10GX4-IPSEC-M) NOTE The MLXe device should have at least one line card through which the external syslog server can be reached.

All of the current limitations of the logging feature on MLX devices, and the limitations of the IPsec security feature apply to the logging of IKE and PKI transaction details.

In addition, there are some limitations specific to the feature for logging IKE and PKI transaction details. The following table lists the current limitations for this feature.

TABLE 31 Limitations for IKE and PKI transaction detail logging

Default and Extended Logging	Description
IKE transaction details (send and receive packets)	The maximum size syslog buffer is 1024 bytes. Logging of packet content in hex format along with other logging parameters in the packet syslog only allow 250 bytes of packet in one syslog. An IKEv2 packet that together with protocol headers totals more than 250 bytes will be logged in multiple syslogs.
Packet and event syslogs for IKE sessions	If a line card on which IPsec tunnels are configured (IPv4 or IPv6 tunnels) is restarted, packet and event data for IKE sessions are lost. The packet and event data are not logged on the local syslog or the remote syslog server.

NOTE
There is important potential impact to the IKE and PKI logging functionality when downgrading the version of NetIron firmware on your system devices. Downgrading the NetIron firmware to a previous version below NI 5.9.00a removes the following logging functionality from the startup configuration:

- IKEv2 extended logging
- PKI standard (default) logging
- PKI extended logging.

Differences in Default and Extended Logging Syslogs

Extended logging of IKE transaction and PKI transaction details enables you to log additional details not included in the default IKE and PKI syslog messages.

NOTE
You must enable extended logging. It is not enabled by default.

Differences in PKI transaction detail syslogs

Nearly all of the information in extended logging syslogs for PKI transactions are also included in default logging syslogs. The extended logging syslogs include a **hex dump** for the message, which is not included in the default logging syslog. The hex dump follows the text of the message.

Differences in IKE transaction detail syslogs

Default logging and extended logging of IKE transaction details provide different information. The information for each logging type includes:

- **Default logging:** The syslogs contain logs of events that occur during the IKE transactions (for example, *Session Established with Peer* and *Session Terminated with Peer*).
- **Extended logging:** The syslogs contain send and receive packets that are exchanged between IKE peers during the IKE transactions. The extended logging syslogs also contain a **hex dump** for the message. The hex dump follows the text of the message.

Logging functionality for VPN gateway certification

Extended logging of IKE transaction and PKI transaction details provides the logging functionality that is required for VPN gateway certification.

The required functionality includes:

- Logging of complete IKE transaction details, including logging of complete IKEv2 packets.
- Logging of PKI transactions between the IPsec device and the certification authority (CA), revocation check responder, or any other external server for importing local certificates or downloading peer certificates. Complete PKI packets are logged.
- Transporting syslogs to an external syslog server over a secure channel using an IPsec tunnel.

    NOTE
    Default logging does not provide this functionality.

IKEv2 traps

IKEv2 traps are supported for IKEv2 error conditions.

The following traps are supported for both IPv4 and IPv6.

- Invalid IKE Message Type Received. The audit log entry for this event will include the Message type, SPI value for IKE SA, date and time received, source address, and destination address in the received packet.
- Invalid IKE Payload Received. The audit log entry for this event will include the Payload type, SPI value for IKE SA, date and time received, source address, and destination address in the received packet.
- Maximum IKE peers limit reached on LP. The audit log entry for this event will include the LP number on which this limit is reached.
- Recovered from maximum IKE peers limit on LP. The audit log entry for this event will include the LP number on which this limit is reached.

IPsec traps

IPsec traps are supported for IPsec error conditions.

The following traps are supported for both IPV4 and IPv6 error conditions:

- No valid Security Association (SA) exists for a session. The audit log entry for this event will include the SPI value, date and time received, source address, destination address, and sequence number.

- A packet offered to ESP for processing appears to be an IP fragment; that is, if the OFFSET field is non-zero or the more fragments flag is set. The audit log entry for this event will include the SPI value, date and time received, source address, destination address, and sequence number.

- Attempt to transmit a packet that would result in sequence number overflow. The audit log entry for this event will include the SPI value, current date and time, source address, destination address, and sequence number.

- The received packet fails the anti-replay checks. The audit log entry for this event will include the SPI value, date and time received, source address, destination address, and sequence number.

- The received packet fails integrity check. The audit log entry for this event will include the SPI value, date and time received, source address, destination address, and sequence number in the received packet.

- The de-encapsulation of received packet failed. The audit log entry for this event will include the SPI value, date and time received, source address, destination address, and sequence number.

- The check on IP packet length failed for the received packet. The audit log entry for this event will include the date and time received, source address, and destination address.

The generation of IPsec traps is controlled; new traps are generated but repeat traps are limited to one trap per minute.

IPsec syslog messages

The following example IPsec syslog messages are supported.

- Sequence number overflow when trying to send packet with SPI <SPI-ID> Source <source-address> Destination <destination-address>

- No IPsec SA found for received packet with Source <source-address> Destination <destination-address> SPI <SPI-ID>

- Received fragmented packet with Source <source-address> Destination <destination-address> SPI <SPI-ID>

- Integrity check failed for received packet with Source <source-address> Destination <destination-address> SPI <SPI-ID> Sequence Number <sequence-number>

- ESP: De-encapsulation failed for received packet with Source <source-address> Destination <destination-address> SPI <SPI-ID> Sequence Number <sequence-number>

- ESP: Length error detected for received packet with Source <source-address> Destination <destination-address>

- IKEv2: Invalid message type received with Source <source-address> Destination <destination-address> SPI <SPI-ID> MessageType <x>

- IKEv2: Invalid payload type received with Source <source-address> Destination address type <type> Destination <destination-address> SPI <SPI-ID> PayloadType <x>

- IKEv2: Maximum IKE peers limit reached on LP <n>

- IKEv2: Recovered from maximum IKE peers limit condition on LP <n>

In the syslog examples, x is the value of an unsupported payload type in an IKEv2 packet. It is a UNIT8 value. The value will not be 0, or 32 through 42, which are current valid payload types. And n is the LP module number, which ranges from 1 through 32.

The generation of IPsec syslog messages is controlled. A new syslog message is generated but repeat syslog messages are limited to one per minute.

The following example show syslog message output.

```
May 20 23:24:22:N:ESP: Length Error Detected for Received Packet with Source 100.1.1.3 Destination
100.1.1.13

May 27 19:49:57:I:ESP:  Received Fragmented Packet with Source 51.54.41.51 Destination 51.54.41.54 SPI
0x2804c68

May 27 10:57:16:I:ESP: Anti-Replay Check Failed for Received Packet with Source 51.54.201.54 Destination
51.54.201.51 SPI 0xd4622ffc Sequence Number 6

May 27 10:56:08:I:ESP: Anti-Replay Check Failed for Received Packet with Source 51.54.201.54 Destination
51.54.201.51 SPI 0xd4622ffc Sequence Number 2

May 27 11:04:53:I:ESP: Integrity Check Failed for Received Packet with Source 51.54.201.54 Destination
51.54.201.51 SPI 0xd4622ffc Sequence Number 23

May 27 14:32:53:I:ESP: No IPsec SA Found for Received Packet with Source 51.54.201.54 Destination
51.54.201.51 SPI 0x93255201

May 28 15:40:22:I:IKEv2: Invalid Payload Type Received with Source 51.54.41.54 Destination 51.54.41.51 SPI
0x97b682b9 PayloadType 53

May 28 15:39:50:I:IKEv2: Invalid Message Type Received with Source 51.54.201.54Destination 51.54.201.51 SPI
0xd251c58f MessageType 0
```

IKE and PKI Default and Extended Logging Syslog Messages

By default, the system automatically generates syslog messages that contain details of events that occur during the IKE and PKI transactions involved in the setup and teardown of IPsec tunnels. If you enable extended logging of IKE and PKI transaction details, the syslogs contain additional details not included in the default logging syslogs.

Format of syslog messages

The format of all IKE or PKI transaction syslog messages saved to the local syslog buffer or sent to an external syslog server over an IPsec tunnel are the same.

The format of the IKE or PKI transaction syslog messages is:

Month Date HH:MM:SS: <message-level>: Protocol: <Event or Packet dump in Hex>

Example Jul 30 01:51:20:I:ESP: De-encapsulation Failed for Received Packet with Source <source address> Destination <destination-address> SPI <SPI-ID> Sequence Number <sequence-number>

> NOTE
> Local logging on the management module is determined by the configuration Syslog feature on the MLXe device. For example, if time and date are configured on the onboard system clock, the date and time are shown in syslogs in the format: mm dd hh:mm:ss. Otherwise, date and time are shown in syslogs in the format: numd numh numm nums.

Basic types of transaction detail syslogs

The different types of IKE and PKI transaction detail syslogs are:

- IKEv2 transaction default logging syslogs
- IKEv2 transaction extended logging syslogs
- PKI transaction default logging syslogs
- PKI transaction extended logging syslogs

IKEv2 transaction default logging syslogs

The syslogs contain logs of events that occur during the IKE transactions. Default logging is enabled by default.

The following table lists the IKEv2 default logging syslog messages.

Type	Description
Session Established	This message includes tunnel ID, tunnel source address and tunnel destination address. **Example** `Aug 12 01:51:20:I: IKEv2: Session Established for TNL <Tunnel ID> with Src <source-address> Dest <destination-address>`
Session Terminated	This message includes tunnel ID, tunnel source address, tunnel destination address, and SPI. **Example** `Aug 12 01:58:25:I: IKEv2: Session Terminated for TNL <Tunnel ID> with Src <source-address> Dest <destination-address> SPI <SPI number>`
IKE Session Rekey	This message includes tunnel ID, tunnel source address, tunnel destination address, and SPI. **Example** `Aug 12 01:55:30:I: IKEv2: Session Rekeyed for TNL <Tunnel ID> with SPI <SPI-ID> Src <source-address> Dest <destination-address>`
IPSec SA Rekey	This message includes tunnel ID, tunnel source address, tunnel destination address, and SPI. **Example** `Aug 12 01:55:30:I: IKEv2: IPSEC SA Rekeyed for TNL <Tunnel ID> with SPI <SPI-ID> Src <source-address> Dest <destination-address>`
Timer Expiration	This message includes tunnel ID, tunnel source address, and tunnel destination address. **Example** `Aug 12 01:56:40:I: IKEv2: Session Timer Expired for TNL <Tunnel ID> with Src <source-address> Dest <destination-address>`
Authentication Failure	This message includes tunnel ID, tunnel source address, and tunnel destination address. **Example** `Aug 12 01:51:30:I: IKEv2: Authentication Failed for TNL <Tunnel ID> with Src <source-address> Dest <destination-address>`

IKEv2 transaction extended logging syslogs

The syslogs contain logs of the send packets and receive packets that occur during the IKE transactions. Extended logging is not enabled by default.

The following table lists the syslog messages that are generated when extended logging for IKE transactions is enabled. In these messages, X represents the exchange type, which can have one of the values:

- IKE_SA_INIT
- IKE_AUTH
- CREATE_CHILD_SA
- INFORMATIONAL are not logged.

> **NOTE**
> Messages with empty payloads, for example, IKEv2 keepalive messages are not logged.

Type	Description
Receive Packets	Packets received by IKE peer during the IKE transaction. **Example** `Aug 12 01:51:20:I: IKEv2: Pkt rcvd with Src <source-address> Dest <destination- address> SPI <SPI-ID> Exchg <X> Pkt len: XX Seq: YY Pkt content <Hex dump of the packet from L2 header>`
Send Packets	Packets sent by IKE peer during the IKE transaction.

Type	Description
	Example `Aug 12 01:51:21:I: IKEv2: Pkt sent with Src <source-address> Dest <destination- address> SPI <SPI-ID> Exchg <X> Pkt len: XX Seq: YY Pkt content <Hex dump of the packet from L2 header>`
Fragmented Packets	Send or Receive packets that were fragmented during the IKE transaction. **Example** `Aug 12 01:51:21:I: IKEv2: Pkt sent with Src <source-address> Dest <destination-address> SPI <SPI-ID> Exchg <X> Pkt len XX Seq:YY Frag:Z Frag len: ZZ Pkt Content <Hex dump of the packet from L2 header>`

PKI transaction default logging syslogs

The syslogs contain logs of events that occur during the PKI transactions. Default logging is enabled by default.

The default logging syslogs contain almost all of the information in extended logging syslogs for PKI transactions. The extended logging syslogs include a hex dump for the message, which is not included in default logging syslogs.

PKI transaction extended logging syslogs

The syslogs contain logs of events that occur during PKI transactions and of packets sent and received during PKI transactions. Extended logging is not enabled by default.

Extended logging syslogs include a hex dump for the message, which is not included in default logging syslogs. The hex dump follows the text of the message.

The types of syslogs are:

- Connection establishment events
- PKI user certificate request packets
- PKI authentication packets
- Certification Revocation List (CRL) and Online Control Status Protocol (OCSP) packets
- Manual import of certificate from external server, and peer certificate download
- Close connection requests

  NOTE
  Information on fragmented packets is displayed only if packet is fragmented to fit in syslog.

The following tables list the PKI extended logging syslog messages.

TABLE 32 Connection establishment events

Event	Description
Connection Request	All connection request events to external entities are logged, including: • Connection request events to Certificate Authority (CA) • Connection request events to external server for to download peer certificate • Connection request events to import local certificates **Example** `Aug 12 10:11:12:I: PKI: connection req is sent to host:<hostname> for trust point<trustpoint_name> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`
Connection Successfully Established	All connection successfully established events. **Example** `Aug 12 10:11:12:I: pki: connection to host:<hostname> is success for trustpoint <trustpoint_name> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`

TABLE 33 PKI user certificate request packets

Type	Description

Enrollment Request	All enrollment request packets are logged. The message includes the trust point name, CA name, and request type. The enrollment request types are:
	
	    • PKCSREQ
	    • GETCERTINITIAL
	
	Example `Aug 12 10:11:12:I: PKI: enrollment req PKCSREQ/ GETCERTINITIAL sent for trust point :< trustpoint_name> to CA :< hostname> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`

Enrollment Response	All enrollment response packets are logged. The message indicates whether the enrollment response is valid or invalid for the trust point (by name). For valid responses, the response status is included. If the response status is successful, the enrollment status is also indicated. If enrollment fails, the reason for failure is given. **Example** `Aug 12 10:11:12:I: PKI: valid/invalid pki enrollment response received for trust point:<trustpoint_name> pki status: success/pending : enrollment status: failure/success. <Failure reason> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`

TABLE 34 PKI authentication packets

Type	Description
Authentication Request	Authentication requests sent to CA and RA for certificates are logged. **Example** `Aug 12 10:11:12:I: PKI: authenticate request sent for trust point :< trustpoint_name> to CA :< hostname> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`
Authentication Reply	Authentication replies from the CA are logged. The reply includes the CA and RA certificates and the trust point authentication status. If authentication fails, the reason for failure is given. **Example** `Aug 12 10:11:12:I: PKI: authentication reply for trustpoint:<trustpoint_name> pki authentication success/failure <failure reason> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`

TABLE 35 Certification Revocation List (CRL) and Online Control Status Protocol (OCSP) packets

Type	Description
Request for CRL or OCSP Packets	PKI requests for CRL or for OCSP packets are logged (the requests are based on revocation check configuration). **Example** `Aug 12 10:11:12 :I: PKI: crl/ocsp req sent for trust point :< trustpoint_name> to CA :< hostname> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`
CA or OCSP Responder Reply	PKI responses to the requests for CRL or for OCSP packets are logged. During this process, the peer certificate is validated and the certificate's revocation status is checked based on the CA or OCSP reply. Validation status is also logged. If validation fails, the reason for failure is given. **Example** `Aug 12 10:11:12 :I: PKI:crl/ocsp reply for trustpoint:<trustpoint_name>. Peer certificate <serial number> validation success/failed <failure reason> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`

The message in the following table apply to the manual import of certificates from external servers (in contrast to being imported from the device's flash memory). No validation checks are performed and no messages related to validation are logged.

TABLE 36 Manual import of certificate from external server, and peer certificate download

Type	Description
Request for Certificate	**Scenario** Manual import of certificate from external server

TABLE 36 Manual import of certificate from external server, and peer certificate download (continued)

Type	Description
	Requests for certificates from external server are logged. The request includes the server URL. **Example** `Aug 12 10:11:12 :I: PKI: certificate request for trustpoint/peer certificate <http url> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`
Certificate Request Reply	**Scenario** Manual import of certificate from external server Replies from external server to request for certificates are logged. The reply includes the certificate. **Example** `Aug 12 10:11:12 :I: PKI: certificate reply for trustpoint/peer certificate from <http url> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:<hex_dump>`

The message in the following table apply to scenarios in which the peer doesn't directly provide the certificate, but instead gives the URL of the certificate. PKI is used to download the certificate. Validation may happen next in a different step, but is not part of the request processing. For this reason, validation results are not logged.

TABLE 37 Manual import of certificate from external server, and peer provides certificate URL

Type	Description
Request for Certificate	**Scenario** Import of certificate from external server based on URL provided by peer Requests for certificates from external server are logged. The request includes the server URL. **Example** `Aug 12 10:11:12 :I: PKI: certificate request for trustpoint/peer certificate <http url> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:hex_dump`
Certificate Request Reply	**Scenario** Import of certificate from external server based on URL provided by peer Replies from external server to request for certificates are logged. The reply includes the certificate. **Example** `Aug 12 10:11:12 :I: PKI: certificate reply for trustpoint/peer certificate from <http url> Event:<event_no>, pkt:<pkt_no>, pkt_len:<pkt_len>, fragment:<frag_no>: frag_len:<frag_len>:hex_dump`

TABLE 38 Close connection requests

Type	Description
Close Connection Request	Requests to close the connection between the trust point and the host are logged. The name of the host is included in the log. **Example** `Aug 12 10:11:12 :I: PKI: close connection request for trustpoint to host :< host_name>`

Generating and deleting a PKI key pair

Public Key Infrastructure (PKI) uses a key pair (one public key and one private key) for data encryption and decryption. A key pair is generated on a device and used to enroll the device with a certification authority (CA).

A cryptographic key pair is not present on a device by default.

- Enter the **crypto key generate** command to generate a cryptographic key pair. The following example shows how to generate a 384-bit Elliptic Curve (EC) key pair with the label ec_key1.

```
device(config)# crypto key generate ec label ec_key1 key-size 384
```

- Enter the **crypto key zeroize** command to delete cryptographic keys. The following example shows how to delete all key pairs on a device.

```
device(config)# crypto key zeroize
```

The following example shows how to delete a specific Elliptic Curve (EC) key pair labeled ec_key1.

```
device(config)# crypto key zeroize ec label ec_key1
```

Configuring a PKI entity

A distinguished name (DN) uniquely identifies a Public Key Infrastructure (PKI) entity.

Perform the following task to create and configure a DN for a PKI entity.

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

   ```
   device# configure terminal
   ```

2. Create a PKI entity and enter configuration mode for the entity.

   ```
   device(config)# pki entity example_entity
   ```

3. Configure a common name for the entity.

   ```
   device(config-pki-entity-example_entity)# common-name mlx2
   ```

4. Configure the name of the organization that the entity belongs to.

   ```
   device(config-pki-entity-example_entity)# org-name example_org
   ```

5. Configure the name of the organization unit that the entity belongs to.

   ```
   device(config-pki-entity-example_entity)# org-unit-name marketing
   ```

6. Configure the state where the entity is located.

   ```
   device(config-pki-entity-example_entity)# state-name ca
   ```

7. Configure the two-character country code for the country where the entity is located.

   ```
   device(config-pki-entity-example_entity)# country-name us
   ```

8. Configure the email ID of the entity.

   ```
   device(config-pki-entity-example_entity)# email ROUTER1@example.com
   ```

9. (Optional) Configure the IPv4 address to be used in the certificate.

   ```
   device(config-pki-entity-example_entity)# ip 192.0.2.2
   ```

10. (Optional) Configure the fully qualified domain name (FQDN) of the entity.

    ```
    device(config-pki-entity-example_entity)# fqdn www.example.com
    ```

11. Configure the entity location.

    ```
    device(config-pki-entity-example_entity)# location sj
    ```

12. (Optional) Configure the subject alternative name of the entity.

The following example shows how to configure a subject alternate name by specifying an email address and URI.

```
device(config-pki-entity-example_entity)# subject-alt-name email:my@other.address,URI:http://
my.url.here/
```

When an IKEv2 peer uses an ID other than the DN, then the ID should be configured as a subject alternate name for the entity.

The following example shows how to configure a PKI entity named example_entity

```
device# configure terminal
device(config)# pki entity example_entity
device(config-pki-entity-example_entity)# common-name mlx2
device(config-pki-entity-example_entity)# org-name example_org
device(config-pki-entity-example_entity)# org-unit-name marketing
device(config-pki-entity-example_entity)# state-name ca
device(config-pki-entity-example_entity)# country-name us
device(config-pki-entity-example_entity)# email ROUTER1@example.com
device(config-pki-entity-example_entity)# ip 192.0.2.2
device(config-pki-entity-example_entity)# fqdn www.example.com
device(config-pki-entity-example_entity)# location sj
device(config-pki-entity-example_entity)# subject-alt-name email:my@other.address,URI:http://my.url.here/
```

Configuring a PKI trustpoint

Trustpoint configuration set parameters for communication with a Public Key Infrastructure (PKI) certification authority (CA).

Before configuring a trustpoint:

- The Elliptic Curve (EC) key pair to use for enrollment with the CA should be generated.
- The PKI entity that is to be enrolled with the CA must be created and configured.

PKI certificates are obtained from a CA. Perform the following task to create a trustpoint on a device and to configure the trustpoint parameters for communication with the CA. Up to five trustpoints can be configured on a device.

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

```
device# configure terminal
```

2. Create a trustpoint.

```
device(config)# pki trustpoint tp1
```

3. (Optional) Configure the device to automatically send enrollment messages to the CA to obtain the CA certificate and local certificates.

```
device(config-pki-trustpoint-tp1)# auto-enroll
```

4. Configure an enrollment profile for the trustpoint.

```
device(config-pki-trustpoint-tp1)# enrollment profile profileA
```

5. (Optional) When the CA does not send the certificate immediately, the device goes into polling mode and retries to obtain the certificate. The following example shows how to set the number of retries to 10 and the polling interval to 2 minutes.

```
device(config-pki-trustpoint-tp1)# enrollment retry-count 10 retry-period 2
```

6. Specify the PKI entity to enroll with the CA.

```
device(config-pki-trustpoint-tp1)# pki-entity example_entity
```

7. (Optional) Configure the revocation check method to be used during certificate validation.

```
device(config-pki-trustpoint-tp1)# revocation-check crl
```

By default, revocation check is disabled.

8. (Optional) When revocation check is enabled by specifying the **crl** option, configure the URL from which the updated certificate revocation list (CRL) can be obtained.

```
device(config-pki-trustpoint-tp1)# crl-query http://WIN-HJ98AK136A0.englab.example.com/CertEnroll/
englab-WIN-HJ98AK136A0-CA-10.crl
```

9. (Optional) When revocation check is enabled by specifying the **crl** option, configure the CRL update period. The following example shows how to set the update period to 24 hours.

```
device(config-pki-trustpoint-tp1)# crl-update-time 24
```

10. (Optional) When revocation check is enabled by specifying the **ocsp** option, configure the Online Certificate Status Protocol (OCSP) server URL.

```
device(config-pki-trustpoint-tp1)# ocsp-url http://abc.com/ocsp
```

11. Configure the EC key pair to use for enrollment with the CA.

```
device(config-pki-trustpoint-tp1)# eckeypair key-label label-1
```

12. Configure the fingerprint of the CA.

```
device(config-pki-trustpoint-tp1)# fingerprint 26:fc:77:6c:b9:02:91:c9:a2:ab:60:28:c9:b9:c1:23:45:0a:
8b:06
```

When a certificate is received from the CA, it is accepted when it matches the fingerprint configured for the trustpoint on the device.

Authenticating a PKI trustpoint

A trustpoint is authenticated by obtaining the certification authority (CA) certificate. Trustpoint authentication must be completed prior to enrolling a client certificate manually.

The commands listed in the following task are not effective when old certificates exist on the device. Any old certificates on the device for the trustpoint should be deleted by using the **no pki authenticate** command prior to completing the following task.

Perform the following task to authenticate a trustpoint on a device.

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

```
device# configure terminal
```

2. Authenticate the trustpoint to the device.

```
device(config)# pki authenticate tp1
```

This command is saved in the device configuration. The trustpoint is authenticated by verifying the fingerprint of self-signed certificate obtained from the CA (trustpoint); the self-signed certificate contains the public key of the CA. You should manually contact the CA to obtain the CA fingerprint.

Creating a PKI enrollment profile

Public Key Infrastructure (PKI) enrollment profile configuration sets parameters for certificate authentication and enrollment on a device.

Perform the following task to create a PKI enrollment profile.

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

   ```
   device# configure terminal
   ```

2. Create an enrollment profile and enter configuration mode for the profile.

   ```
   device(config)# pki profile-enrollment profileA
   ```

3. Configure the URL of the authentication server.

   ```
   device(config-pki-enrollment-profile-profileA)# authentication-url http://192.0.2.1/CertSrv/mscep/
   mscep.dll
   ```

   When Simple Certificate Enrollment Protocol (SCEP) is used for enrollment, the URL format must be `http://CA_name`, where `CA_name` is the host Domain Name System (DNS) name or the IP address of the CA. When the domain-name format is specified, the domain name must be mapped to the IP address of the CA by using the **ip host** command in global configuration mode.

   When Trivial File Transfer Protocol (TFTP) is used for enrollment, the URL format must be `tftp://certserver/file_specification`. When the URL does not include a file specification, the fully qualified domain name (FQDN) of the router is used.

4. (Optional) Configure the authentication command to use during authentication.

   ```
   device(config-pki-enrollment-profile-profileA)# authentication-command GET/certs/cacert.der
   ```

5. Configure the URL of the enrollment server.

   ```
   device(config-pki-enrollment-profile-profileA)# enrollment-url http://entrust:81/cda-cgi/
   clientcgi.exe
   ```

   When Simple Certificate Enrollment Protocol (SCEP) is used for enrollment, the URL format must be `http://CA_name`, where `CA_name` is the host Domain Name System (DNS) name or the IP address of the CA. When the domain-name format is specified, the domain name must be mapped to the IP address of the CA by using the **ip host** command in global configuration mode.

   When Trivial File Transfer Protocol (TFTP) is used for enrollment, the URL format must be `tftp://certserver/file_specification`. When the URL does not include a file specification, the fully qualified domain name (FQDN) of the router is used.

6. (Optional) When the **auto-enroll** option is configured for a trustpoint, the SCEP-challenge password is used in the issue and revocation of client certificates. Configure an SCEP-challenge password.

   ```
   device(config-pki-enrollment-profile-profileA)# password F1D43B07E91C82DE
   ```

The following example shows how to configure an enrollment profile named profileA.

```
device# configure terminal
device(config)# pki profile-enrollment profileA
device(config-pki-enrollment-profile-profileA)# authentication-url http://192.0.2.1/CertSrv/mscep/mscep.dll
device(config-pki-enrollment-profile-profileA)# authentication-command GET/certs/cacert.der
device(config-pki-enrollment-profile-profileA)# enrollment-url http://entrust:81/cda-cgi/clientcgi.exe
device(config-pki-enrollment-profile-profileA)# password F1D43B07E91C82DE
device(config-pki-enrollment-profile-profileA)# end
```

Obtaining a PKI client certificate

A Public Key Infrastructure (PKI) client certificate is obtained by requesting a certification authority to issue certificates for each cryptographic key pair that exists on a device.

Before generating a client certificate request, a cryptographic key pair must be generated on the device by using the **crypto key generate** command.

Perform the following task to generate a PKI client certificate request.

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

   ```
   device# configure terminal
   ```

2. Generate a client certificate request.

   ```
   device(config)# pki enroll tp1
   ```

   The certificates received from the CA are saved to the device. The command is not saved to the device configuration.

3. Verify that the certificates are saved to the device.

   ```
   device# show pki certificates local

   ----------------PKI TRUSTPOINT CERTIFICATE ENTRY----------------
   CA: tp1
   Certificate:
       Data:
           Version: 3 (0x00000002)
           Serial Number:
               12:bd:f7:01:24:27:ec:bf:4e:b5:c5:89:15:71:27:9f
       Signature Algorithm: ecdsa-with-SHA384
           Issuer: DC=com, DC=brocade, DC=englab, CN=englab-WIN-HJ98AK136A0-CA-7
           Validity
               Not Before: Apr 1 10:31:03 2015 GMT
               Not After : Apr 1 10:41:01 2020 GMT
           Subject: DC=com, DC=brocade, DC=englab, CN=englab-WIN-HJ98AK136A0-CA-7
   ```

4. Validate the trustpoint certificate on the device.

   ```
   device(config)# pki cert-validate tp1
   ```

   The **pki cert-validate** command is used to verify that the local certificate was requested from and granted by the CA, and that the certificate is currently valid.

The following example shows how to request client certificates for a trustpoint named tp1, view the certificates obtained from the CA that are saved on the device, and to validate the local certificate.

```
device# configure terminal
device(config)# pki enroll tp1
device# show pki certificates trustpoint tp1
device(config)# pki cert-validate tp1
```

Exporting PKI keys, certificates, and CRLs manually

Public Key Infrastructure (PKI) cryptographic keys, certificates, and certificate revocation lists (CRLs) can be exported to flash memory for import after a device is rebooted.

Before executing any of the following tasks, the relevant cryptographic keys, certificate, or CRL must exist on the device. A cryptographic key pair is generated by using the **crypto key generate** command. Certificates and CRLs are imported by using the **pki import** command.

- The following example shows how to export a key pair to flash memory.

    ```
    device(config)# pki export key keylabel1 password key1
    ```

- The following example shows how to export a client certificate from a trustpoint named tp1 to a file in flash memory. The file is saved in privacy-enhanced mail (PEM) format.

    ```
    device(config)#pki export tp1 pem url router1.crt
    ```

- The following example shows how export a CRL file from a trustpoint named tp1 to a file in flash memory.

    ```
    device(config)# pki export crl tp1 url ca10.crl
    ```

Authenticating a PKI trustpoint manually

A Public Key Infrastructure (PKI) trustpoint that is configured on a device can be manually authenticated when the certification authority (CA) server is offline.

The trustpoint must be configured on the device before it can be authenticated.

Prior to completing the following task, enrollment terminal must be enabled for the trustpoint by issuing the **enrollment terminal** command .

Perform the following task to manually authenticate a trustpoint that is configured on an MLX Series device.

1. Copy the certificate file from the CA server.
2. Take the copy of the CA certificate file to the MLX Series device and load it onto the device either by using a file copy protocol such as Trivial File Transfer Protocol (TFTP) if that is supported or by pasting onto the device terminal.

3. Authenticate the trustpoint.

   The following example shows how to authenticate the trustpoint by entering a chain of certificates.

```
device(config)# pki authenticate example_tp

Enter certificate in base64 encoded format.
End with a blank line.

-----BEGIN CERTIFICATE-----
MIIDEzCCArqgAwIBAgIJAMYEw0fMnyOpMAoGCCqGSM49BAMCMIGMMQswCQYDVQQG
EwJVUzELMAkGA1UECBMCTUQxFDASBgNVBAcTC0NhdG9uc3ZpbGxlMQwwCgYDVQQK
EwNHU1MxKzApBgkqhkiG9w0BCQEWHHJvb3RjYS5lY2RzYUBnb3NzYW1lcnNlYy5j
b20xHzAdBgNVBAMTFkdvc3NhbWVyIEVDRFNBIFJvb3QgQ0EwHhcNMTUwNjE5MTY0
NjI4WhcNMjUwNjE2MTY0NjI4WjCBjDELMAkGA1UEBhMCVVMxCzAJBgNVBAgTAk1E
MRQwEgYDVQQHEwtDYXRvbnN2aWxsZTEMMAoGA1UEChMDR1NTMSswKQYJKoZIhvcN
AQkBFhxyb290Y2EtZWNkc2FAZ29zc2FtZXJzZWMuY29tMR8wHQYDVQQDExZHb3Nz
YW1lciBFQ0RTQSBSb290IENBMFkwEwYHKoZIzj0CAQYIKoZIzj0DAQcDQgAEIKpj
7bRTde6IjMB/5I0YfQA5tbugpeMFOSQv7FwiGS9wH+Ie49s5N6AcRZnqDxjabEaD
qdroVG+tv0PEuaWd7KOCAQEwgf4wHQYDVR0OBBYEFOyftp0utn4K/0BY0xxT5yWP
1MxYMIHBBgNVHSMEgbkwgbaAFOyftp0utn4K/0BY0xxT5yWP1MxYoYGSpIGPMIGM
MQswCQYDVQQGEwJVUzELMAkGA1UECBMCTUQxFDASBgNVBAcTC0NhdG9uc3ZpbGxl
MQwwCgYDVQQKEwNHU1MxKzApBgkqhkiG9w0BCQEWHHJvb3RjYS5lY2RzYUBnb3Nz
YW1lcnNlYy5jb20xHzAdBgNVBAMTFkdvc3NhbWVyIEVDRFNBIFJvb3QgQ0GCCQDG
BMNHzJ8jqTAMBgNVHRMEBTADAQH/MAsGA1UdDwQEAwIBBjAKBggqhkjOPQQDAgNH
ADBEAiAUXYCIZRZaHztfQ/YHKSaDH49IsOqNxFTkm9ojz7QAEgIgGMua7Ww53dXJ
Tc8S1Bkbkh5GN/zuybCMHQ1ioqSME1w=
-----END CERTIFICATE-----

The self signed certificate has fingerprint:c5:15:e8:27:7e:5d:1b:e0:66:df:ca:3e:18:40:ff:
9e:f1:9d:a3:ec
Do you accept this certificate as root CA certificate? (enter 'y' or 'n'): y
you accepted the certificate.

Do you wish to enter more ca certificates? (enter 'y' or 'n'): y
Enter certificate in base64 encoded format.
End with a blank line.

-----BEGIN CERTIFICATE-----
MIIDfTCCAySgAwIBAgIBAjAKBggqhkjOPQQDAjCBjDELMAkGA1UEBhMCVVMxCzAJ
BgNVBAgTAk1EMRQwEgYDVQQHEwtDYXRvbnN2aWxsZTEMMAoGA1UEChMDR1NTMSsw
KQYJKoZIhvcNAQkBFhxyb290Y2EtZWNkc2FAZ29zc2FtZXJzZWMuY29tMR8wHQYD
VQQDExZHb3NzYW1lciBFQ0RTQSBSb290IENBMB4XDTE1MDYxOTE2NDYyOVoXDTI1
MDYxNjE2NDYyOVowgZYxCzAJBgNVBAYTAlVTMQswCQYDVQQIEwJNRDEUMBIGA1UE
BxMLQ2F0b25zdmlsbGUxDDAKBgNVBAoTA0dTUzEkMCIGA1UEAxMbR29zc2FtZXIg
RUNEU0EtRUNEU0EgU3ViIENBMTAwLgYJKoZIhvcNAQkBFiFzdWJjYS1lY2RzYS1l
Y2RzYUBnb3NzYW1lcnNlYy5jb20wWTATBgcqhkjOPQIBBggqhkjOPQMBBwNCAAQR
Bq93D2C7q+w/vBRgn90h363E7SRhheFDpB6YIDDSp9S5dH501WTelqsAUzagrip1
sl8HDeqMmVuZUzSNZV0vo4IBaTCCAWUwHQYDVR0OBBYEFD1qxubtQ0QYnlmY/MLS
4a2qNuUfMIHBBgNVHSMEgbkwgbaAFOyftp0utn4K/0BY0xxT5yWP1MxYoYGSpIGP
MIGMMQswCQYDVQQGEwJVUzELMAkGA1UECBMCTUQxFDASBgNVBAcTC0NhdG9uc3Zp
bGxlMQwwCgYDVQQKEwNHU1MxKzApBgkqhkiG9w0BCQEWHHJvb3RjYS5lY2RzYUBn
b3NzYW1lcnNlYy5jb20xHzAdBgNVBAMTFkdvc3NhbWVyIEVDRFNBIFJvb3QgQ0GC
CQDGBMNHzJ8jqTAMBgNVHRMEBTADAQH/MAsGA1UdDwQEAwIBBjAyBgNVHR8EKzAp
MCegJaAjhiFodHRwOi8vMTAuMC4wLjM0L3Jvb3RjYS5lY2RzYS5jcmwwMQYIKwYB
BQUHAQEEJTAjMCEGCCsGAQUFBzABhhVodHRwOi8vMTAuMC4wLjM0Ojc3NzgwCgYI
KoZIzj0EAwIDRwAwRAIgemQ2bHWJT+8EH/ol9qDY8QGmKSb/EbKSGUnVVVrkqBIC
IEZLG8jwpnAVaT7S+Sdc2nGf600CCdcVKQp2q8QOjpmE
-----END CERTIFICATE-----

Do you wish to enter more ca certificates? (enter 'y' or 'n'): n
PKI: Trustpoint example_tp, authentication: Success

device(config)#
```

When authentication is successful, the CA certificates are saved as PEM-encoded files in flash memory.

Generating a CSR for manual submission to a CA

When a certification authority (CA) server is offline, a Public Key Infrastructure (PKI) certificate signing request (CSR) can be generated and manually transferred to the CA server to obtain a client certificate.

Before completing the following task, the trustpoint (for which the CSR is to be generated) must be configured and authenticated on the MLX Series device.

When the CA server is offline, perform the following task to generate a CSR and manually submit it to the CA.

1. From global configuration mode, enter trustpoint configuration mode.
   The following example shows how to enter configuration mode for a trustpoint named example_tp.

   ```
   device(config)# pki trustpoint example_tp
   device(config-pki-trustpoint-example_tp)#
   ```

2. Enable enrollment terminal for the trustpoint.

   ```
   device(config-pki-trustpoint-example_tp)# enrollment terminal
   device(config-pki-trustpoint-example_tp)# exit
   ```

   Enabling enrollment terminal causes the **pki enroll** command to display the CSR on the terminal for copy and manual transfer to an offline CA. When enrollment terminal is disabled issuing the **pki enroll** command causes a CSR to be sent to the CA server over HTTP and requires the CA server to be online.

3. Generate a CSR for the trustpoint.
   The following example shows how to generate a CSR for the trustpoint named example_tp.

   ```
   device(config)# pki enroll example_tp

   PKI: Certificate Request for trustpoint example_tp saved in file pki_certreq_example_tp and
   displayed below:

   -----BEGIN CERTIFICATE REQUEST-----
   MIIBkzCCARsCAQAwWjELMAkGA1UEBhMCSU4xDDAKBgNVBAgTA0tBUjEMMAoGA1UE
   BxMDQkxSMQ0wCwYDVQQKEwRCUkNEMQswCQYDVQQLEwJOSTETMBEGA1UEAxMKY2hh
   aXRhbmlhMTB2MBAGByqGSM49AgEGBSuBBAAiA2IABLOY6MslGw/juGnz5hG0H7tu
   fhSwM/k1QoX13Vg7RRguD2HdnRvjfzUcY/g4Fsq8nE9xAvuDwvxgA3TbCrZJys36
   Rv0jNJWJ7TUfelUq4hbmzW9vCCsmbPZeH6caRWFMDqBCMB8GCSqGSIb3DQEJBzES
   ExA5MjE4QjdGQ0ZFNUU4REJBMB8GCSqGSIb3DQEJDjESMBAwDgYDVR0PAQH/BAQD
   AgXgMAkGByqGSM49BAEDZwAwZAIwOSUGx0Z+IOixgVJ8QMARr4r5cpTr+xTLXaNR
   zNxEfGHq/tZzKm/zYBsNhBaCwJ7JAjA+lz5Ub12bqJX688yBaYToo9lTeho0BilC
   7bCvKKVWckgOWQKys85AJtf4PS1nA8U=
   -----END CERTIFICATE REQUEST-----

   device(config)#
   ```

   The CSR is displayed on the device terminal in PKCS #10 format. A copy of the CSR is also saved to a PEM-encoded file named pki_certreq_example_tp in flash memory.

4.   **NOTE**
       When the operating environment permits the use of a file copy protocol such as Secure Copy Protocol (SCP) or Trivial File Transfer Protocol (TFTP) that are supported by the MLX Series device, the CSR file in flash memory may be copied. Alternatively, the CSR displayed on the device terminal can be copied.

   Copy the CSR file to the CA server.
   The following example shows a CSR file named cert_req_ET that is copied into a Notepad file.
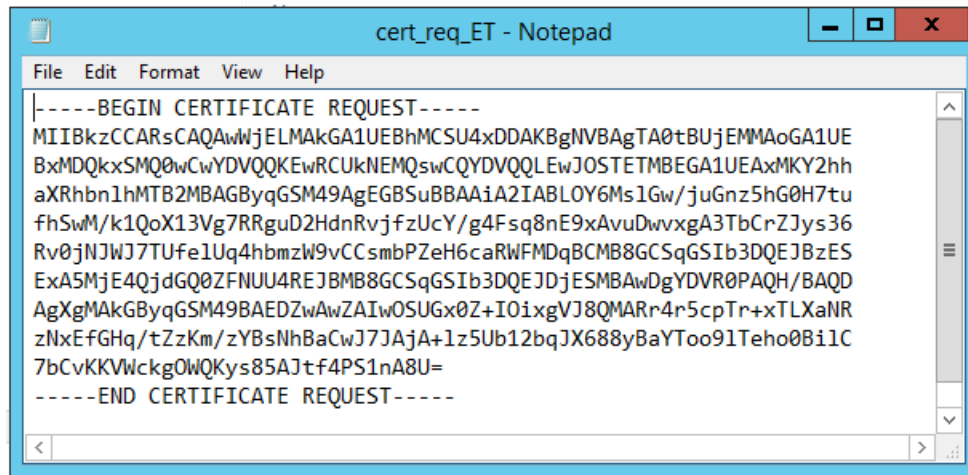
5.   Submit the CSR file by way of the CA console as shown in the following figure.
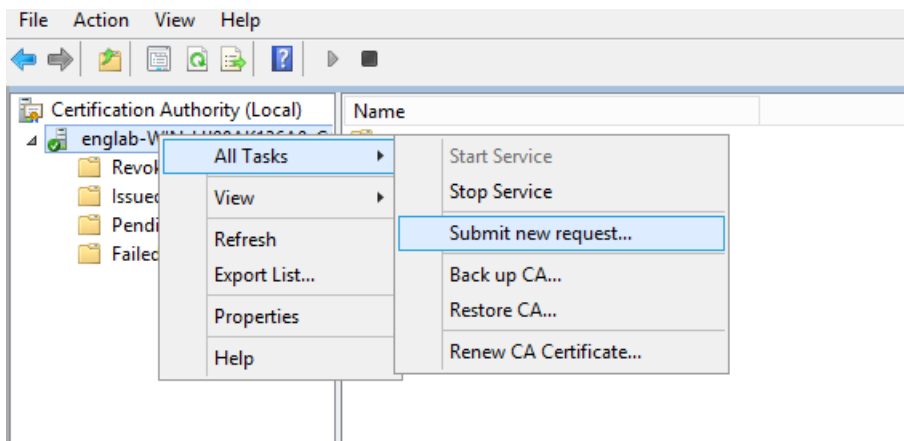
After the client certificate is generated on the CA, it can be manually imported into the MLX Series device by using the **pki import** command.

Importing a client certificate manually by using TFTP

When a certification authority (CA) server is offline, certificates can be copied from the CA server and manually imported into the trustpoint on a device by using Trivial File Transfer Protocol (TFTP).

Typically, when the CA server is offline, a certificate signing request (CSR) is generated on the MLX Series device, transported to the CA, and manually submitted to generate the client certificate. Therefore, the task of generating a CSR for manual submission to a CA should be completed before performing the following task.

NOTE
The trustpoint must be authenticated prior to manually importing a client certificate by using
TFTP.

When the CA server is offline and the operating environment permits the use of Trivial File Transfer Protocol (TFTP), perform the following task to manually import a client certificate into a trustpoint on an MLX Series device by using TFTP. Only one client certificate can be manually imported per trustpoint.

1. Copy the client certificate from the CA server in privacy-enhanced mail (PEM) format.
2. Take the copied client certificate file to the MLX Series device and copy it into flash memory on the device.
3. Manually import the client certificate file into the trustpoint on the MLX Series device by using TFTP.
   The following example shows how to import a PEM-encoded client certificate file named client_cert from flash memory into a trustpoint named example_tp.

   ```
   device# pki import example_tp pem url flash: client_cert
   ```

4. Verify that the client certificate is now imported.

   ```
   device# show pki certificates trustpoint example_tp
   ```

Importing a client certificate manually by way of the device terminal

When a certification authority (CA) server is offline, certificates can be copied from the CA server and manually imported into the trustpoint on a device by pasting onto the device terminal.

Typically, when the CA server is offline, a certificate signing request (CSR) is generated on the MLX Series device, transported to the CA, and manually submitted to generate the client certificate. Therefore, the task of generating a CSR for manual submission to a CA should be completed before performing the following task.

NOTE
The trustpoint must be authenticated prior to manually importing the client certificate by way of the device
terminal.

When the CA server is offline and the operating environment does not permit the use of Trivial File Transfer Protocol (TFTP), perform the following task to manually import a client certificate into a trustpoint on an MLX Series device by pasting it onto the device terminal. Only one client certificate can be manually imported per trustpoint.

1. Copy the client certificate from the CA server in privacy-enhanced mail (PEM) format.
2. Take the copied client certificate file to the MLX Series device.
3. From global configuration mode, enter configuration mode for the trustpoint associated with the client certificate that is to be imported.
   The following example shows how to enter configuration mode for a trustpoint named example_tp.

   ```
   device(config)# pki trustpoint example_tp
   device(config-pki-trustpoint-example_tp)#
   ```

4. Enable enrollment terminal for the trustpoint.

   ```
   device(config-pki-trustpoint-example_tp)# enrollment terminal
   device(config-pki-trustpoint-example_tp)# exit
   ```

5. Manually import the client certificate file into the trustpoint on the device by copying the client certificate file transported from the CA and pasting it onto the device terminal.

   The following example shows how to import a client certificate file named client_cert into a trustpoint named example_tp.

```
device# pki import mytp pem terminal client_cert

Enter certificate in base64 encoded format.
End with a blank line.

-----BEGIN CERTIFICATE-----
MIIDgDCCAyigAwIBAgITGAAAACtOFPtBtvv8ywAAAAAAKzAJBgcqhkjOPQQBMCcx
JTAjBgNVBAMTHGVuZ2xhYi1XSU4tTjZDM1IwTFVEQUotQ0EtNDUwHhcNMTUxMTE0
MDgyODU4WhcNMTYxMDAxMTQ1MzIzWjBZMQswCQYDVQQGEwJJTjEMMAoGA1UECBMD
S0FSMQwwCgYDVQQHEwNCTFIxDTALBgNVBAoTBEJSQ0QxCzAJBgNVBAsTAk5JMRIw
EAYDVQQDEwljaGFpcGFueWEwdjAQBgcqhkjOPQIBBgUrgQQAIgNiAAQgDXauakg4
CIimoWBMW0xoak+aopUtOFNrywqlaTuDfGP8Mc7As2C85kBIYawjAbKcZ8dDj0zc
Widy6fGB/jQqzMCcv9FjthK7/JOSjxbNhHX4hW/mMMlMtEcFla/HGK6jggHjMIIB
3zAOBgNVHQ8BAf8EBAMCBeAwHQYDVR0OBBYEFCJnYr45VOWbfJs9wIpeY5kaVuHX
MB8GA1UdIwQYMBaAFPOwj03oTL9+tBSm3lxcvPuDAa3TMGYGA1UdHwRfMF0wW6BZ
oFeGVWh0dHA6Ly9XSU4tTjZDM1IwTFVEQUouZW5nbGFiLmJyb2NhZGUuY29tL0Nl
cnRFbnJvbGwvZW5nbGFiLVdJTi1ONkMzUjBMVURBSi1DQS00NS5jcmwwgdUGCCsG
AQUFBwEBBIHIMIHFMIGGBggrBgEFBQcwAoZ6ZmlsZTovLy8vV0lOLU42QzNSMExV
REFKLmVuZ2xhYi5icm9jYWRlLmNvbS9DZXJ0RW5yb2xsL1dJTi1ONkMzUjBMVURB
Si5lbmdsYWIuYnJvY2FkZS5jb21fZW5nbGFiLVdJTi1ONkMzUjBMVURBSi1DQS00
NS5jcnQwOgYIKwYBBQUHMAGGLmh0dHA6Ly9XSU4tTjZDM1IwTFVEQUouZW5nbGFi
LmJyb2NhZGUuY29tL29jc3AwPwYJKwYBBAGCNxQCBDIeMABJAFAAUwBFAEMASQBu
AHQAZQByAG0AZQBkAGkAYQB0AGUATwBmAGYAbABpaG4AZTAMBgNVHRMBAf8EAjAA
MAkGByqGSM49BAEDRwAwRAIgAz+NzKDsi2+apIZ3A+PfYT2Z1p4BiXjd3IoKoeGE
xkkCIDnEJCxJ6XOc3CF+LTYCJs7MvpIBBIsEzC7fveKq53fo
-----END CERTIFICATE-----

device#
```

6. Verify that the client certificate is now imported.

```
device# show pki certificates trustpoint example_tp
```

Importing a CRL manually by using TFTP

When a CRL distribution point is offline, a Public Key Infrastructure (PKI) certificate revocation list (CRL) can be manually imported into a trustpoint on a device by using Trivial File Transfer Protocol (TFTP).

Before completing the following task, the trustpoint (to which the CRL is to be imported) must be configured on the MLX Series device.

When the CRL distribution point is offline and the operating environment supports the use of TFTP, perform the following task to manually import a CRL into a trustpoint on an MLX Series device.

1. From the CRL distribution point server, make a copy the CRL in either DER or PEM format.
2. Take the CRL file to the MLX Series device and copy it into flash memory.

3. Import the CRL file into the trustpoint.
   The following example shows how to import a DER-encoded CRL file (CA-10.crl) in flash memory into a trustpoint named example_tp.

   ```
   device# pki import example_tp crl der url flash CA-10.crl

   Crl imported successfully.

   device#
   ```

   The following example shows how to import a PEM-encoded CRL file (CA-10_out.pem) in flash memory into a trustpoint named example_tp.

   ```
   device# pki import example_tp crl pem url flash CA-10_out.pem

   Crl imported successfully.

   device#
   ```

4. Verify that the CRL is now imported.

   ```
   device# show pki crls example_tp
   ```

Importing a CRL manually by way of device terminal

When a CRL distribution point is offline, a Public Key Infrastructure (PKI) certificate revocation list (CRL) can be manually imported into a trustpoint on a device by pasting the CRL file onto the device terminal.

Before completing the following task, the trustpoint (to which the CRL is to be imported) must be configured on the MLX Series device.

When the CRL distribution point is offline and the operating environment does not support the use of TFTP, perform the following task to manually import a CRL into a trustpoint on an MLX Series device.

1. Take a copy the CRL from the CRL distribution point server in PEM format.
2. Take the copy of the CRL file to the MLX Series device.

3. Manually import the CRL file into the trustpoint by pasting the copied file onto the device terminal.
    The following example shows how to import a CRL file into a trustpoint named example_tp and specify "1" as the label for the CRL.

    ```
    device# pki import example_tp crl pem terminal 1

    Enter CRL in base64 encoded format.
    End with a blank line.

    -----BEGIN X509 CRL-----
    MIICGTCCAcECAQEwCQYHKoZIzj0EATAnMSUwIwYDVQQDExxlbmdsYWItV0lOLUhK
    OThBSzEzNkEwLUNBLTEwFw0xNTExMTYxMjMwNDhaFw0xNjExMTcwMDUwNDhaMCYw
    JAITGgAAAAe02n9njdqiRgAAAAAABxcNMTUxMTE2MTI0MDAwWqCCAUAwggE8MB8G
    A1UdIwQYMBaAFJ3/I+5GucF3N48ryw116Pt2l5wdMBAGCSsGAQQBgjcVAQQDAgEA
    MAoGA1UdFAQDAgEDMBwGCSsGAQQBgjcVBAQPFw0xNjExMTYxMjQwNDhaMIHcBgkr
    BgEEAYI3FQ4Egc4wgcswgciggcWggcKGgb9sZGFwOi8vL0NOPWVuZ2xhYi1XSU4t
    SEo5OEFLMTM2QTAtQ0EtMTAsQ049V0lOLUhKOThBSzEzNkEwLENOPUNEUCxDTj1Q
    dWJsaWMlMjBLZXklMjBTZXJ2aWNlcyxDTj1TZXJ2aWNlcyxEQz1VbmF2YWlsYWJs
    ZUNvbmZpZ0ROP2NlcnRpZmljYXRlUmV2b2NhdGlvbkxpc3Q/YmFzZT9vYmplY3RD
    bGFzcz1jUkxEaXN0cmlidXRpb25Qb2ludDAJBgcqhkjOPQQBA0cAMEQCIHK4cITc
    lHRR6oqhnDxNc2nspJhkdfFK3qTP3ahIk6yaAiAPxVuIyfXIEaxjRvIqiix4T3VL
    07GN+2Xe/cerwkuN7Q==
    -----END X509 CRL-----

    CRL imported successfully.

    device#
    ```

4. Verify that the CRL is now imported.

    ```
    device# show pki crls example_tp
    ```

Enabling hash and URL encoding of X.509 certificates

Hash and URL encoding of X.509 certificates avoids fragmentation of User Datagram Protocol (UDP) datagrams in Internet Key Exchange version 2 (IKEv2) negotiations.

To use hash and URL encoding, you need to:

*   Configure the URL where the hash version of the X.509 certificate is to be stored by using the **local-certificate url** command.
*   Prepare and upload the hash version of the certificate to the configured URL.
*   Enable hash and URL encoding for IKEv2 negotiations by using the **ikev2 http-url-cert** command.

Perform the following task to enable hash and URL encoding of X.509 certificates.

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

    ```
    device# configure terminal
    ```

2. Enter PKI trustpoint configuration mode.

    ```
    device(config)# pki trustpoint tp1
    device(config-pki-tp1)#
    ```

3. Specify the URL where the hash version of the local certificate will be placed.

    ```
    device(config-pki-tp1)# local-certificate url http://example.com/pki_local_cert
    ```

4. Repeat the preceding steps for all trustpoints that are configured on the device.

5. (Optional) When the X.509 certificate file is in PEM-encoded format, use OpenSSL to convert it to DER-encoded format. The following example shows how to convert a certificate file (peer1.cert) from PEM-encoded format to DER-encoded format.

```
openssl x509 -in peer1_cert.pem -inform PEM -out peer1_cert.der -outform DER
```

6. Obtain a SHA1 hash of the local X.509 certificate file data. The following example shows how to obtain a hash for a certificate file named peer1_cert.der by using OpenSSL.

```
openssl dgst -sha1 peer1_cert.der
```

A hash value (394ceefaef48af8394d9a0e63d74cc56a4117a23) is returned.

7. Rename the peer1_cert.der file to the hash value obtained in the preceding step.

```
mv ca.der 394ceefaef48af8394d9a0e63d74cc56a4117a23
```

8. Upload the renamed file (394ceefaef48af8394d9a0e63d74cc56a4117a23) to the URL previously configured by using the **local-certificate url** command (http://example.com/pki_local_cert) .

   The local-certificate URL and the SHA1 hash of the DER-encoded certificate are concatenated for use in IKEv2 negotiations; for example http://example.com/pki_local_cert/394ceefaef48af8394d9a0e63d74cc56a4117a23.

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

```
device# configure terminal
```

10. Enable hash and URL encoding for X.509 certificates.

```
device(config)# ikev2 http-url-cert
```

11. Confirm that hash and URL encoding is enabled.

```
device(config)# show ikev2

IKEv2 Global data:
Retry Count           : 5            Max Exchange Time        : 60
Retransmit Interval   : 5            Cookie Challenge Number  : 0
Max Sa In Nego per LP : 256          Max Sa per LP            : 256
Allow Duplicate       : False        Http Cert Enable         : True
Total Peers           : 1            Total Ipsec Intf         : 1
Total Ike Sa          : 1            Total Ipsec Sa           : 2
NAT-T enabled         : True             NAT-T keep-alive time : 5 sec
Sync In Progress      : 0            Time to syn peer to LP   : 0
Pending Job           : 0            Sw pending time          : 0
```

Clearing PKI CRLs

You can delete the Public Key Infrastructure (PKI) certificate revocation list (CRL) database from a certificate authority (CA).

- Enter the **clear pki crl** command to delete the CRL database from a specific trustpoint. The following example shows how to clear the CRL database from trustpoint1.

```
device(config)# clear pki crl trustpoint1
```

Clearing PKI counters

You can clear the Public Key Infrastructure (PKI) counters on a device.

- Enter the **clear pki counters** command to clear the PKI counters.

```
device(config)# clear pki counters
```

Displaying PKI configuration information

Various commands can be used to display Public Key Infrastructure (PKI) configuration information.

PKI must be configured before displaying the following information.

- Enter the **show pki certificates** command to display certificate information associated with a trustpoint or the local device.

  The following example displays trustpoint certificate information.

```
device# show pki certificates trustpoint

----------------PKI TRUSTPOINT CERTIFICATE ENTRY----------------
CA: R3-Trustpoint
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 1 (0x1)
    Signature Algorithm: ecdsa-with-SHA1
        Issuer: C=US, ST=KS, L=KS, O=ST, OU=Eng, CN=Eng/emailAddress=st@ca.com
        Validity
            Not Before: Apr 29 02:15:15 2016 GMT
            Not After : Apr 29 02:15:15 2018 GMT
        Subject: C=US, ST=CA, L=SFO, O=SPT, OU=Eng, CN=Eng/emailAddress=spt@ca.com

----------------PKI TRUSTPOINT CERTIFICATE ENTRY----------------
CA: R3-Trustpoint
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            f2:c4:16:d6:bc:9a:76:54
    Signature Algorithm: ecdsa-with-SHA1
        Issuer: C=US, ST=KS, L=KS, O=ST, OU=Eng, CN=Eng/emailAddress=st@ca.com
        Validity
            Not Before: Apr 29 02:12:42 2016 GMT
            Not After : Apr 29 02:12:42 2021 GMT
        Subject: C=US, ST=KS, L=KS, O=ST, OU=Eng, CN=Eng/emailAddress=st@ca.com
```

  The following example displays certificate information for the local device.

```
device# show pki certificates local

----------------PKI LOCAL CERTIFICATE ENTRY----------------
CA: R3-Trustpoint
Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number: 1 (0x1)
    Signature Algorithm: ecdsa-with-SHA1
        Issuer: C=US, ST=CA, L=SFO, O=SPT, OU=Eng, CN=Eng/emailAddress=spt@ca.com
        Validity
            Not Before: Apr 29 02:16:43 2016 GMT
            Not After : Jan 19 02:16:43 2018 GMT
        Subject: C=US, ST=CA, L=SD, O=BR, OU=NI, CN=SQA/emailAddress=ni@br.com
```

The following example displays detailed certificate information for the local device.

```
device# show pki certificates local detail

----------------PKI LOCAL CERTIFICATE ENTRY-----------------
CA: auto
Certificate:
    Data:
        Version: 3 (0x00000002)
        Serial Number:
            1d:00:00:01:6e:fb:90:2f:19:d6:22:8f:e1:00:00:00:00:01:6e
    Signature Algorithm: ecdsa-with-SHA384
        Issuer: DC=com, DC=example, DC=englab, CN=englab-WIN-HJ98AK136A0-CA-7
        Validity
            Not Before: Jun 3 13:25:54 2015 GMT
            Not After : Jun 3 13:35:54 2016 GMT
        Subject: C=US, ST=CA, L=SJ, O=BRCD, OU=NI, CN=mlx2

Subject Public Key Info:
        Public Key Algorithm: id-ecPublicKey
            Public-Key: (384 bit)
            pub:
                04:c2:9d:56:3d:08:46:8b:d7:50:99:c2:38:41:66:
                0f:43:cb:3c:ac:f1:38:12:87:d4:cd:0e:10:a0:f6:
                e6:5b:4e:2f:96:55:46:fd:94:05:f5:89:bf:79:46:
                66:36:06:bf:a8:98:95:0c:0f:0b:d3:08:ef:68:90:
                07:d3:d1:47:a7:4a:1a:e9:4b:13:b0:d9:e7:a6:f1:
                c5:a1:8d:e1:b0:b4:c4:f3:c9:77:d8:a6:43:b0:f1:
                e1:5d:28:7b:e9:ba:d3
            ASN1 OID: secp384r1
        X509v3 extensions:
            X509v3 Key Usage: critical
                Digital Signature, Non Repudiation, Key Encipherment
            X509v3 Subject Key Identifier:
                21:BB:59:75:F5:EB:4D:8E:FB:E4:21:7E:3D:D1:A4:C0:CE:2C:23:97
            X509v3 Authority Key Identifier:
                keyid:90:07:32:9E:D3:C8:73:EC:9B:04:6B:22:60:84:B0:81:36:D5:9C:05

X509v3 CRL Distribution Points:

URI:file:////WIN-HJ98AK136A0.englab.example.com/CertEnroll/englab-WIN-HJ98AK136A0-CA-7.crl
URI: URI:http://WIN-HJ98AK136A0.englab.example.com/CertEnroll/englab-WIN-HJ98AK136A0-CA-7.crl

Authority Information Access:
Auth CA Issuers - URI:file:////WIN-HJ98AK136A0.englab.example.com/CertEnroll/WINHJ98AK136A0.
englab.example.com_englab-WIN-HJ98AK136A0-CA-7.crt
Auth CA Issuers - URI:http://WIN-HJ98AK136A0.englab.example.com/CertEnroll/WINHJ98AK136A0.
englab.example.com_englab-WIN-HJ98AK136A0-CA-7.crt

1.3.6.1.4.1.311.20.2:
1.3. .0.I.P.S.E.C.I.n.t.e.r.m.e.d.i.a.t.e.O.f.f.l.i.n.e
1.3. X509v3 Basic Constraints: critical
1.3. CA:FALSE
    Signature Algorithm: ecdsa-with-SHA384
        30:64:02:30:60:e0:ab:aa:9e:61:84:fe:e3:c9:11:44:f1:48:
        ea:45:7a:2c:4e:56:de:cc:f9:cc:8d:2e:2b:2b:b3:66:d9:fc:
        f5:d9:c9:e6:bf:6d:bd:5a:64:e3:2c:48:28:09:7f:0d:02:30:
        6a:65:67:3b:46:9e:84:62:8a:d3:da:69:85:6d:17:82:05:4f:
        b9:b1:f4:a8:ac:08:73:f7:28:ec:e8:36:29:c3:d8:03:28:7d:
        f9:2f:60:31:c7:39:8f:45:7b:d8:87:f3
```

- Enter the **show pki trustpoint** command to display information about a trustpoint.

```
devuce# show pki trustpoint auto

----------------PKI TRUSTPOINT ENTRY-----------------
CA: auto
VRF:
Entity Name: auto
  Common Name: mlx2
```

```
    Organization Name: BRCD
    Organization Unit Name: NI
    State Name: CA
    Country Name: US
    Location: SJ
Key Information:
    The key label is auto
    Public-Key: (384 bit)
    pub:
        04:c2:9d:56:3d:08:46:8b:d7:50:99:c2:38:41:66:
        0f:43:cb:3c:ac:f1:38:12:87:d4:cd:0e:10:a0:f6:
        e6:5b:4e:2f:96:55:46:fd:94:05:f5:89:bf:79:46:
        66:36:06:bf:a8:98:95:0c:0f:0b:d3:08:ef:68:90:
        07:d3:d1:47:a7:4a:1a:e9:4b:13:b0:d9:e7:a6:f1:
        c5:a1:8d:e1:b0:b4:c4:f3:c9:77:d8:a6:43:b0:f1:
        e1:5d:28:7b:e9:ba:d3
    ASN1 OID: secp384r1
Configured Fingerprint for authentication:
    26:fc:77:6c:b9:02:91:c9:a2:ab:60:28:c9:b9:c1:23:45:0a:8b:06
Configured local certificate url to send in IKE:
    http://192.0.2.1:8000/pki_local_cert_auto_mlx15
Enrollment Protocol: SCEP
Enrollment Profile: auto
```

- Enter the **show pki crls** command to display information about the current certificate revocation list (CRL).

```
device# show pki crls

----------------TRUST POINT auto CRL ENTRY-----------------
Certificate Revocation List (CRL):
        Version 2 (0x00000001)
    Signature Algorithm: ecdsa-with-SHA384
        Issuer: /DC=com/DC=example/DC=englab/CN=englab-WIN-HJ98AK136A0-CA-7
        Last Update: May 29 09:08:39 2015 GMT
        Next Update: Jun 5 21:28:39 2015 GMT
            CRL extensions:
                X509v3 Authority Key Identifier:
                    keyid:90:07:32:9E:D3:C8:73:EC:9B:04:6B:22:60:84:B0:81:36:D5:9C:05
1.3.6.1.4.1.311.21.1:
1.3. ...
1.3. X509v3 CRL Number:
1.3. 11
1.3. 1.3.6.1.4.1.311.21.4:
150605091839Z .
1.3. 1.3.6.1.4.1.311.21.14:
1.3. 0..0...........ldap:///CN=englab-WIN-HJ98AK136A0-CA-7,CN=WINHJ98AK136A0,
CN=CDP,CN=PublicKeyServices,CN=Services,DC=UnavailableConfigDN?certificateRevocationList?base?obje
ctClass=cRLDistributionPoint
Revoked Certificates:
    Serial Number: 1D0000001F3323F90AD37DD3A800000000001F
        Revocation Date: Apr 7 12:42:00 2015 GMT
            CRL entry extensions:
                X509v3 CRL Reason Code:
                    Certificate Hold
Signature Algorithm: ecdsa-with-SHA384
        30:66:02:31:00:98:47:03:5d:38:8d:73:1c:c6:da:b9:1a:a2:
        03:77:43:e5:a3:8f:07:12:a5:95:18:62:ec:1c:72:dc:80:9e:
        83:5e:8f:bc:0c:87:0f:39:87:ae:3d:20:94:23:a0:ca:2e:02:
        31:00:e6:e2:2c:17:da:de:e7:45:2c:db:75:26:f9:f2:6a:83:
        00:de:1d:a9:1a:33:a3:3d:1e:19:1f:cd:f8:2f:39:e8:86:94:
        4c:3c:60:9e:14:c5:cb:94:36:ee:3f:16:05:28
```

- Enter the **show pki counters** command to display the current PKI counter information.

```
device# show pki counters

------------PKI-COUNTERS----------
PKI Sessions Started: 335
PKI Sessions Ended: 0
PKI Sessions Active: 17
```

```
      Successful Validations: 0
      Failed Validations: 0
      Bypassed Validations: 0
      Pending Validations: 0
      CRLs checked: 0
      CRL - fetch attempts: 0
      CRL - failed attempts: 0
```

• Enter the **show pki enrollment-profile** command to display information about the PKI enrollment profiles.

```
device# show pki enrollment-profile

----------------PKI ENROLLMENT PROFILE ENTRY-----------------
Enrollment Profile: auto
Authentication Command: WIN-HJ98AK136A0.englab.example.com_englab-WIN-HJ98AK136A0-CA-7
Authentication URL: http://WIN-HJ98AK136A0.englab.example.com/CertSrv/mscep/mscep.dll
Enrollment URL: http://WIN-HJ98AK136A0.englab.example.com/CertSrv/mscep/mscep.dll
SCEP password: F1D43B07E91C82DE
```

• Enter the **show pki entity** command to display PKI entity information. The following example displays information about a PKI entity named example-entity.

```
device# show pki entity

----------------PKI ENTITY ENTRY-----------------
Entity Name: auto
 Common Name: mlx2
 Organization Name: BRCD
 Organization Unit Name: NI
 State Name: CA
 Country Name: US
 Location: SJ
```

• Enter the **show pki key mypubkey** command to display information about the current public keys on the router. Information can be displayed for generated keys only, manual keys only, or all keys. The following example shows information for all public keys on the router.

```
device# show pki key mypubkey ec all

----------------PKI PUBLIC KEY ENTRY-----------------
Public key of generated EC key pair:
  The key label is auto_ocsp
  Public-Key: (384 bit)
  pub:
      04:1b:24:46:20:7d:82:18:5b:a5:7b:4a:a5:ff:5e:
      55:61:0c:06:68:a2:1e:52:c5:77:bb:be:81:ed:1b:
      09:5f:f8:93:95:95:12:c6:61:96:c7:26:36:f5:3a:
      ad:78:76:bc:9e:d9:61:c8:82:d7:f0:c2:34:e8:f3:
      59:ae:8c:e5:91:f1:ad:ff:5d:ef:be:0c:47:54:82:
      ce:5e:a3:30:2a:ce:f4:82:aa:7e:f3:62:8d:d0:ba:
      c1:dd:cb:a3:e1:c5:f4
      ASN1 OID: secp384r1
```

Configuring global parameters for IKEv2

Global Internet Key Exchange version 2 (IKEv2) parameters are configured independently of peer configurations.

When you need to change the default values, use the following commands in global configuration mode to configure global parameters for IKEv2.

For further information and the default values for specific commands, refer to the Command Reference Guide.

- Use the **ikev2 cookie-challenge** command to enable the cookie-challenge option. The following example enables the cookie-challenge option when the maximum number of half-open IKEv2 security associations exceeds 1000.

```
device(config)# ikev2 cookie-challenge 1000
```

- Use the **ikev2 exchange-max-time** command to configure the maximum setup time, in seconds, for an IKEv2 message exchange. The following example shows how to set the maximum setup time for IKEv2 message exchange to 50 seconds.

```
device(config)# ikev2 exchange-max-time 50
```

- Use the **ikev2 http-url-support** command to enable sending of the hash and URL for the X509v2 certificate in the IKEv2 packet (instead of the actual X509v3 certificate).

```
device(config)# ikev2 http-url-support
```

- Use the **ikev2 limit max-in-negotiation-sa** command to configure the maximum number of in-negotiation IKEv2 SA sessions per LP. The following example shows how to set the maximum number of in-negotiation IKEv2 SA sessions to 10.

```
device(config)# ikev2 limit max-in-negotiation-sa 10
```

- Use the **ikev2 limit max-sa** command to configure the maximum number of IKEv2 SA sessions per LP. The following example shows how to set the maximum number of IKEv2 SA sessions to 200.

```
device(config)# ikev2 limit max-sa 200
```

- Use the **ikev2 nat-enable** to enable Network Address Translation Traversal (NAT-T).

```
device(config)# ikev2 nat-enable
```

- Use the **ikev2 nat-keepalive** to configure the keepalive time for NAT mappings. The following example set the keepalive time to nine seconds.

```
device(config)# ikev2 nat-keepalive 9
```

- Use the **ikev2 retransmit-interval** command to configure the delay time, in seconds, for resending IKEv2 messages. The following example shows how to set a delay time of 20 seconds.

```
device(config)# ikev2 retransmit-interval 20
```

- Use the **ikev2 retry-count** command to configure the number of attempts to retransmit an IKEv2 message. The following example shows how to set the number of retransmit attempts to 15.

```
device(config)# ikev2 retry-count 15
```

- Use the **show ikev2** command to display the configuration of the global IKEv2 parameters.

```
device(config)# show ikev2

IKEv2 Global data:
Retry Count         : 5          Max Exchange Time      : 60
Retransmit Interval : 5          Cookie Challenge Number : 0
Max Sa In Nego per LP: 256       Max Sa per LP          : 256
Allow Duplicate     : False      Http Cert Enable       : False
Total Peers         : 1          Total Ipsec Intf       : 1
Total Ike Sa        : 1          Total Ipsec Sa         : 2
NAT-T enabled       : True       NAT-T keep-alive time  : 9 sec
Sync In Progress    : 0          Time to syn peer to LP : 0
Pending Job         : 0          Sw pending time        : 0
```

The following example shows how to configure various global IKEv2 parameters.

```
device# configure terminal
device(config)# ikev2 cookie-challenge 1000
device(config)# ikev2 exchange-max-time 50
device(config)# ikev2 http-url-support
device(config)# ikev2 limit max-in-negotiation-sa 10
device(config)# ikev2 limit max-sa 200
device(config)# ikev2 nat-enable
device(config)# ikev2 nat-keepalive 9
device(config)# ikev2 retransmit-interval 20
device(config)# ikev2 retry-count 15
```

Configuring an IKEv2 proposal

Internet Key Exchange version 2 (IKEv2) proposal configuration sets parameters that are exchanged in the first phase of IKEv2 peer negotiations. After configuration, an IKEv2 proposal must be attached to an IKEv2 policy for use in IKEv2 negotiations.

There is a default IKEv2 proposal (def-ike-prop) that does not require configuration and has the following settings:

- Encryption algorithm: AES-CBC-256
- PRF algorithm: SHA-384
- Integrity algorithm: SHA-384
- DH group: 20

When the default values are not acceptable, perform the following task to configure an IKEv2 proposal. You only need to complete the steps for settings that you want to change.

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

   ```
   device# configure terminal
   ```

2. Create an IKEv2 proposal and enter configuration mode for the proposal.

   ```
   device(config)# ikev2 proposal prop_RTB
   ```

3. Configure an encryption algorithm for the proposal.

   ```
   device(config-ike-proposal-prop_RTB)# encryption aes-cbc-128
   ```

   This step adds the AES-CBC-128 algorithm to the encryption algorithms configured for prop_RTB. Because the AES-CBC-256 algorithm is configured by default, both the AES-CBC-256 and AES-CBC-128 algorithms are configured for prop_RTB after executing this step. Configuration of multiple encryption algorithms is allowed.

   When you want to configure the AES-CBC-128 algorithm only for the proposal, you must first add the AES-CBC-128 algorithm and then remove the default algorithm by using the **no encryption aes-cbc-256** command.

4. Configure an integrity algorithm for the proposal.

   ```
   device(config-ike-proposal-prop_RTB)# integrity sha256
   ```

   This step adds the SHA-256 algorithm to the integrity algorithms configured for prop_RTB. Because the SHA-384 algorithm is configured by default, both the SHA-384 and SHA-256 algorithms are configured for prop_RTB after executing this step. Configuration of multiple integrity algorithms is allowed.

   When you want to configure the SHA-256 algorithm only for the proposal, you must first add the SHA-256 algorithm and then remove the default algorithm by using the **no integrity sha384** command.

5. Configure a pseudorandom function (PRF) for the proposal.

   ```
   device(config-ike-proposal-prop_RTB)# prf sha256
   ```

   This step adds the SHA-256 algorithm to the PRF algorithms configured for prop_RTB. Because the SHA-384 algorithm is configured by default, both the SHA-384 and SHA-256 algorithms are configured for prop_RTB after executing this step. Configuration of multiple PRF algorithms is allowed.

   When you want to configure the SHA-256 algorithm only for the proposal, you must first add the SHA-256 algorithm and then remove the default algorithm by using the **no prf sha384** command.

6. Configure a DH group for the proposal.

   ```
   device(config-ike-proposal-prop_RTB)# dhgroup 19
   ```

   This step adds DH group 19 to the DH groups configured for prop_RTB. Because DH group 20 is configured by default, both DH groups (19 and 20) are configured for prop_RTB after executing this step. Configuration of multiple DH groups is allowed.

   When you want to configure DH group 19 only for the proposal, you must first add DH group 19 and then remove the default DH group by using the **no dhgroup 20** command.

7. Return to privileged EXEC mode.

   ```
   device(config-ike-proposal-prop_RTB)# end
   ```

8. Verify the IKEv2 proposal configuration.

   ```
   device# show ikev2 proposal prop_RTB

   ==========================================================================
   Name        : prop_RTB
   Encryption  : AES-CBC-256,AES-CBC-128
   Integrity   : sha384,sha256
   PRF         : sha384,sha256
   DH Group    : 384_ECP/Group 20,256_ECP/Group 19
   Ref Count   : 0
   ```

The following example shows how to create and configure an IKEv2 proposal named prop-RTB. This example also shows how to remove default configurations; that is, by first configuring an alternate algorithm or DH group and then removing the default configuration.

```
device# configure terminal
device(config)# ikev2 proposal prop_RTB
device(config-ike-proposal-prop_RTB)# encryption aes-cbc-128
device(config-ike-proposal-prop_RTB)# no encryption aes-cbc-256
device(config-ike-proposal-prop_RTB)# integrity sha256
device(config-ike-proposal-prop_RTB)# no intergrity sha384
device(config-ike-proposal-prop_RTB)# prf sha256
device(config-ike-proposal-prop_RTB)# no prf sha384
device(config-ike-proposal-prop_RTB)# dhgroup 19
device(config-ike-proposal-prop_RTB)# no dhgroup 20
device(config-ike-proposal-prop_RTB)# end
```

To use the IKEv2 proposal in IKEv2 negotiations, attach it to an IKEv2 policy by using the **proposal** command in IKEv2 policy configuration mode.

Configuring an IKEv2 policy

Internet Key Exchange version 2 (IKEv2) policy configuration specifies the IKEv2 proposal to be used by an IKEv2 policy and sets match parameters for the policy. An IKEv2 policy is used to protect IKEv2 peer negotiations.

Before configuring an IKEv2 policy, any IKEv2 proposal that is to be associated with the policy must be configured. There is an example at the end of this task that shows all the configuration steps in order.

There is a default IKEv2 policy (def-ike-policy) that does not require configuration. The default policy uses the default IKEv2 proposal (def-ike-prop) and matches:

- All local addresses because a local address to match is not configured for the policy
- Any VRF because a front-door VRF (FVRF) to match is not configured for the policy

    NOTE
    You should not configure overlapping policies. When multiple possible policy matches are configured, the policy that was created most recently is selected.

When the default policy is not acceptable, perform the following task to configure an IKEv2 policy. You only need to complete the steps for settings that you want to change.

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

    ```
    device# configure terminal
    ```

2. Create an IKEv2 policy and enter configuration mode for the policy.

    ```
    device(config)# ikev2 policy pol_RTB
    ```

3. (Optional) An IKEv2 policy must contain at least one IKEv2 proposal. By default, **def-ike-prop** is attached to an IKEv2 policy. Use the **proposal** command to configure an alternate IKEv2 proposal for the IKEv2 policy.

    ```
    device(config-ike-policy-pol_RTB)# proposal prop_RTB
    ```

    This example specifies using an IKEv2 proposal named prop_RTB for the policy.

4. (Optional) When a local match address is not specified, the IKEv2 policy matches all local addresses. Use the **match address-local** command to specify a local IP address as a match parameter for the IKEv2 policy.

    ```
    device(config-ike-policy-pol_RTB)# match address-local 10.3.3.3 255.255.255.0
    ```

    This example matches the IKEv2 policy pol_RTB based on a specific local IPv4 address (10.3.3.3 255.255.255.0) and the source address of the IPsec tunnel.

5. (Optional) When a front-door VRF (FVRF) to match is not specified, packets that match the local IP addresses specified for the policy are matched to any VRF. Use the **match fvrf** command to specify a front-door VRF for the policy.

    ```
    device(config-ike-policy-pol_RTB)# match fvrf vrf-name example_vrf
    ```

    This example specifies that the IKEv2 policy pol_RTB is matched when the local IPv4 address and the source address of the IPsec tunnel match and when the base VRF for the IPsec tunnel matches a VRF named example_vrf.

6. Return to privileged EXEC mode.

    ```
    device(config-ike-policy-pol_RTB)# end
    ```

7.   Verify the IKEv2 policy configuration.

```
device# show ikev2 policy pol_RTB

=========================================================================
Name               : pol_RTB
Vrf                : example_vrf
Local address/Mask : 10.3.3.3/255.255.255.0
Proposal           : prop_RTB
Ref Count          : 0
```

The following example configures an IKEv2 proposal named prop_RTB. It then configures an IKEv2 policy named pol_RTB using the **proposal** command to attach the IKEv2 proposal (prop_RTB) to the IKEv2 policy.

```
device# configure terminal
device(config)# ikev2 proposal prop_RTB
device(config-ike-proposal-prop_RTB)# encryption aes-cbc-128
device(config-ike-proposal-prop_RTB)# integrity sha256
device(config-ike-proposal-prop_RTB)# prf sha256
device(config-ike-proposal-prop_RTB)# dhgroup 19
device(config-ike-proposal-prop_RTB)# end

device# configure terminal
device(config)# ikev2 policy pol_RTB
device(config-ike-policy-pol_RTB)# proposal prop_RTB
device(config-ike-policy-pol_RTB)# match address-local 10.3.3.3 255.255.255.0
device(config-ike-policy-pol_RTB)# match fvrf vrf-name example_vrf
device(config-ike-policy-pol_RTB)# end
```

Configuring an IKEv2 authentication proposal

Internet Key Exchange version 2 (IKEv2) authentication proposal configuration sets parameters that are used to authenticate IKEv2 peer devices. After configuration, an IKEv2 authentication proposal must be attached to an IKEv2 profile for use in IKEv2 negotiations.

There is a default IKEv2 authentication proposal (def-ike-auth-prop) that does not require configuration and has the following settings:

*   Method for local device authentication: pre-shared
*   Method for remote device authentication: pre-shared
*   Pre-shared key: $QG5HTT1Ebk1TVW5NLWlhVW5ATVMhLS0rc1VA

When the default IKEv2 authentication proposal is not acceptable, perform the following task to configure an IKEv2 authentication proposal.

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

```
device# configure terminal
```

2.   Create an IKEv2 authentication proposal and enter configuration mode for the proposal.

```
device(config)# ikev2 auth-proposal auth_blue
```

3.   Specify an authentication method for local device authentication.

```
device(config-ike-auth-proposal-auth_blue)# method local pre-shared
```

This example specifies using a pre-shared key for local device authentication.

4.   Specify an authentication method for remote device authentication.

```
device(config-ike-auth-proposal-auth_blue)# method remote pre-shared
```

This example specifies using a pre-shared key for remote device authentication.

5.

   NOTE
   This step should only be completed when using PKI-based authentication.

   Specify the trusted certificate authority. The authority used to verify a certificate can differ from the authority used to sign the certificate.
   The following example show how to specify using tp1 only for signing of the certificate.

   ```
   device(config-ike-auth-proposal-auth_blue)# pki-trustpoint tp1 sign
   ```

   The following example show how to specify using tp2 only for verifying of the certificate.

   ```
   device(config-ike-auth-proposal-auth_blue)# pki-trustpoint tp2 verify
   ```

6. (Optional) There is a default pre-shared key that is assigned to an IKEv2 authentication proposal. Use the **pre-shared-key** command to specify an alternate pre-shared key. The following example configures a text-based pre-shared key (ps_key) for the proposal.

   ```
   device(config-ike-auth-proposal-auth_blue)# pre-shared-key ps_key
   ```

   The following example specifies a hexadecimal-based pre-shared key for the proposal.

   ```
   device(config-ike-auth-proposal-auth_blue)# pre-shared-key hex-based Abcxyz1290
   ```

7. Return to privileged EXEC mode.

   ```
   device(config-ike-auth-proposal-auth_blue)# end
   ```

8. Verify the IKEv2 authentication proposal configuration.

   ```
   device# show ikev2 auth-proposal auth_blue

   ========================================================================
   Ikev2 Auth-Proposal : auth_blue
   Local Auth Method   : pre_shared
   Remote Auth Method  : pre_shared
   pre-share-key       : $cTJkQ1x4Wnx7UQ==
   ```

   The encrypted form of the pre-shared key is displayed in the output of the **show ikev2 auth-proposal** command.

The following example creates and configures an IKEv2 authentication proposal named auth_blue.

```
device# configure terminal
device(config)# ikev2 auth-proposal auth_blue
device(config-ike-auth-proposal-auth_blue)# method local pre-shared
device(config-ike-auth-proposal-auth_blue)# method remote pre-shared
device(config-ike-auth-proposal-auth_blue)# pre-shared-key ps_key
device(config-ike-auth-proposal-auth_blue)# end
```

To use the IKEv2 authentication proposal in IKEv2 negotiations, attach it to an IKEv2 profile by using the **authentication** command in IKEv2 profile configuration mode.

Configuring an IKEv2 profile

Internet Key Exchange version 2 (IKEv2) profile configuration sets parameters that are exchanged in the second phase of IKEv2 peer negotiation. An IKEv2 profile specifies match identity criteria and the authentication proposal that is to be applied to an incoming connection. An IKEv2 profile may be used to protect a single VRF or all VRFs.

An IKEv2 session is a unique combination of local IP address, remote IP address, and IKEv2 profile. The IKEv2 profile determines the local identifier that is used for the session.

An IKEv2 profile is applied to an incoming IPsec connection by using match identity criteria presented by incoming IKEv2 connections such as IP address, fully qualified domain name (FQDN), and so on.

For an outgoing connection, the IKEv2 profile is determined by the IPsec profile used for the virtual tunnel interface (VTI).

A complete IKEv2 profile configuration contains a local identity, remote identity, local match identity, and remote match identity. When an IKEv2 profile configuration is incomplete, it is not used.

An IKEv2 VRF matches the forwarding VRF for the VTI.

Before configuring an IKEv2 profile, define and configure the IKEv2 authentication proposal that is to be associated with the profile. There is an example at the end of this task that shows all the configuration steps in order.

There is a default IKEv2 profile (def-ike-profile) that does not require configuration and that has the following settings:

- Authentication proposal: def-ike-auth-prop
- Local identifier address: 0.0.0.0.
- Lifetime period for an IKEv2 security association: 2592000 seconds
- Keepalive interval (the interval between sending IKEv2 messages to detect if a peer is still alive): 300 seconds

    NOTE
    The default IKEv2 profile protects any VRF.

    NOTE
    The method of authentication you select for IKEv2 transactions affects some IKEv2 profile configuration options. When configuring the **local-identifier**, **remote-identifier** and **match-identity**, it is recommended that you select:

    - Distinguished Name (DN), when using PKI-based authentication.
    - Fully Qualified Domain Name (FQDN), when using pre-shared key authentication.

When the default profile is not acceptable, perform the following task to configure an IKEv2 profile.

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

    ```
    device# configure terminal
    ```

2. Create an IKEv2 profile and enter configuration mode for the profile.

    ```
    device(config)# ikev2 profile prof_blue
    ```

3. (Optional) When an IKEv2 profile is created, the default IKEv2 authentication proposal (def-ike-auth-prop) is attached to the profile. Use the **authentication** command to attach an alternate authentication proposal to the profile.

    ```
    device(config-ike-profile-prof_blue)# authentication auth_blue
    ```

    This example attaches the auth-blue authentication proposal to the profile.

4. Specify a local system ID to be sent with the payload during peer negotiation.

    ```
    device(config-ike-profile-prof_blue)# local-identifier address 10.2.2.1
    ```

5. (Optional) To configure sending an initial contact message after reboot, enable the **initial-contact-payload** option.

    ```
    device(config-ike-profile-prof_blue)# initial-contact-payload
    ```

    When a device reboots, peer devices may have security associations (SAs) that are no longer valid. The initial contact message is sent so that peer devices delete the old SAs.

6. Specify a remote system ID to be sent with the payload during peer negotiation.

```
device(config-ike-profile-prof_blue)# remote-identifier address 10.3.3.3
```

7. An IKEv2 profile must contain at least one remote identity to match; it is not mandatory to specify a local identity to match. Specify a match identify for the peer device.
The following example shows how to specify matching on a remote identity (IP address 10.3.3.3).

```
device(config-ike-profile-prof_blue)# match-identity remote address 10.3.3.3
```

The following example shows how to specify matching on a local identity (IP address 10.3.3.3).

```
device(config-ike-profile-prof_blue)# match-identity local address 10.2.2.1
```

When multiple match identities are configured, a match occurs when any statement is matched with a peer's local identity or remote identity in remote or local match identity respectively.

8. Specify the name of the VRF that is to be protected.

```
device(config-ike-profile-prof_blue)# protected blue
```

This example specifies that the IKEv2 profile protects a VRF named blue.

9. Return to privileged EXEC mode.

```
device(config-ike-profile-prof_blue)# end
```

10. Verify the IKEv2 profile configuration.

```
device# show ikev2 profile prof_blue

========================================================================
IKEv2 Profile      : prof_blue
Auth Profile       : auth_blue
Match Criteria     :
 Inside VRF        : blue
  Local:
    address 10.2.2.1
  Remote:
    address 10.3.3.3
Local Identifier   : address 10.2.2.1
Remote Identifier  : address 10.3.3.3
Lifetime           : 2592000 sec
Keepalive Check    : 300 sec
Ref Count          : 0
```

The following example configures an IKEv2 authentication proposal named auth_blue. It then configures an IKEv2 profile named prof_blue using the **authentication** command to attach the IKEv2 authentication proposal to the IKEv2 profile.

```
device# configure terminal
device(config)# ikev2 auth-proposal auth_blue
device(config-ike-auth-proposal-auth_blue)# method local pre-shared
device(config-ike-auth-proposal-auth_blue)# method remote pre-shared
device(config-ike-auth-proposal-auth_blue)# pre-shared-key ps_key
device(config-ike-auth-proposal-auth_blue)# end

device# configure terminal
device(config)# ikev2 profile prof_blue
device(config-ike-profile-prof_blue)# authentication auth_blue
device(config-ike-profile-prof_blue)# local-identifier address 10.2.2.1
device(config-ike-profile-prof_blue)# remote-identifier address 10.3.3.3
device(config-ike-profile-prof_blue)# match-identity local address 10.2.2.1
device(config-ike-profile-prof_blue)# protected blue
device(config-ike-profile-prof_blue)# match-identity remote address 10.3.3.3
device(config-ike-profile-prof_blue)# end
```

To use the IKEv2 profile for outgoing connections, attach it to an IPsec profile by using the **ike-profile** command in IPsec profile configuration mode.

Configuring an IPsec proposal

IPsec proposal configuration sets encryption parameters for IPsec. An IPsec proposal is activated by attaching it to an IPsec profile.

When IPsec is initialized, a default IPsec proposal (def-ipsec-prop) is created with the following settings:

- Transform type: ESP
- Encapsulation mode: Tunnel
- Encryption algorithm: AES-GCM-256 (Both authentication and encryption are performed by this algorithm.)
- Extended sequence numbers: Disabled

When the default proposal is not acceptable, perform the following task to configure an IPsec proposal.

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

   ```
   device# configure terminal
   ```

2. Create an IPsec proposal and enter configuration mode for the proposal.

   ```
   device(config)# ipsec proposal prop_blue
   ```

3. **NOTE**
   (Optional). By default, the encapsulation mode is set to tunnel. Because this is the only mode currently available, this step is not required.

   Specify the encapsulation mode for the proposal.

   ```
   device(config-ipsec-proposal-prop_blue)# encapsulation-mode tunnel
   ```

4. **NOTE**
   (Optional). By default, the transform type is set to ESP. Because this is the only mode currently available, this step is not required.

   Specify the transform type for the proposal.

   ```
   device(config-ipsec-proposal-prop_blue)# transform esp
   ```

5. **NOTE**
   If you choose to use the AES-GCM-128 algorithm for encryption and decryption of IP packets transmitted across the tunnel, make sure that:
   - The tunnel end nodes (both local and remote) both use NetIron 5.9.0a or later. If the tunnel end nodes are not running the same version, the tunnel will not come up successfully. To prevent this issue, upgrade any tunnel end node that is using an earlier version (such as 5.9.0) to ensure that both nodes are using NetIron version 5.9.0a or later.
   - The remote device has consistent configuration parameter settings for the tunnel.

   Specify an encryption algorithm for the proposal.

   ```
   device(config-ipsec-proposal-prop_blue)# encryption-algorithm aes-gcm-128
   ```

   This step adds the AES-GCM-128 algorithm to the encryption algorithms configured for prop_blue. Because the AES-GCM-256 algorithm is configured by default, after executing this step both the AES-GCM-256 and AES-GCM-128 algorithms are configured for prop_blue. Configuration of multiple encryption algorithms is allowed.

   When you want to configure the AES-CBC-128 algorithm only for the proposal, you must first add the AES-CBC-128 algorithm and then remove the default algorithm by using the **no encryption-algorithm aes-cbc-256** command.

   For an IPsec tunnel to come up successfully, IPsec peer devices must be configured with a common encryption algorithm.

6. (Optional) Enable extended sequence number (ESN).

   ```
   device(config-ipsec-proposal-prop_blue)# esn-enable
   ```

   ESN is enabled when an IPsec proposal is to be attached to an IPsec profile with the **replay-protection** option enabled.

   For an IPsec tunnel to operate correctly, the **esn** and **replay-protection** options must have the same configuration, that is; both settings must be enabled or both settings must be disabled. In addition, the configuration must be identical on each of the tunnel endpoints.

7. Return to privileged EXEC mode.

   ```
   device(config-ipsec-proposal-prop_blue)# end
   ```

8. Verify the IPsec proposal configuration.

   ```
   device# show ipsec proposal prop_blue

   ========================================================================
   Name               : prop_blue
   Protocol           : ESP
   Encryption         : aes-gcm-256,aes-gcm-128
   Authentication     : NULL
   ESN                : Disable
   Mode               : Tunnel
   Ref Count          : 0
   ```

The following example creates and configures an IPsec proposal named prop_blue.

```
device# configure terminal
device(config)# ipsec proposal prop_blue
device(config-ipsec-proposal-prop_blue)# encapsulation-mode tunnel
device(config-ipsec-proposal-prop_blue)# transform esp
device(config-ipsec-proposal-prop_blue)# encryption-algorithm aes-gcm-128
device(config-ipsec-proposal-prop_blue)# esn-enable
device(config-ipsec-proposal-prop_blue)# end
```

To use the IPsec proposal to encrypt data in an IPsec tunnel, attach it to an IPsec profile by using the **proposal** command in IPsec profile configuration mode.

Configuring an IPsec profile

IPsec profile configuration sets parameters used to encrypt data between IPsec peer devices. After configuration, an IPsec profile is activated by attaching it to a virtual tunnel interface (VTI).

Before configuring an IPsec profile, any IKE profile or IPsec proposal that is to be attached to it must be configured. There is a configuration example at the end of this task that shows all the steps in order.

You can configure an IPsec profile by completing the following task.

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

   ```
   device# configure terminal
   ```

2. Create an IPsec profile and enter configuration mode for the profile.

   ```
   device(config)# ipsec profile prof_blue
   ```

   When an IPsec profile is created, the default IKEv2 profile (def-ike-prof) and the default IPsec proposal (def-ipsec-prop) are automatically attached to the profile.

3. (Optional) Use the **ike-profile** command to attach an alternate IKEv2 profile to the IPsec profile.

   ```
   device(config-ipsec-profile-prof_blue)# ike-profile prof_blue
   ```

   This example attaches the prof-blue IKEv2 profile to the IPsec profile.

4. (Optional) Use the **proposal** command to attach an alternate IPsec proposal to the IPsec profile.

   ```
   device(config-ipsec-profile-prof_blue)# proposal prop_blue
   ```

   This example attaches the prop-blue IPsec proposal to the IPsec profile.

5. (Optional) Configure the lifetime of an IPsec security association (SA).
   The following example shows how to set the IPsec SA lifetime to 240 minutes (14400 seconds).

   ```
   device(config-ipsec-profile-prof_blue)# lifetime 240
   ```

6. (Optional) By default, replay protection is disabled. Use the **replay-protection** command to enable it.

   ```
   device(config-ipsec-profile-prof_blue)# replay-protection
   ```

   When the **replay-protection** is enabled for an IPsec profile, the IPsec proposal that is attached to the profile should be configured with the **esn** option enabled.

   For an IPsec tunnel to operate correctly, the **esn** and **replay-protection** options must have the same configuration, that is; both settings must be enabled or both settings must be disabled. In addition, the configuration must be identical on each of the tunnel endpoints.

7. Return to privileged EXEC mode.

   ```
   device(config-ipsec-profile-prof_blue)# end
   ```

8. Verify the IPsec profile configuration.

```
device# show ipsec profile prof_blue

=======================================================================
Name                 : prof_blue
Ike Profile          : prof_blue
Lifetime             : 14400 sec
Anti-Replay Service  : Enabled
Replay Window Size   : 64
DH Group             : None
Proposal             : prop_blue
```

The following example shows how to configure an IKEv2 authorization proposal, an IKEv2 profile, an IPsec proposal and an IPsec profile. The IKEv2 authorization proposal is used in the configuration of the IKEv2 profile. The IKEv2 profile and IPsec proposal are used in the configuration of the IPsec profile.

```
device# configure terminal
device(config)# ikev2 auth-proposal auth_blue
device(config-ike-auth-proposal-auth_blue)# method local pre-shared
device(config-ike-auth-proposal-auth_blue)# method remote pre-shared
device(config-ike-auth-proposal-auth_blue)# pre-shared-key ps_key
device(config-ike-auth-proposal-auth_blue)# end

device# configure terminal
device(config)# ikev2 profile prof_blue
device(config-ike-profile-prof_blue)# authentication auth_blue
device(config-ike-profile-prof_blue)# local-identifier address 10.2.2.1
device(config-ike-profile-prof_blue)# remote-identifier address 10.3.3.3
device(config-ike-profile-prof_blue)# match-identity local address 10.2.2.1
device(config-ike-profile-prof_blue)# protected blue
device(config-ike-profile-prof_blue)# match-identity remote address 10.3.3.3
device(config-ike-profile-prof_blue)# end

device# configure terminal
device(config)# ipsec proposal prop_blue
device(config-ipsec-proposal-prop_blue)# transform esp
device(config-ipsec-proposal-prop_blue)# encryption-algorithm aes-gcm-256
device(config-ipsec-proposal-prop_blue)# end

device# configure terminal
device(config)# ipsec profile prof_blue
device(config-ipsec-profile-prof_blue)# ike-profile prof_blue
device(config-ipsec-profile-prof_blue)# proposal prop_blue
device(config-ipsec-profile-prof_blue)# end
```

To activate the IPsec profile, bind it to a virtual tunnel interface (VTI) by using the **tunnel protection ipsec profile** command in tunnel interface configuration mode.

Activating an IPsec profile on a VTI

An IPsec profile is activated by binding it to a virtual tunnel interface (VTI) that is configured as an IPsec VTI.

The IPsec profile must be defined and configured before binding to the VTI. The tunnel interface must be configured also. There is an example at the end of this task that shows the configuration steps in order.

You can activate an IPsec profile on a VTI by performing the following task.

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

```
device# configure terminal
```

2. Enter configuration mode for an IPsec tunnel.

    ```
    device (config)# interface tunnel 1
    ```

3. Set the mode of the tunnel to IPsec.

    ```
    device(config-tnif-1)# tunnel mode ipsec ipv4
    ```

4. Specify the IPsec protection profile for the tunnel.

    ```
    device(config-tnif-1)# tunnel protection ipsec profile prof_blue
    ```

5. Return to privileged EXEC mode.

    ```
    device(config-tnif-1)# end
    ```

6. Verify that the IPsec profile is attached to the VTI.

    ```
    device# show running-config interface tunnel 1
    !
    interface tunnel 1
     tunnel mode ipsec ipv4
     tunnel protection ipsec profile prof_blue
    !
    ```

The following example shows how to configure a VTI, set the mode of the tunnel to IPsec, and bind an IPsec profile to the VTI.

```
device# configure terminal
device (config)# interface tunnel 1
device(config-tnif-1)# vrf forwarding blue
device(config-tnif-1)# tunnel source ethernet 1/1
device(config-tnif-1)# tunnel destination 10.2.2.1
device(config-tnif-1)# ip address 11.1.1.1/24

device(config-tnif-1)# tunnel mode ipsec ipv4
device(config-tnif-1)# tunnel protection ipsec profile prof_blue
```

Routing traffic over IPsec using static routing

Traffic can be routed over an IPsec tunnel by configuring a static route.

To steer traffic over an IPsec tunnel by configuring a static route, complete the following task.

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

    ```
    device# configure terminal
    ```

2. Define a static route.

    ```
    device(config)# ip route 10.157.23.0/24 10.4.4.4
    ```

    This example defines a static route with destination IP address 10.157.23.0/24 and where the next-hop address 10.4.4.4 is reachable through the IPsec tunnel.

    Alternatively, you can define the static route by specifying the tunnel itself as the outgoing interface.

    ```
    device(config)# ip route 10.157.23.0/24 tunnel 2
    ```

    This example defines a static route with destination IP address 10.157.23.0/24 and specifies tunnel 2 as the outgoing interface.

The following example shows how to configure an IPsec VTI and how to steer traffic over the tunnel by configuring a static route.

```
device# interface tunnel 2
device(config-tnif-2)# vrf forwarding blue
device(config-tnif-2)# tunnel source ethernet 1/2
device(config-tnif-2)# tunnel destination 10.2.2.1
device(config-tnif-2)# tunnel mode ipsec ipv4
device(config-tnif-2)# tunnel protection ipsec profile prof-blue
device(config-tnif-2)# ip address 10.4.4.4/24
device(config-tnif-2)# exit

device(config)# ip route vrf blue 10.157.23.0/24 tunnel 2
```

Routing traffic over IPsec using PBR

Traffic can be configured to route over an IPsec tunnel by using policy-based routing (PBR).

Before configuring traffic to route over an IPsec tunnel, the virtual tunnel interface (VTI) must be configured. There is an example at the end of this task that shows the configuration steps in order.

To route traffic over an IPsec tunnel using PBR, complete the following task.

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

   ```
   device# configure terminal
   ```

2. Define an access control list (ACL).

   ```
   device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
   ```

3. Create a route map to accept this traffic and enter route map configuration mode.

   ```
   device(config)# route-map crypto-route permit 99
   ```

   This example creates a route map called "crypto-route".

4. Specify the values to match for the route map.

   ```
   device(config-routemap crypto-route)# match ip address 99
   ```

   This example specifies matching based on the IP address in ACL 99.

5. Configure matching packets to route over the IPsec tunnel.

   ```
   device(config-routemap crypto-route)# set ip next-hop 10.4.4.4
   ```

   This example configures IPsec VTI 10.4.4.4 as the next-hop address for matching packets.
   Alternatively, you can configure the tunnel itself as the next-hop address for matching packets.

   ```
   device(config-routemap crypto-route)# set next-hop-ip-tunnel 2
   ```

   This example configures tunnel 2 as the next-hop address for matching packets.

6. Enter configuration mode on the interface where you want to apply the route map.

   ```
   device(config-routemap crypto-route)# interface ethernet 1/3
   ```

   This example enters configuration mode on Ethernet interface 1/3.

7. Enable policy-based routing on the interface and specify the route map to be used.

   ```
   device(config-if-e1000-1/3)# ip policy route-map crypto-route
   ```

   This example specifies using the "crypto-route" route map for PBR on Ethernet interface 1/3.

8. Return to global configuration mode.

```
device(config-if-e1000-1/3)# end
```

The following example shows how to configure an IPsec VTI and how to steer traffic over the tunnel by using PBR.

```
device# interface tunnel 1
device(config-tnif-1)# vrf forwarding blue
device(config-tnif-1)# tunnel source ethernet 1/1
device(config-tnif-1)# tunnel destination 10.2.2.1
device(config-tnif-1)# tunnel mode ipsec ipv4
device(config-tnif-1)# tunnel protection ipsec profile prof-blue
device(config-tnif-1)# ip address 10.4.4.4/24
device(config-tnif-1)# exit

device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
device(config)# route-map crypto-route permit 99
device(config-routemap crypto-route)# match ip address 99
device(config-routemap crypto-route)# set ip next-hop 10.4.4.1 vrf blue
device(config-routemap crypto-route)# end
device(config)# interface ethernet 1/3
device(config-if-e1000-1/3)# vrf forwarding blue
device(config-if-e1000-1/3)# ip policy route-map crypto-route
device(config-if-e1000-1/3)# end
```

Re-establishing SAs

An IKEv2 SA or IPsec SA is re-established after the SA is cleared. When the SA is cleared, any existing child SAs are also cleared and re-established.

- Enter the **clear ikev2 sa** command to clear and re-establish an IKEv2 SA. The following example shows how to clear an IKEv2 SA (and any existing child SAs) by specifying the local IPv4 address for the SA.

```
device# clear ikev2 sa local 10.1.1.1
```

- Enter the **clear ipec sa** command to clear and re-establish an IPsec SA. The following example shows how to clear an IPsec SA by specifying the peer address for the SA.

```
device# clear ipsec sa peer 10.2.2.2
```

NOTE
An IKEv2 SA or an IPsec SA is also re-established when the lifetime period configured for the SA expires.

Enabling learning of self MAC addresses

When encrypted or decrypted IP packets are looped back to the system (device) for an additional level of encryption or decryption, the packets are not sent to the CPU to learn the self MAC addresses on the device.

> **NOTE**
> By default the learning of self MAC addresses for IPsec tunnels is disabled. After enablement, it should be disabled when it is no longer needed.

- Enter the **ipsec-sa-learning-enable** command to enable the learning of self MAC addresses for all configured IPv4 and IPv6 tunnels.

  ```
  device# configure terminal
  device(config)# ipsec-sa-learning-enable
  ```

- Enter the **no** version of the command to disable the learning of self MAC addresses for all configured IPv4 and IPv6 tunnels.

  ```
  device(config)# no ipsec-sa-learning-enable
  ```

Configuring PIM on an IPsec VTI

Protocol Independent Multicast (PIM) can be configured over an IPsec virtual tunnel interface (VTI).

Prior to configuring PIM, an IPsec VTI must be configured and an IPsec profile must be configured and bound to the IPsec VTI. There is an example at the end of this task that shows all the configuration steps in order.

Enable PIM on a VTI by performing the following task.

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

   ```
   device# configure terminal
   ```

2. Enter configuration mode for an IPsec tunnel.

   ```
   device (config)# interface tunnel 1
   ```

3. Enable PIM on the interface.

   ```
   device(config-tnif-1)# ip pim-sparse
   ```

   This example enables PIM Sparse Mode (PIM-SM) on the VTI interface.

4. Return to privileged EXEC mode.

   ```
   device(config-tnif-1)# end
   ```

5. Verify that PIM is configured on the VTI.

   ```
   device# show running-config interface tunnel 1
   !
   interface tunnel 1
   ...                     output truncated for brevity
    ip pim-sparse
   !
   ```

The following example shows how to configure a VTI, set the mode of the tunnel to IPsec, bind an IPsec profile named prof_blue to the VTI, and enable PIM Sparse Mode (PIM-SM) on the VTI.

```
device# configure terminal
device (config)# interface tunnel 1
device(config-tnif-1)# vrf forwarding blue
device(config-tnif-1)# tunnel source ethernet 1/1
device(config-tnif-1)# tunnel destination 10.2.2.1
device(config-tnif-1)# ip address 11.1.1.1/24

device(config-tnif-1)# tunnel mode ipsec ipv4
device(config-tnif-1)# tunnel protection ipsec profile prof_blue

device(config-tnif-1)# ip pim-sparse
```

Configuring extended logging for IKEv2 and PKI

The logging of IKEv2 and PKI transaction details can be extended to record more detailed information. The logging infrastructure must be set up prior to enabling extended logging for IKEv2 and PKI.

Before you can configure IKEv2 and PKI extended logging, you must make sure that:

- The logging infrastructure required to generate and transport IKEv2 and PKI syslogs is completely set up and ready. For example, the commands used to disable syslog server, disable message-level logging must not be configured. If the disable commands are configured, syslogs cannot be sent to the external server or saved to the management module.
- Any ACL or PBR configuration involving the syslog server will not prevent the system from sending IKEv2 or PKI syslogs to the external server over the IPsec tunnel. If needed, modify the ACL or PBR configuration.
- A dedicated IPsec tunnel exists through which the external syslog server can be reached.
- When operating in Federal Information Processing Standard (FIPS) mode, the dedicated IPsec tunnel to the external syslog server must be the only communication channel that can be used to reach the server. This can be done by adding a static route with a higher priority.

> NOTE
> VPN Gateway Certification requires that a dedicated IPsec tunnel is configured to transport syslogs to the external syslog server.

You have the option of enabling extended logging for both IKEv2 and PKI, IKEv2 only, or PKI only.

Perform the following task to configure extended logging for IKEv2 and PKI.

1. Enter the **logging** command to enable extended logging for the external syslog server. (This ensures that the transaction detail syslogs are sent to the external syslog server.)

   ```
   device(config)#logging  host 10.2.2.1 udp-port 3
   ```

   > NOTE
   > The IP address of the syslog server can be an IPv4 or IPv6 address.

2. By default, extended logging for IKEv2 is disabled. Enter the **logging enable** command specifying the **ikev2 extended** options to enable extended logging for IKEv2.

   ```
   device(config)#logging enable ikev2 extended
   ```

3. By default, extended logging for PKI is disabled. Enter the **logging enable** command specifying the **pki pki-extended** options to enable extended logging for PKI.

```
device(config)#logging enable pki pki-extended
```

Disabling traps and syslog messages for IKEv2 and IPsec

You can disable error traps and syslog messages for Internet Key Exchange version 2 (IKEv2) and IPsec. By default, traps and syslog messages for both IKEv2 and IPsec are enabled.

Use the following commands to disable error traps and syslog messages for IKEv2 and IPsec:.

- The **no snmp-server enable traps** command disables error traps for IKEv2 and IPsec.

  The following example shows how to disable error traps for IKEv2.

  ```
  device(config)# no snmp-server enable traps ikev2
  ```

  The following example shows how to disable error traps for IPsec.

  ```
  device(config)# no snmp-server enable traps ipsec
  ```

- The **no logging enable ikev2** command disables syslog messages for IKEv2.

  ```
  device(config)# no logging enable ikev2
  ```

- The **no loging enable ipsec** command disables syslog messages for IPsec.

  ```
  device(config)# no logging enable ipsec
  ```

Displaying IPsec module information

Some show commands can be used to display information about the status of an installed IPsec interface module and about the current utilization of the module.

- Enter the **show module** command to display module status information.

  ```
  device# show module

     Module                                      Status             Ports Starting MAC
  M1 (left ):NI-XMR-MR Management Module         Active
  M2 (right):BR-MLX-MR2-X Management Module      Standby(Ready State)
  F1: NI-X-HSF Switch Fabric Module              Active
  F2: NI-X-HSF Switch Fabric Module              Active
  F3:
  S1: BR-MLX-10Gx4-M-IPSEC 4-port 1/10GbE
      and 4-port 1GbE Module                     CARD_STATE_UP      8     cc4e.241e.c100
  S2:
  S3: BR-MLX-1GCx24-X 24-port 10/100/1000
      Base-T Copper Module                       CARD_STATE_UP      24    cc4e.241e.c160
  S4:
  ```

Displaying IKEv2 configuration information

Various show commands can be used to display information about IKEv2 configurations.

IKEv2 must be configured before displaying this information.

- Enter the show ikev2 command to display the current IKEv2 configuration.

```
device(config)# show ikev2

IKEv2 Global data:
Retry Count          : 5          Max Exchange Time       : 60
Retransmit Interval  : 5          Cookie Challenge Number : 0
Max Sa In Nego per LP: 256        Max Sa per LP           : 256
Allow Duplicate      : False      Http Cert Enable        : False
Total Peers          : 1          Total Ipsec Intf        : 1
Total Ike Sa         : 1          Total Ipsec Sa          : 2
NAT-T enabled        : True          NAT-T keep-alive time   : 5 sec
Sync In Progress     : 0          Time to syn peer to LP  : 0
Pending Job          : 0          Sw pending time         : 0
```

- Enter the **show ikev2 proposal** command to display information about IKEv2 proposal configurations. The Ref Count shows the number of uses of this IKEv2 proposal.

```
device# show ikev2 proposal
======================================================================================
Name               : def-ike-prop
Encryption         : aes-cbc-256
Integrity          : sha384
Prf                : sha384
DH Group           : 384_ECP/Group 20
Ref Count          : 3
======================================================================================
```

- Enter the **show ikev2 policy** command to display information about IKEv2 policy configurations.

```
device# show ikev2 policy

======================================================================================
Name               : def-ike-v6-policy
Vrf                : any
Local Address/Prefix: ::/0
Proposal           : def-ike-prop
Ref Count          : 0
======================================================================================
Name               : def-ike-policy
Vrf                : any
Local Address/Mask : 0.0.0.0/0.0.0.0
Proposal           : def-ike-prop
Ref Count          : 12
```

- Enter the **show ikev2 auth-proposal** command to display information about IKEv2 authentication proposal configurations.

```
device# show ike2 auth-proposal

======================================================================================
Ikev2 Auth-Proposal : def-ike-auth-prop
Local Auth Method   : pre_shared
Remote Auth Method  : pre_shared
pre-share-key(text) : $QG5HTT1Ebk1TVW5NLWIhVW5ATVMhLS0rc1VA
```

- Enter the **show ikev2 profile** command to display information about IKEv2 profile configurations.

```
device# show ike profile

======================================================================================
IKEv2 Profile      : def-ike-profile
```

```
Auth Profile        : def-ike-auth-prop
Match Criteria      :
 Inside VRF         : any
   Local:
   Remote:
Local Identifier    : address 0.0.0.0
Lifetime            : 2592000 sec
Keepalive Check     : 300 sec
Ref Count           : 0
================================================================================
IKEv2 Profile       : test
Auth Profile        : def-ike-auth-prop
Match Criteria      :
 Inside VRF         : any
   Local:
   Remote:
Local Identifier    : address 1.1.1.2
Lifetime            : 2592000 sec
Keepalive Check     : 300 sec
Ref Count           : 0
```

- Enter the **show ikev2 sa** command to display summary information about IKEv2 security association (SA) configurations.

```
device# show ikev2 sa

tnl-id    local             remote        Status  vrf(i)      vrf(f)
--------------------------------------------------------------------------------
tnl 2    1.2.10.1/500      1.2.10.2/500  rdy     Blue        Default
```

- Enter the **show ikev2 sa detail** command to display detailed information about IKEv2 SA configurations.

```
device# show ikev2 sa detail

Total SA : 1
Active SA: 1   : Constructing SA:0      : Dying SA:0
-----------------------------------------------------------------------------------------------------
---------------
tnl 23        100.23.23.2/500                              100.23.23.1/500       active
vrf3              default-vrf
-----------------------------------------------------------------------------------------------------
----------------
Role                : Responder
Local SPI           : 0x2708deebeec67f20    Remote SPI: 0x84628d4dea442668
Ike Profile         : 23
Ike Policy          : 251
Auth Proposal       : 23
NAT-T is detected
Local node behind NAT: True
Remote node behind NAT: False
```

- Enter the **show ikev2 session** command to display summary information about IKEv2 session configurations.

```
device# show ikev2 session

IKE count:1, CHILD count:1
Tunnel-id       Local             Remote              Status      vrf(i) vrf(f)
----------------------------------------------------------------------
Tnl 2          1.2.10.1/500     1.2.10.2/500        rdy|in-use  Blue   Default
child sa:
 id 1
       local selector  0.0.0.0/0 - 255.255.255.255/65535
       remote selector 0.0.0.0/0 - 255.255.255.255/65535
       ESP spi in/out: 0x0000004b/0x0000005e
       Encryption: aes-gcm-256, ICV Size: 16 octects, Esp_hmac: null
       Authetication: null  DH Group:none , Mode: tunnel
```

- Enter the **show ikev2 session detail** command to display detailed information about IKEv2 session configurations.

```
device# show ikev2 session detailed

IKE count:1, CHILD count:1

Tunnel-id  Local                    Remote                  Status        vrf(p) vrf(f)
-------------------------------------------------------------------------------
2          1.2.10.1/500             1.2.10.2/500            rdy|in-use Blue    Default
        Encr: aes-cbc-256, Hash: sha384, DH Grp:384_ECP/Group 20, Auth: not supported
     Life/Active Time: 86400/361 sec
     Status Description: Negotiation done
        Local spi: f7c029048eb25082     Remote spi: 56b8735e2f6afbde
        Local id : address 1.2.45.2     Remote id : address 1.2.45.1
        No Exchange in Progress
        Next Request Message id=29
          Total Keepalive sent: 0     Total Keepalive Received: 0
        Time Past Since Last Msg: 60

child sa:
 id 1
        local selector  0.0.0.0/0 - 255.255.255.255/65535
        remote selector 0.0.0.0/0 - 255.255.255.255/65535
        ESP spi in/out: 0x0000004b/0x0000005e
        Encryption: aes-gcm-256, ICV Size: 16 octects, Esp_hmac: null
        Authetication: null  DH Group:none , Mode: tunnel
```

Displaying IPsec configuration information

Various show commands can be used to display IPsec configuration information.

IPsec must be configured before displaying this information.

- Enter the **show ipsec proposal** command to display information about IPsec proposal configurations.

```
device# show ipsec proposal

    Name          : prop_red
    Protocol      : ESP
    Encryption    : aes-gcm-256
    Authentication: NULL
    ESN           : Enable
    Mode          : Tunnel
```

- Enter the **show ipsec profile** command to display information about IPsec profile configurations.

```
device# show ipsec profile

Name                : red
Ike Profile         : red
Lifetime            : 28800
Anti-replay service : Enabled
 Replay window size : 64
DH group            : None
Proposal            : red
```

- Enter the **show ipsec sa** command to display summary information about IPsec security association (SA) configurations.

```
device# show ipsec sa
            IPSEC Security Association Database(Entries:2)
SPDID(vrf:if) Dir Encap SPI        Destination  AuthAlg  EncryptAlg  Status Mode
0:v2          out ESP   400        ::           sha1     Null        ACT    TRAN
0:v2          in  ESP   400        FE80::       sha1     Null        ACT    TRAN
1:Tun1        in  ESP   0xBD481319 1.2.10.2     Null     AES-GCM-256 ACT    TNL
1:Tun1        out ESP   0x9EAB77D6 1.2.10.2     Null     AES-GCM-256 ACT    TNL
```

- Enter the **show ipsec sa address** command to display detailed information about a specific IPsec SA by specifying the local address of the SA.

```
device# show ipsec sa address 1.2.10.2 detail

Total ipsec SAs: 2

0:
    interface          : tnl 1
     Local address: 1.2.45.1/500, Remote address: 1.2.45.2/500
     Inside vrf: default-vrf
     Local identity (addr/mask/prot/port): address(0.0.0.0/0/0/0)
     Remote identity(addr/mask/prot/port): address(0.0.0.0/0/0/0)
     DF-bit: clear
     Profile-name: red
     DH group: none
     Direction: inbound, SPI: 0x0000004b
     Mode: tunnel,
     Protocol: esp, Encryption: gcm-256, Authentication: null
     ICV size: 16 bytes
     lifetime(sec): Expiring in (4606816/3576)
     Anti-replay service: Enabled, Replay window size: 0
     Status: ACTIVE
     slot Assigned 0
     nht_index 0000ffff
     Is tunnel NHT: false
1:
    interface          : tnl 1
     Local address: 1.2.45.1/500, Remote address: 1.2.45.2/500
     Inside vrf: default-vrf
     Local identity (addr/mask/prot/port): address(0.0.0.0/0/0/0)
     Remote identity(addr/mask/prot/port): address(0.0.0.0/0/0/0)
     DF-bit: clear
     Profile-name: red
     DH group: none
     Direction: inbound, SPI: 0x0000009c
     Mode: tunnel,
     Protocol: esp, Encryption: gcm-256, Authentication: null
     ICV size: 16 bytes
     lifetime(k/sec): Expiring in (4606816/3576)
     Anti-replay service: Enabled, Replay window size: 0
     Status: ACTIVE
     slot Assigned 0
     nht_index 00000004
     Is tunnel NHT: true
```

Displaying and clearing statistics for IKEv2 and IPsec

Various commands can be used to display and clear statistical information for IKEv2 and IPsec.

- Enter the **show ikev2 statistics** command to display IKEv2 statistics.

```
device# show ikev2 statistics

Total IKEv2 SA Count   : 1 active: 1 negotiating: 0
Incoming IKEv2 Requests: 0 accepted: 0 rejected: 0
Outgoing IKEv2 Requests: 1 accepted: 1 rejected: 0
Rejected IKEv2 Requests: 0
Incoming IKEV2 Cookie Challenged Requests: 0
accepted: 0 rejected: 0 rejected no cookie: 0
IKEv2  Packet Statistics:
    Total Packets Received     : 57
    Total  Packets Transmitted : 57
    Total Packets Retransmitted: 0
    Total Keepalive Received   : 10
    Total Keepalive Transmitted: 10
```

```
IKEv2 Error Statistics:
    Unsupported Payload  : 0     Invalid IKE SPI  : 0
    Invalid Version      : 0     Invalid Syntax   : 0
    Proposal Mismatch    : 0     Invalid Selectors: 0
    Authentication Failed : 0    Others           : 0
```

- Enter the **clear ikev2 statistics** command to reset or clear IKEv2 statistics.

```
device# clear ikev2 statistics
```

- Enter the **show ipsec statistics** command to display IPsec statistics.

```
device# show ipsec statistics

                      IPSecurity Statistics
ipsecEspCurrentInboundSAs 1          ipsecEspTotalInboundSAs:  1
ipsecEspCurrentOutboundSA 1          ipsecEspTotalOutboundSAs: 1
                   IPSecurity Packet Statistics
ipsecEspTotalInPkts:      0          ipsecEspTotalInPktsDrop:  0
ipsecEspTotalOutPkts:     7
                   IPSecurity Error Statistics
ipsecAuthenticationErrors 0
ipsecReplayErrors:        0          ipsecPolicyErrors:        0
ipsecOtherReceiveErrors:  0          ipsecSendErrors:          0
ipsecUnknownSpiErrors:    0
```

- Enter the **clear ipsec statistics** command to reset or clear IPsec statistics.

```
device# clear ipsec statistics
```

- Enter the **show statistics tunnel** command to display tunnel statistics.

```
device# show statistics tunnel

IP GRE Tunnels
     Tunnel Status  Packet Received  Packet Sent  KA recv  KA sent

IPSEC Tunnels
     Tunnel Status  Packet Received  Packet Sent  Bytes Received  Bytes Sent
  9  down/down      0                0            0               0
 10  up/up          50               16442474     7300            9372173444
```

- Enter the **clear statistics tunnel** command to reset or clear tunnel statistics. The following example shows how to clear statistics for tunnel 10.

```
device# clear statistics tunnel 10
```

Configuration example for an IPv4 IPsec tunnel using manual PKI

Public Key Infrastructure (PKI) can be used for Internet Key Exchange version 2 (IKEv2) authentication when the certification authority (CA) server is offline by manually configuring PKI.

FIGURE 5 IPv4 IPsec tunnel

First, download the certificate files and cryptographic keys (in Privacy Enhanced Mail (PEM) format) from the CA server to flash memory on Router1 and Router2 by using Trivial File Transfer Protocol (TFTP):

Router1

- router1.crt—X509v3 certificate for Router1
- CA31.crt—Intermediate Level 3 Sub CA X509v3 certificate
- CA21.crt—Intermediate Level 2 Sub CA X509v3 certificate
- CA1.crt—Root CA X509v3 certificate
- router1_eckey.pem—Router1 keys

Router2

- router2.crt—X509v3 certificate for Router1
- CA31.crt—Intermediate Level 3 Sub CA X509v3 certificate
- CA21.crt—Intermediate Level 2 Sub CA X509v3 certificate
- CA1.crt—Root CA X509v3 certificate
- router2_eckey.pem—Router2 keys

Then, manually configure PKI by completing the following tasks on the routers at each end of the IPv4 IPsec tunnel:

- Configuring a PKI entity
- Configuring the PKI trustpoint
- Importing the CA certificates from flash memory
- Authenticating the PKI trustpoint
- Configuring the ethernet interface 1/1 with an IPv4 address.
- Configuring an IKEv2 authentication proposal that specifies using the PKI trustpoint for IKEv2 authentication
- Configuring an IKEv2 profile for the tunnel with local and remote identifiers and corresponding match conditions.

- Creating an IPsec profile that uses the IKEv2 profile.
- Configuring the tunnel and setting the mode of the tunnel to ipsec.
- Configuring an IPsec protection profile and the IPv4 address for the tunnel interface.

The following examples show how to configure Router1 and Router2.

Router1

```
Router1(config)# pki entity router1
Router1(config-pki-entity-router1)# common-name ROUTER1
Router1(config-pki-entity-router1)# org-unit-name Routing
Router1(config-pki-entity-router1)# org-name example_company
Router1(config-pki-entity-router1)# state-name CA
Router1(config-pki-entity-router1)# country-name US
Router1(config-pki-entity-router1)# email-id ROUTER1@example_company.com
Router1(config-pki-entity-router1)# location SJ
Router1(config-pki-entity-router1)# exit

Router1(config)# pki trustpoint CA1
Router1(config-pki-trustpoint-CA1)# pki-entity router1
Router1(config-pki-trustpoint-CA1)# eckeypair key-label router1
Router1(config-pki-trustpoint-CA1)# fingerprint e1:9f:e9:13:67:65:e9:73:bf:5d:04:36:0f:b1:20:96:83:23:c9:66
!Note: the configured fingerprint must match the fingerprint contained in the CA1 certificate
Router1(config-pki-trustpoint-CA1)# end

Router1# pki import CA1 pem url flash: router1.crt
Router1# pki import CA1 pem url flash: CA31.crt
Router1# pki import CA1 pem url flash: CA21.crt
Router1# pki import CA1 pem url flash: CA1.crt
Router1# pki import key ec router1 pem url flash: router1_eckey.pem

Router1# configure terminal
Router1(config)# pki authenticate CA1

Router1(config)# ikev2 auth-proposal ipsectun100
Router1(config-ike-auth-proposal-ipsectun100)# method remote ecdsa384
Router1(config-ike-auth-proposal-ipsectun100)# method local ecdsa384
Router1(config-ike-auth-proposal-ipsectun100)# pki-trustpoint CA1 sign
Router1(config-ike-auth-proposal-ipsectun100)# pki-trustpoint CA1 verify

Router1(config)# interface ethernet 1/1
Router1(config-if-e10000-1/1)# enable

Router1(config)# ikev2 profile ipsectun100
Router1(config-ike-profile-ipsectun100)# authentication ipsectun100
Router1(config-ike-profile-ipsectun100)# local-identifier address 192.168.10.41
Router1(config-ike-profile-ipsectun100)# remote-identifier address 192.168.10.43
Router1(config-ike-profile-ipsectun100)# match-identity local address 192.168.10.41
Router1(config-ike-profile-ipsectun100)# match-identity remote address 192.168.10.43
Router1(config-ike-profile-ipsectun100)# exit

Router1(config)# ipsec profile ipsectun100
Router1(config-ipsec-profile-ipsectun100)# ike-profile ipsectun100
Router1(config-ipsec-profile-ipsectun100)# exit

Router1(config)# interface tunnel 100
Router1(config-tnif-100)# tunnel mode ipsec ipv4
Router1(config-tnif-100)# tunnel protection ipsec profile ipsectun100
Router1(config-tnif-100)# tunnel source 192.168.10.41
Router1(config-tnif-100)# tunnel destination 192.168.10.43
Router1(config-tnif-100)# ip address 10.41.43.41/24
Router1(config-tnif-100)# end
```

Router2

```
Router2(config)# pki entity router2
Router2(config-pki-entity-router1)# common-name ROUTER2
Router2(config-pki-entity-router1)# org-unit-name Routing
Router2(config-pki-entity-router1)# org-name example_company
Router2(config-pki-entity-router1)# state-name CA
Router2(config-pki-entity-router1)# country-name US
Router2(config-pki-entity-router1)# email-id ROUTER2@example_company.com
Router2(config-pki-entity-router1)# location SJ
Router2(config-pki-entity-router1)# exit

Router2(config)# pki trustpoint CA1
Router2(config-pki-trustpoint-CA1)# pki-entity router2
Router2(config-pki-trustpoint-CA1)# eckeypair key-label router2
Router2(config-pki-trustpoint-CA1)# fingerprint e1:9f:e9:13:67:65:e9:73:bf:5d:04:36:0f:b1:20:96:83:23:c9:66
!Note: the fingerprint configured above must match the fingerprint contained in the CA1 certificate
Router2(config-pki-trustpoint-CA1)# end

Router2# pki import CA1 pem url flash: router2.crt
Router2# pki import CA1 pem url flash: CA31.crt
Router2# pki import CA1 pem url flash: CA21.crt
Router2# pki import CA1 pem url flash: CA1.crt
Router2# pki import key ec router1 pem url flash: router1_eckey.pem

Router2# configure terminal
Router2(config)# pki authenticate CA1

Router2(config)# ikev2 auth-proposal ipsectun100
Router2(config-ike-auth-proposal-ipsectun100)# method remote ecdsa384
Router2(config-ike-auth-proposal-ipsectun100)# method local ecdsa384
Router2(config-ike-auth-proposal-ipsectun100)# pki-trustpoint CA1 sign
Router2(config-ike-auth-proposal-ipsectun100)# pki-trustpoint CA1 verify

Router2(config)# interface ethernet 1/1
Router2(config-if-e10000-1/1)# enable

Router2(config)# ikev2 profile ipsectun100
Router2(config-ike-profile-ipsectun100)# authentication ipsectun100
Router2(config-ike-profile-ipsectun100)# local-identifier address 192.168.10.43
Router2(config-ike-profile-ipsectun100)# remote-identifier address 192.168.10.41
Router2(config-ike-profile-ipsectun100)# match-identity local address 192.168.10.43
Router2(config-ike-profile-ipsectun100)# match-identity remote address 192.168.10.41
Router2(config-ike-profile-ipsectun100)# exit

Router2(config)# ipsec profile ipsectun100
Router2(config-ipsec-profile-ipsectun100)# ike-profile ipsectun100
Router2(config-ipsec-profile-ipsectun100)# exit

Router2(config)# interface tunnel 100
Router2(config-tnif-100)# tunnel mode ipsec ipv4
Router2(config-tnif-100)# tunnel protection ipsec profile ipsectun100
Router2(config-tnif-100)# tunnel source 192.168.10.43
Router2(config-tnif-100)# tunnel destination 192.168.10.41
Router2(config-tnif-100)# ip address 10.41.43.43/24
Router2(config-tnif-100)# end
```

Configuration example for an IPv4 IPsec tunnel using dynamic PKI

Dynamic Public Key Infrastructure (PKI) can be used for Internet Key Exchange version 2 (IKEv2) authentication when the certification authority (CA) server is online.

FIGURE 6 IPv4 IPsec tunnel

Dynamic PKI is set up by completing the following tasks on the routers at each end of the IPv4 IPsec tunnel:

- Configuring a PKI entity
- Configuring a PKI enrollment profile that includes:
  - The URL of the Simple Certificate Enrollment Protocol (SCEP) or CA server to use for authentication
  - The URL of the SCEP or CA server to use for enrollment
  - The password from the SCEP server that is used in the certificate request from the local router during enrollment
- Generating Elliptic Curve Digital Signature Algorithm (ECDSA) 384 key pairs used to obtain the local router certificate
- Configuring the PKI trustpoint
- Importing the CA certificates from flash memory
- Authenticating the PKI trustpoint
- Configuring the ethernet interface 1/1 with an IPv4 address.
- Configuring an IKEv2 authentication proposal that specifies using the PKI trustpoint for IKEv2 authentication
- Configuring an IKEv2 profile for the tunnel with local and remote identifiers and corresponding match conditions.
- Creating an IPsec profile that uses the IKEv2 profile.
- Configuring the tunnel and setting the mode of the tunnel to ipsec.
- Configuring an IPsec protection profile and the IPv4 address for the tunnel interface.

The following examples show how to configure Router1 and Router2.

Router1

```
Router1(config)# pki entity router1
Router1(config-pki-entity-router1)# common-name ROUTER1
Router1(config-pki-entity-router1)# org-unit-name Routing
Router1(config-pki-entity-router1)# org-name example_company
Router1(config-pki-entity-router1)# state-name CA
Router1(config-pki-entity-router1)# country-name US
Router1(config-pki-entity-router1)# email-id ROUTER1@example_company.com
Router1(config-pki-entity-router1)# location SJ
Router1(config-pki-entity-router1)# exit
```

```
Router1(config)# ip host example.com 192.0.2.1
Router1(config)# pki profile-enrollment Router1
Router1(config-pki-profile-enrollment-Router1)# authentication-url http://example.com/CertSrv/mscep/
mscep.dll
Router1(config-pki-profile-enrollment-Router1)# authentication-command GET/certs/cacert.der
Router1(config-pki-profile-enrollment-Router1)# enrollment-url http://example.com/CertSrv/mscep/
mscep.dll
Router1(config-pki-profile-enrollment-Router1)# password 9218B7FCFE5E8DBA
Router1(config-pki-profile-enrollment-Router1)# exit

Router1(config)# crypto key generate ec label label-1

Router1(config)# pki trustpoint R1
Router1(config-pki-trustpoint-R1)# enrollment profile Router1
Router1(config-pki-trustpoint-R1)# pki-entity router1
Router1(config-pki-trustpoint-R1)# eckeypair key-label label-1
Router1(config-pki-trustpoint-R1)# fingerprint e1:9f:e9:13:67:65:e9:73:bf:5d:04:36:0f:b1:20:96:83:23:c9:66
! Note: the fingerprint should be obtained from the root CA certificate of the SCEP or CA server
Router1(config-pki-trustpoint-R1)# end

Router1# configure terminal
Router1(config)# pki authenticate R1

Router1(config)# ikev2 auth-proposal ipsectun100
Router1(config-ike-auth-proposal-ipsectun100)# method remote ecdsa384
Router1(config-ike-auth-proposal-ipsectun100)# method local ecdsa384
Router1(config-ike-auth-proposal-ipsectun100)# pki-trustpoint R1 sign
Router1(config-ike-auth-proposal-ipsectun100)# pki-trustpoint R1 verify

Router1(config)# interface ethernet 1/1
Router1(config-if-e10000-1/1)# enable

Router1(config)# ikev2 profile ipsectun100
Router1(config-ike-profile-ipsectun100)# authentication ipsectun100
Router1(config-ike-profile-ipsectun100)# local-identifier address 192.168.10.41
Router1(config-ike-profile-ipsectun100)# remote-identifier address 192.168.10.43
Router1(config-ike-profile-ipsectun100)# match-identity local address 192.168.10.41
Router1(config-ike-profile-ipsectun100)# match-identity remote address 192.168.10.43
Router1(config-ike-profile-ipsectun100)# exit

Router1(config)# ipsec profile ipsectun100
Router1(config-ipsec-profile-ipsectun100)# ike-profile ipsectun100
Router1(config-ipsec-profile-ipsectun100)# exit

Router1(config)# interface tunnel 100
Router1(config-tnif-100)# tunnel mode ipsec ipv4
Router1(config-tnif-100)# tunnel protection ipsec profile ipsectun100
Router1(config-tnif-100)# tunnel source 192.168.10.41
Router1(config-tnif-100)# tunnel destination 192.168.10.43
Router1(config-tnif-100)# ip address 10.41.43.41/24
Router1(config-tnif-100)# end
```

Router2

```
Router2(config)# pki entity router2
Router2(config-pki-entity-router1)# common-name ROUTER2
Router2(config-pki-entity-router1)# org-unit-name Routing
Router2(config-pki-entity-router1)# org-name example_company
Router2(config-pki-entity-router1)# state-name CA
Router2(config-pki-entity-router1)# country-name US
Router2(config-pki-entity-router1)# email-id ROUTER2@example_company.com
Router2(config-pki-entity-router1)# location SJ
Router2(config-pki-entity-router1)# exit

Router2(config)# ip host example.com 192.0.2.1
Router2(config)# pki profile-enrollment Router2
Router2(config-pki-profile-enrollment-Router1)# authentication-url http://example.com/CertSrv/mscep/
mscep.dll
```

```
Router2(config-pki-profile-enrollment-Router1)# authentication-command GET/certs/cacert.der
Router2(config-pki-profile-enrollment-Router1)# enrollment-url http://example.com/CertSrv/mscep/
mscep.dll
Router2(config-pki-profile-enrollment-Router1)# password 9218B7FCFE5E8DBA
Router2(config-pki-profile-enrollment-Router1)# exit

Router2(config)# crypto key generate ec label label-1

Router2(config)# pki trustpoint R1
Router2(config-pki-trustpoint-R1)# enrollment profile Router2
Router2(config-pki-trustpoint-R1)# pki-entity router2
Router2(config-pki-trustpoint-R1)# eckeypair key-label router2
Router2(config-pki-trustpoint-R1)# fingerprint e1:9f:e9:13:67:65:e9:73:bf:5d:04:36:0f:b1:20:96:83:23:c9:66
! Note: the fingerprint should be obtained from the root CA certificate of the SCEP or CA server
Router2(config-pki-trustpoint-R1)# end

Router2# pki import CA1 pem url flash: router2.crt
Router2# pki import CA1 pem url flash: CA31.crt
Router2# pki import CA1 pem url flash: CA21.crt
Router2# pki import CA1 pem url flash: CA1.crt
Router2# pki import key ec router1 pem url flash: router1_eckey.pem

Router2# configure terminal
Router2(config)# pki authenticate CA1

Router2(config)# ikev2 auth-proposal ipsectun100
Router2(config-ike-auth-proposal-ipsectun100)# method remote ecdsa384
Router2(config-ike-auth-proposal-ipsectun100)# method local ecdsa384
Router2(config-ike-auth-proposal-ipsectun100)# pki-trustpoint CA1 sign
Router2(config-ike-auth-proposal-ipsectun100)# pki-trustpoint CA1 verify

Router2(config)# interface ethernet 1/1
Router2(config-if-e10000-1/1)# enable

Router2(config)# ikev2 profile ipsectun100
Router2(config-ike-profile-ipsectun100)# authentciation ipsectun100
Router2(config-ike-profile-ipsectun100)# local-identifier address 192.168.10.43
Router2(config-ike-profile-ipsectun100)# remote-identifier address 192.168.10.41
Router2(config-ike-profile-ipsectun100)# match-identity local address 192.168.10.43
Router2(config-ike-profile-ipsectun100)# match-identity remote address 192.168.10.41
Router2(config-ike-profile-ipsectun100)# exit

Router2(config)# ipsec profile ipsectun100
Router2(config-ipsec-profile-ipsectun100)# ike-profile ipsectun100
Router2(config-ipsec-profile-ipsectun100)# exit

Router2(config)# interface tunnel 100
Router2(config-tnif-100)# tunnel mode ipsec ipv4
Router2(config-tnif-100)# tunnel protection ipsec profile ipsectun100
Router2(config-tnif-100)# tunnel source 192.168.10.43
Router2(config-tnif-100)# tunnel destination 192.168.10.41
Router2(config-tnif-100)# ip address 10.41.43.43/24
Router2(config-tnif-100)# end
```

Minimum configuration example for an IPv4 IPsec tunnel

An IPv4 IPsec tunnel can be set up by performing a minimum number of configuration tasks.

FIGURE 7 IPv4 IPsec tunnel

The preceding figure shows an IPv4 IPsec tunnel topology. The following examples show how to configure Router1 and Router2 and include:

- Configuring the ethernet interface 1/1 with an IPv4 address.
- Configuring an IKEv2 profile for the tunnel with local and remote identifiers and corresponding match conditions.
- Creating an IPsec profile that uses the IKEv2 profile.
- Configuring the tunnel and setting the mode of the tunnel to ipsec.
- Configuring an IPsec protection profile and the IPv4 address for the tunnel interface.

Router1

```
Router1(config)# interface ethernet 1/1
Router1(config-if-e10000-1/1)# enable

Router1(config)# ikev2 profile ipsectun100
Router1(config-ike-profile-ipsectun100)# local-identifier address 192.168.10.41
Router1(config-ike-profile-ipsectun100)# remote-identifier address 192.168.10.43
Router1(config-ike-profile-ipsectun100)# match-identity local address 192.168.10.41
Router1(config-ike-profile-ipsectun100)# match-identity remote address 192.168.10.43
Router1(config-ike-profile-ipsectun100)# exit

Router1(config)# ipsec profile ipsectun100
Router1(config-ipsec-profile-ipsectun100)# ike-profile ipsectun100
Router1(config-ipsec-profile-ipsectun100)# exit

Router1(config)# interface tunnel 100
Router1(config-tnif-100)# tunnel mode ipsec ipv4
Router1(config-tnif-100)# tunnel protection ipsec profile ipsectun100
Router1(config-tnif-100)# tunnel source 192.168.10.41
Router1(config-tnif-100)# tunnel destination 192.168.10.43
Router1(config-tnif-100)# ip address 10.41.43.41/24
Router1(config-tnif-100)# end
```

Router2

```
Router2(config)# interface ethernet 1/1
Router2(config-if-e10000-1/1)# enable
```

```
Router2(config)# ikev2 profile ipsectun100
Router2(config-ike-profile-ipsectun100)# local-identifier address 192.168.10.43
Router2(config-ike-profile-ipsectun100)# remote-identifier address 192.168.10.41
Router2(config-ike-profile-ipsectun100)# match-identity local address 192.168.10.43
Router2(config-ike-profile-ipsectun100)# match-identity remote address 192.168.10.41
Router2(config-ike-profile-ipsectun100)# exit

Router2(config)# ipsec profile ipsectun100
Router2(config-ipsec-profile-ipsectun100)# ike-profile ipsectun100
Router2(config-ipsec-profile-ipsectun100)# exit

Router2(config)# interface tunnel 100
Router2(config-tnif-100)# tunnel mode ipsec ipv4
Router2(config-tnif-100)# tunnel protection ipsec profile ipsectun100
Router2(config-tnif-100)# tunnel source 192.168.10.43
Router2(config-tnif-100)# tunnel destination 192.168.10.41
Router2(config-tnif-100)# ip address 10.41.43.43/24
Router2(config-tnif-100)# end
```

Minimum configuration example for an IPv6 IPsec tunnel

An IPv6 IPsec tunnel can be set up by performing a minimum number of configuration tasks.

FIGURE 8 IPv6 IPsec tunnel

The preceding figure shows an IPv6 IPsec tunnel topology. The following examples show how to configure Router1 and Router2 and include:

- Configuring the ethernet interface 1/1 with an IPv6 address.
- Configuring an IKEv2 profile for the tunnel with local and remote identifiers and corresponding match conditions.
- Creating an IPsec profile that uses the IKEv2 profile.
- Configuring the tunnel and setting the mode of the tunnel to ipsec.
- Configuring an IPsec protection profile and the IPv6 address for the tunnel interface.

Router1

```
Router1(config)# interface ethernet 1/1
Router1(config-if-e10000-1/1)# enable

Router1(config)# ikev2 profile ipsectun150
Router1(config-ike-profile-ipsectun150)# local-identifier address 2001:db8:41:43:150::41
Router1(config-ike-profile-ipsectun150)# remote-identifier address 2001:db8:41:43:150::43
Router1(config-ike-profile-ipsectun150)# match-identity local address 2001:db8:41:43:150::41
Router1(config-ike-profile-ipsectun150)# match-identity remote address 2001:db8:41:43:150::43
```

```
Router1(config-ike-profile-ipsectun150)# exit

Router1(config)# ipsec profile ipsectun150
Router1(config-ipsec-profile-ipsectun150)# ike-profile ipsectun150
Router1(config-ipsec-profile-ipsectun150)# exit

Router1(config)# interface tunnel 150
Router1(config-tnif-150)# tunnel mode ipsec ipv6
Router1(config-tnif-150)# tunnel protection ipsec profile ipsectun150
Router1(config-tnif-150)# tunnel source 2001:db8:41:43:150::41
Router1(config-tnif-150)# tunnel destination 2001:db8:41:43:150::43
Router1(config-tnif-150)# ipv6 address 2001:db8:150:41:43::41/64
Router1(config-tnif-150)# end
```

Router2

```
Router2(config)# interface ethernet 1/1
Router2(config-if-e10000-1/1)# enable

Router2(config)# ikev2 profile ipsectun150
Router2(config-ike-profile-ipsectun150)# local-identifier address 2001:db8:41:43:150::43
Router2(config-ike-profile-ipsectun150)# remote-identifier address 2001:db8:41:43:150::41
Router2(config-ike-profile-ipsectun150)# match-identity local address 2001:db8:41:43:150::43
Router2(config-ike-profile-ipsectun150)# match-identity remote address 2001:db8:41:43:150::41
Router2(config-ike-profile-ipsectun150)# exit

Router2(config)# ipsec profile ipsectun150
Router2(config-ipsec-profile-ipsectun150)# ike-profile ipsectun150
Router2(config-ipsec-profile-ipsectun150)# exit

Router2(config)# interface tunnel 150
Router2(config-tnif-150)# tunnel mode ipsec ipv6
Router2(config-tnif-150)# tunnel protection ipsec profile ipsectun150
Router2(config-tnif-150)# tunnel source 2001:db8:41:43:150::43
Router2(config-tnif-150)# tunnel destination 2001:db8:41:43:150::41
Router2(config-tnif-150)# ipv6 address 2001:db8:150:41:43::43/64
Router2(config-tnif-150)# end
```

Configuration example for an IPsec tunnel

An IPsec tunnel is configured by binding an IPsec profile to the virtual tunnel interface (VTI) at each end of the IPsec tunnel.

In the following example, RouterA and RouterB are devices at the endpoints of an IPsec tunnel. On each device, an IPsec profile (red) is created and bound to the VTI by using the **tunnel protection ipsec profile** command.

FIGURE 9 Network diagram showing IPsec tunnels
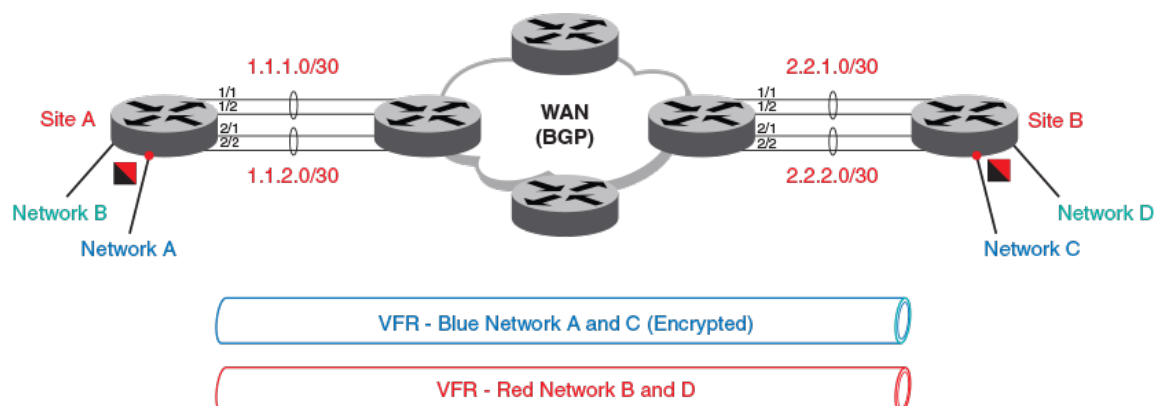
> **NOTE**
> This example configuration is for IPsec IPv4. The configuration for IPsec IPv6 is similar and uses many of the same commands.

RouterA

```
RouterA# configure terminal
RouterA(config)# ikev2 proposal red
RouterA(config-ike-proposal-red)# exit

RouterA(config)# ikev2 auth-proposal ap1
RouterA(config-ike-auth-proposal-ap1)# method remote pre-shared
RouterA(config-ike-auth-proposal-ap1)# method local pre-shared
RouterA(config-ike-auth-proposal-ap1)# pre-shared-key test
RouterA(config-ike-auth-proposal-ap1)# exit

RouterA(config)# ikev2 auth-proposal red
RouterA(config-ike-auth-proposal-red)# method remote pre-shared
RouterA(config-ike-auth-proposal-red)# method local pre-shared
RouterA(config-ike-auth-proposal-red)# pre-shared-key red
RouterA(config-ike-auth-proposal-red)# exit

RouterA(config)# ikev2 policy red
RouterA(config-ike-policy-red)# proposal red
RouterA(config-ike-policy-red)# match address-local 1.2.45.1 255.255.255.255
RouterA(config-ike-policy-red)# exit

RouterA(config)# ikev2 profile red
RouterA(config-ike-profile-red)# authentication red
RouterA(config-ike-profile-red)# local-identifier address 1.2.45.1
RouterA(config-ike-profile-red)# remote-identifier address 1.2.45.2
RouterA(config-ike-profile-red)# match-identity local address 1.2.45.1
RouterA(config-ike-profile-red)# match-identity remote address 1.2.45.2
RouterA(config-ike-profile-red)# exit

RouterA(config)# ipsec proposal red
RouterA(config-ipsec-proposal-red)# exit

RouterA(config)# ipsec profile red
RouterA(config-ipsec-profile-red)# proposal red
RouterA(config-ipsec-profile-red)# ike-profile red
RouterA(config-ipsec-profile-red)# exit

RouterA(config)# router ospf
RouterA(config-ospf-router)# area 0
RouterA(config-ospf-router)# exit

RouterA(config)# interface loopback 1
RouterA(config-lbif-1)# ip ospf area 0
RouterA(config-lbif-1)# ip address 1.1.1.1/24
RouterA(config-lbif-1)# exit

RouterA(config)# interface ethernet 4/5
RouterA(config-if-e10000-4/5)# enable
RouterA(config-if-e10000-4/5)# ip ospf area 0
RouterA(config-if-e10000-4/5)# ip address 1.2.45.1/24
RouterA(config-if-e10000-4/5)# ipv6 address 1:2:45::1/64
RouterA(config-if-e10000-4/5)# exit

RouterA(config)# interface tunnel 1
RouterA(config-tnif-1)# tunnel mode ipsec ipv4
RouterA(config-tnif-1)# tunnel mtu 1431
RouterA(config-tnif-1)# tunnel protection ipsec profile red
RouterA(config-tnif-1)# tunnel source ethernet 4/5
RouterA(config-tnif-1)# tunnel destination 1.2.45.2
RouterA(config-tnif-1)# ip address 1.1.2.1/24
RouterA(config-tnif-1)# end
```

RouterB

```
RouterB# configure terminal
RouterB(config)# ikev2 proposal red
RouterB(config-ike-proposal-red)# exit

RouterB(config)# ikev2 auth-proposal ap1
RouterB(config-ike-auth-proposal-ap1)# method remote pre-shared
RouterB(config-ike-auth-proposal-ap1)# method local pre-shared
RouterB(config-ike-auth-proposal-ap1)# pre-shared-key test
RouterB(config-ike-auth-proposal-ap1)# exit

RouterB(config)# ikev2 auth-proposal green
RouterB(config-ike-auth-proposal-red)# method remote pre-shared
RouterB(config-ike-auth-proposal-red)# method local pre-shared
RouterB(config-ike-auth-proposal-red)# pre-shared-key green
RouterB(config-ike-auth-proposal-red)# exit

RouterB(config)# ikev2 policy red
RouterB(config-ike-policy-red)# proposal red
RouterB(config-ike-policy-red)# match address-local 2.2.2.2 255.255.255.0
RouterB(config-ike-policy-red)# exit

RouterB(config)# ikev2 profile red
RouterB(config-ike-profile-red)# authentication red
RouterB(config-ike-profile-red)# responder only
RouterB(config-ike-profile-red)# local-identifier address 1.2.45.2
RouterB(config-ike-profile-red)# remote-identifier address 1.2.45.1
RouterB(config-ike-profile-red)# match-identity local address 1.2.45.2
RouterB(config-ike-profile-red)# match-identity remote address 1.2.45.1
RouterB(config-ike-profile-red)# exit

RouterB(config)# ipsec proposal red
RouterB(config-ipsec-proposal-red)# exit

RouterB(config)# ipsec profile red
RouterB(config-ipsec-profile-red)# proposal red
RouterB(config-ipsec-profile-red)# ike-profile red
RouterB(config-ipsec-profile-red)# exit

RouterB(config)# interface ethernet 3/5
RouterB(config-if-e10000-4/5)# enable
RouterB(config-if-e10000-4/5)# ip ospf area 0
RouterB(config-if-e10000-4/5)# ip address 1.2.45.2/24
RouterB(config-if-e10000-4/5)# exit

RouterB(config)# interface tunnel 1
RouterB(config-tnif-1)# tunnel mode ipsec ipv4
RouterB(config-tnif-1)# tunnel source ethernet 3/5
RouterB(config-tnif-1)# tunnel destination 1.2.45.1
RouterB(config-tnif-1)# ip address 1.1.2.2/24
RouterB(config-tnif-1)# exit

RouterB(config)# interface tunnel 1
RouterB(config-tnif-1)# tunnel mode ipsec ipv4
RouterB(config-tnif-1)# protection ipsec profile red
RouterB(config-tnif-1)# tunnel source ethernet 3/5
RouterB(config-tnif-1)# tunnel destination 1.2.45.1
RouterB(config-tnif-1)# ip address 1.1.2.2/24
RouterB(config-tnif-1)# end
```

The following configuration example shows how to run OSPF and Border Gateway Protocol (BGP) over the IPsec tunnel.

RouterA

```
RouterA# configure terminal
RouterA(config)# router ospf
RouterA(config-ospf-router)# area 0
```

```
RouterA(config-ospf-router)# exit

RouterA(config)# interface tunnel 1
RouterA(config-tnif-1)# tunnel mode ipsec ipv4
RouterA(config-tnif-1)# tunnel mtu 1476
RouterA(config-tnif-1)# tunnel protection ipsec profile red
RouterA(config-tnif-1)# tunnel source ethernet 4/5
RouterA(config-tnif-1)# tunnel destination 1.2.45.2
RouterA(config-tnif-1)# ip address 1.1.2.1/24
RouterA(config-tnif-1)# ip ospf area 0              !Enabling ospf on the VTI Interface
RouterA(config-tnif-1)# exit

!Configuring BGP over the IPSEC tunnel

RouterA(config)# router bgp
RouterA(config-bgp)# local-as 100
RouterA(config-bgp)# neighbor 1.1.2.2 remote-as 100

RouterA(config-bgp)# address-family ipv4 unicast
RouterA(config-bgp)# exit-address-family

RouterA(config-bgp)# address-family ipv4 multicast
RouterA(config-bgp)# exit-address-family

RouterA(config-bgp)# address-family ipv6 unicast
RouterA(config-bgp)# exit-address-family

RouterA(config-bgp)# address-family ipv6 multicast
RouterA(config-bgp)# exit-address-family

RouterA(config-bgp)# address-family vpnv4 unicast
RouterA(config-bgp)# exit-address-family

RouterA(config-bgp)# address-family vpnv6 unicast
RouterA(config-bgp)# exit-address-family

RouterA(config-bgp)# address-family l2vpn vpls
RouterA(config-bgp)# exit-address-family
RouterA(config-bgp)# end
```

RouterB

```
RouterB# configure terminal
RouterB(config)# router ospf
RouterB(config-ospf-router)# area 0
RouterB(config-ospf-router)# exit

RouterB(config)# interface tunnel 1
RouterB(config-tnif-1)# tunnel mode ipsec ipv4
RouterB(config-tnif-1)# tunnel source ethernet 3/5
RouterB(config-tnif-1)# tunnel destination 1.2.45.1
RouterB(config-tnif-1)# ip address 1.1.2.2/24
RouterB(config-tnif-1)# ip ospf area 0
RouterB(config-tnif-1)# exit

!Configuring BGP over the IPSEC tunnel

RouterB(config)# router bgp
RouterB(config-bgp)# local-as 100
RouterB(config-bgp)# neighbor 1.1.2.1 remote-as 100

RouterB(config-bgp)# address-family ipv4 unicast
RouterB(config-bgp)# exit-address-family

RouterB(config-bgp)# address-family ipv4 multicast
RouterB(config-bgp)# exit-address-family

RouterB(config-bgp)# address-family ipv6 unicast
RouterB(config-bgp)# exit-address-family
```

```
RouterB(config-bgp)# address-family ipv6 multicast
RouterB(config-bgp)# exit-address-family

RouterB(config-bgp)# address-family vpnv4 unicast
RouterB(config-bgp)# exit-address-family

RouterB(config-bgp)# address-family vpnv6 unicast
RouterB(config-bgp)# exit-address-family

RouterB(config-bgp)# address-family l2vpn vpls
RouterB(config-bgp)# exit-address-family
RouterB(config-bgp)# end
```

In the following configuration example, the IPsec tunnel source and destination addresses are in the default VRF while the IPsec tunnel address is in a different VRF; a data packet received for IP address 12.1.0.0/24 on the VRF named blue and the VRF named green is forwarded on the IPsec tunnel after IPsec encryption.

RouterA

```
RouterA# configure terminal
RouterA(config)# ip route vrf blue 1.2.45.1 255.255.0.0 next-hop-vrf default-vrf 10.1.1.2
RouterA(config)# ip route 2.2.2.1 255.255.0.0 next-hop-vrf blue 10.1.1.2

RouterA(config)# interface tunnel 1
RouterA(config-tnif-1)# vrf forwarding blue
RouterA(config-tnif-1)# tunnel mode ipsec ipv4
RouterA(config-tnif-1)# tunnel source ethernet 3/5
RouterA(config-tnif-1)# tunnel destination 1.2.45.1
RouterA(config-tnif-1)# tunnel protection ipsec profile prof-blue
RouterA(config-tnif-1)# ip address 1.1.2.2/24
RouterA(config-tnif-1)# exit

RouterA(config)# interface tunnel 1
RouterA(config-tnif-1)# tunnel mode ipsec ipv4
RouterA(config-tnif-1)# tunnel source ether 1/1
RouterA(config-tnif-1)# tunnel destination 2.2.2.1
RouterA(config-tnif-1)# tunnel vrf blue
RouterA(config-tnif-1)# tunnel protection ipsec profile prof-green
RouterA(config-tnif-1)# ip address 10.1.1.1/24
RouterA(config-tnif-1)# end
```

Configuration example for running PIM over IPsec tunnels

To secure multicast traffic, Protocol Independent Multicast (PIM) may be configured over an IPsec virtual tunnel interface (VTI).

The following example shows how to configure PIM to run over IPsec tunnels.

This example shows how to configure two IPsec tunnels between the tunnel endpoints on RouterA and RouterB; tunnel1 is for the blue VRF and tunnel2 is for the default VRF. An IPsec profile named prof-blue protects tunnel1 and an IPsec profile named prof-green protects tunnel2.

RouterA

```
RouterA# configure terminal
RouterA(config)# router pim vrf blue

RouterA(config)# Interface tunnel1
```

```
RouterA(config-tnif-1)# vrf forwarding blue
RouterA(config-tnif-1)# tunnel source ether 1/1
RouterA(config-tnif-1)# tunnel destination 2.2.2.1
RouterA(config-tnif-1)# tunnel mode ipsec ipv4
RouterA(config-tnif-1)# tunnel protection ipsec profile prof-blue
RouterA(config-tnif-1)# ip address 10.1.1.1/24
RouterA(config-tnif-1)# ip pim-sparse !Enabling PIM-SM on the VTI Interface
RouterA(config-tnif-1)# end

RouterA# Interface tunnel2
RouterA(config-tnif-2)# tunnel source ether 1/1
RouterA(config-tnif-2)# tunnel destination 2.2.2.1
RouterA(config-tnif-2)# tunnel mode ipsec ipv4
RouterA(config-tnif-2)# tunnel protection ipsec profile prof-green
RouterA(config-tnif-2)# ip address 10.1.1.1/24
RouterA(config-tnif-2)# ip pim-sparse !Enabling PIM-SM on the VTI Interface
```

RouterB

```
RouterB# configure terminal
RouterB(config)# router pim vrf blue
RouterB(config)# interface tunnel1
RouterB(config-tnif-1)# vrf forwarding blue
RouterB(config-tnif-1)# tunnel source ether 1/1
RouterB(config-tnif-1)# tunnel destination 1.1.1.1
RouterB(config-tnif-1)# tunnel mode ipsec ipv4
RouterB(config-tnif-1)# tunnel protection ipsec profile prof-blue
RouterB(config-tnif-2)# ip address 10.1.1.2/24
RouterB(config-tnif-1)# ip pim-sparse !Enabling PIM-SM on the VTI interface
RouterB(config-tnif-1)# end

RouterB# Interface tunnel2
RouterB(config-tnif-2)# tunnel source ether 1/1
RouterB(config-tnif-2)# tunnel destination 1.1.1.1
RouterB(config-tnif-2)# tunnel mode ipsec ipv4
RouterB(config-tnif-2)# tunnel protection ipsec profile prof-green
RouterB(config-tnif-2)# ip address 10.1.1.2/24
RouterB(config-tnif-1)# ip pim-sparse !Enabling PIM-SM on the VTI interface
```

IP Source Guard

IP source guard

IP Source Guard permits traffic from only valid source IP addresses. IP source guard is used on client ports to prevent IP source address spoofing.

IP source guard is used together with DHCP snooping and Dynamic ARP Inspection on untrusted ports. When IP Source Guard is first enabled, only DHCP packets are allowed and all other IP traffic is blocked. IP Source Guard uses IP or MAC bindings inside the ARP Inspection table to detect a valid IP address. When the system learns a valid IP address on the port, the client port then allows IP traffic to enter. If the source IP address of a packet does not match any of the IP addresses inside the ARP Inspection table, the packet is dropped. Only traffic with valid source IP addresses are permitted. The system learns of a valid IP address from ARP.

> **NOTE**
> When IP source guard is enabled on a port it must have the same configuration as the primary port, otherwise it will not be implemented as IP source guarded.

Enabling IP source inspection on a VLAN

IP Source Guard configuration is enabled on ports per vlan. When IP Source Guard is enabled on a vlan, by default all ports inside the vlan are set as "unguarded". You can selectively turn on which ports inside the vlan to be set as "guarded". Initially, when the vlan port is IP Source Guarded, only DHCP packets are allowed to get through. However, as IP or MAC binding is learned from DHCP snooping, or if it is manually configured, only packets with valid source IP address are allowed through.

There are two modes for IP Source Guard; strict mode and loose mode. You can configure either strict or loose mode during IP Source Guard vlan configuration. In a strict mode, the IP source address is bound to a particular port and vlan. Only packets with an IP address coming from a particular vlan port is considered valid. If the same source IP address is coming from a different port, then it is considered an attack and is dropped. The strict mode provides more security, but it does not allow for a layer 2 occurrence in a vlan. In a loose mode, the IP source address is bound to a vlan. Only packets with IP source addresses that come from ports within the vlan are considered valid.

To enable IP Source Inspection for a VLAN or a range of VLANs, enter the following command.

```
device(config)# ip source-inspection vlan 2
```

Syntax: [no] ip source-inspection vlan *vlan_number* [to *vlan_number*] [strict]

The source IP addresses for VLAN IP packets are inspected for any port when IP Source Guard is enabled.

The *vlan_number* variable specifies the ID of a configured VLAN.

If the strict option is enabled, then valid IP source address is bound to a particular source port. This configuration can be learned from a DHCP reply, or manually configured.

> **NOTE**
> The strict mode requires DHCP relay-information insertion to be turned on.

Displaying IP source inspection status and ports

To display the IP Source Guard status for a VLAN, and the guarded or unguarded ports in the VLAN, enter the following command.

```
device(config)#sh ip source-inspection vlan 10
IP Source Inspection configuration for VLAN 10:
Inspection mode: loose
un-guarded ports:
 ethe 1/4 ethe 1/18
guarded ports:
 ethe 1/20
```

The **show ip source-inspection vlan** command displays IP Source inspection configuration for VLAN 10 in loose mode.

Syntax: **show ip source-inspection** [**vlan** *vlan_id*]

The *vlan_id* variable specifies the ID of a configured vlan.

> NOTE
> This command is also available for debugging purposes on the Interface Module.

Enabling IP source guard

IP Source Guard is enabled on a port in interface configuration mode.

DHCP Snooping should be configured before you enable IP Source Guard.

When IP source guard is enabled on a port it must have the same configuration as the primary port, otherwise it will not be implemented as IP source-guarded.

1. Enter global configuration mode.

       device# configure terminal

2. Configure an Ethernet interface.

       device(config)# interface ethernet 2/2

3. Enable the interface as an IP source-guarded port.

       device(config-if-e1000-2/2)# source-guard

In the following example. IP source guard is enabled on Ethernet interface 2/2.

```
device# configure terminal
device(config)# interface ethernet 2/2
device(config-if-e1000-2/2)# source-guard
```

IP source guard CAM

The Extreme device configuration uses a layer 4 ACL CAM to implement IP Source guard. When IP or MAC binding is learned or configured on an IP Source Guarded vlan-port, a layer 4 ACL CAM is programmed to allow valid source IP addresses.

When ACL is manually configured, a configuration conflict occurs with IP Source Guard, because it uses a layer 4 ACL CAM. The Extreme device gives user ACL configuration a higher priority. When both IP Source Guard and user ACL is configured, the user ACL configuration takes precedence over IP Source Guard.

IP Source Guard uses layer 4 ACL CAM to check layer 2 switched traffic. When IP Source Guard is configured, the layer 3 port check flag is turned on. When IP Source Guard is configured, all traffic from the same physical port is subject to a layer 4 ACL check.

Configuring IP source guard CAM partition

IP Source Guard creates two CAM sub-partitions. The CAM sub-partitions include IP_SOURCE_GUARD_PERMIT and IP_SOURCE_GUARD_DENY. All CAM entries that are permitted, go to the IP_SOURCE_GUARD_PERMIT sub-partition. All CAM entries that are denied, go to the IP_SOURCE_GUARD_DENY sub-partition. The **system-max ip-source-guard-cam** command allows you to control the size of both IP_SOURCE_GUARD_PERMIT and IP_SOURCE_GUARD_DENY sub-partitions.

To specify a partition size of the IP Source Guard CAM, enter the following command.

```
device(config)#system-max ip-source-guard-cam 1008
```

Syntax: [no] system-max [**ip-source-guard-cam decimal**]

By default, **no** system-max is configured.

The *decimal* variable specifies the range that is supported for configuring IP Source Guard CAM sub-partitions. The decimal range is from 0 to 131072. The default is 0.

Media Access Control Security (MACsec)

MACsec overview

Media Access Control Security (MACsec) is a Layer 2 security technology that provides point-to-point security on Ethernet links between nodes.

MACsec, defined in the IEEE 802.1AE-2006 standard, is based on symmetric cryptographic keys. MACsec Key Agreement (MKA) protocol, defined as part of the IEEE 802.1x-2010 standard, operates at Layer 2 to generate and distribute the cryptographic keys used by the MACsec functionality installed in the hardware. MACsec key-enabled security can be deployed on a point-to-point LAN between two connected devices over interfaces that share the pre-configured static key, the Connectivity Association Key (CAK).

As a hop-to-hop Layer 2 security feature, MACsec can be combined with Layer 3 security technologies such as IPsec for end-to-end data security.

> NOTE
> MACsec on MLXe devices can interoperate with MACsec on ICX devices.

The following cards support MACsec configuration:
- BR-MLX-10Gx20 20-port 1/10GbE Module
- BR-MLX-10Gx4-M-IPSEC 4-port 1/10GbE and 4-port 1GbE Module

MACsec RFCs and standards

FastIron MACsec complies with the following industry standards:
- IEEE Std 802.1X-2010: Port-Based Network Access Control
- IEEE Std 802.1AE-2006: Media Access Control (MAC) Security
- RFC 3394: Advanced Encryption Standard (AES) Key Wrap Algorithm
- RFC 5649: Advanced Encryption Standard (AES) Key Wrap with Padding Algorithm

Refer to "802.1x Port-Based Authentication" for information on other IEEE 802.1X features.

How MACsec works

MACsec capabilities prevent Layer 2 security threats, such as passive wiretapping, denial of service, intrusion, man-in-the-middle, and playback attacks.

MACsec protects communications using several configurable techniques. Data origin is authenticated and data is transported over secured channels. Frames are validated as MACsec Ethernet frames. The integrity of frame content is verified on receipt. Frame sequence is monitored using an independent replay protection counter. Invalid frames are discarded or monitored.

Data traffic carried within the MACsec frame is encrypted and decrypted using an industry-standard cipher suite.

How MACsec handles data and control traffic

All traffic is controlled on an active MACsec port; that is, data is encrypted, or its integrity is protected, or both. If a MACsec session cannot be secured, all data and control traffic is dropped.

When MACsec is active on a port, the port blocks the flow of data traffic. Data traffic is not forwarded by the port until a MACsec session is secured. If an ongoing session is torn down, traffic on the port is again blocked until a new secure session is established.

Control traffic (such as STP, LACP, or UDLD traffic) is not transmitted by an active MACsec port until a MACsec session is secured. While a session is being established, only 802.1x protocol packets are transmitted from the port. Once a secure session is established, control traffic flows normally through the port.

MACsec Key Agreement protocol

MACsec Key Agreement (MKA) protocol installed on a device relies on an IEEE 802.1X Extensible Authentication Protocol (EAP) framework to establish communication.

MACsec peers on the same LAN belong to a unique connectivity association. Members of the same connectivity association identify themselves with a shared Connectivity Association Key (CAK) and Connectivity Association Key Name (CKN). The CAK is a static key that is preconfigured on each MACsec-enabled interface. MACsec authentication is based on mutual possession and acknowledgment of the preconfigured CAK and Connectivity Association Key Name (CKN).

Each peer device establishes a single unidirectional secure channel for transmitting MACsec frames (Ethernet frames with MACsec headers that usually carry encrypted data) to its peers within the connectivity association. A connectivity association consists of two secure channels, one for inbound traffic, and one for outbound traffic. All peers within the connectivity association use the same cipher suite, Galois/Counter Mode Advanced Encryption Standard 128 (GCM-AES-128), for MACsec-authenticated security functions.

MACsec Key Agreement (MKA) protocol uses the Connectivity Association Key to derive transient session keys called Secure Association Keys (SAKs). SAKs and other MKA parameters are required to sustain communication over the secure channel and to perform encryption and other MACsec security functions. SAKs, along with other essential control information, are distributed in MKA protocol control packets, also referred to as MKPDUs.

MKA message exchange between two switches

When two MACsec peers confirm possession of a shared CAK and CKN, MKA protocol initiates key-server election.

The key-server is responsible for determining whether MACsec encryption is used and what cipher suite is used to encrypt data. The key-server is also responsible for generating Secure Association Keys (SAKs) and distributing them to the connected device. Once a SAK is successfully installed, the two devices can exchange secure data.

The following figure shows the message flow between two switches during MACsec communication.

FIGURE 10 MKA pre-shared key and key name exchange between two switches

Secure channels

Communication on each secure channel takes place as a series of transient sessions called secure associations. These sessions can only be established with a unique Secure Association Key (SAK) assigned to the session.

Secure associations expire and must be re-established after transmission of a certain number of frames, or after a peer disconnects and reconnects.

The secure association is designated by a Secure Association Identifier (SAI), formed from the Secure Channel Identifier (SCI) combined with an Association Number (AN). When a MACsec frame is received by a peer interface, the device identifies the session key by mapping to the default SCI configured and uses the key to decrypt and authenticate the received frame.

MACsec frame format

When MACsec is enabled, the hardware transforms each Ethernet frame by adding a security tag (secTAG) to the frame.

The following figure shows how the Ethernet frame is converted into a MACsec frame.

FIGURE 11 MACsec frame format
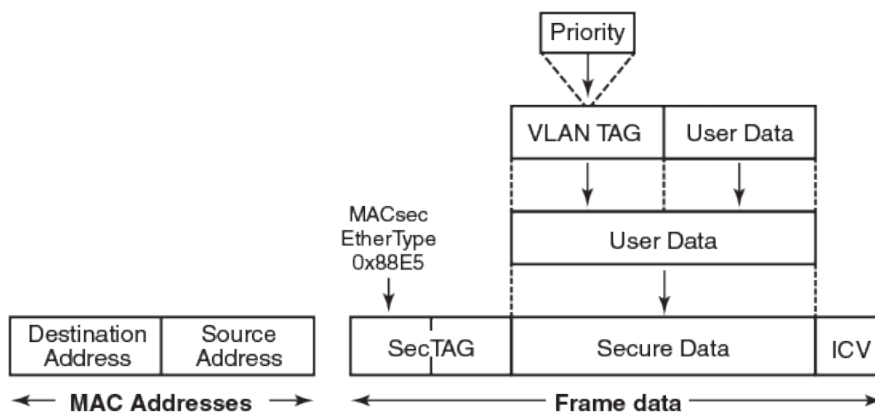
The security tag passes MACsec-related information to MACsec peers.

The following figure defines the fields in a security tag.

FIGURE 12 MACsec security tag format

Processing incoming frames

The hardware processes each MACsec frame received or transmitted based on the information in the MACsec security tag.

The device first confirms the Ethertype on incoming frames as MACsec and then processes incoming MACsec frames as illustrated in the following figure.

FIGURE 13 MACsec incoming frames
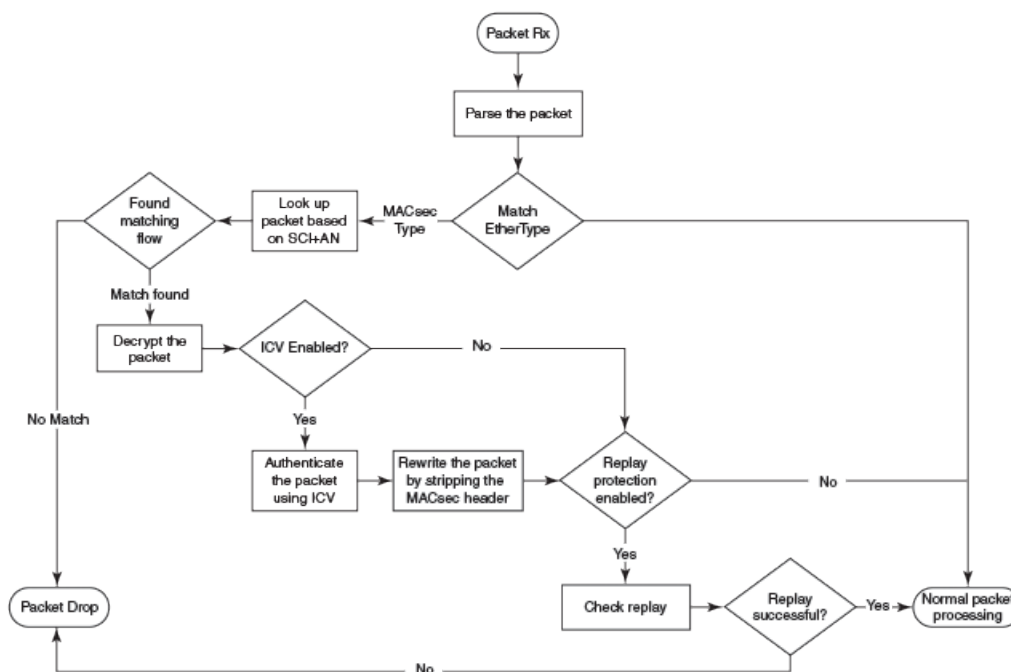
Processing outgoing frames

The device parses each outgoing frame and, if MACsec is enabled, processes the outgoing MACsec frame to apply configured MACsec options.

The following figure shows how the device applies configured MACsec options before transmitting the frames.

FIGURE 14 MACsec outgoing frames
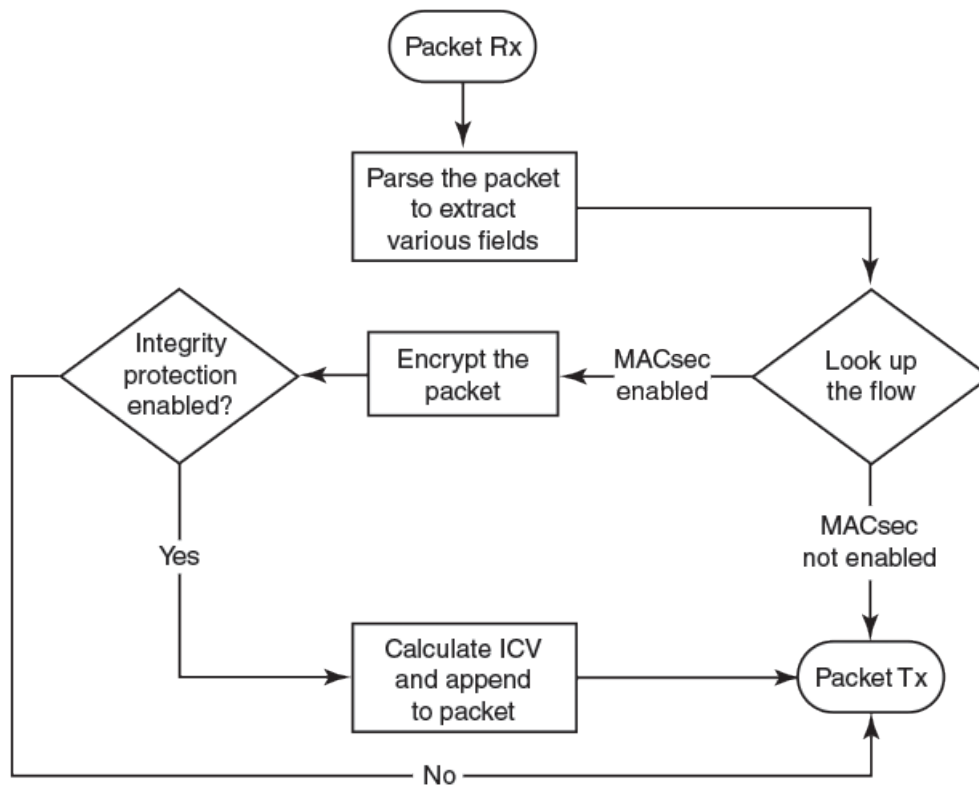
Configuring MACsec

These steps are required to configure MACsec security on a link or a group of connected ports:

1. Enter the dot1x-mka level from the global configuration level, and enable MACsec for the device.
2. Configure the MACsec Key Agreement (MKA) group.
3. Configure required parameters for the group, including frame validation, confidentiality, and replay protection.
4. Enable MKA on each participating interface.
5. Apply the configured MKA group on the participating interface.
6. Configure Connectivity Association Key (CAK) and Connectivity Association Key Name (CKN) on each interface.

Enabling MACsec and configuring group parameters

Enable MACsec globally on the device, and configure the MACsec Key Agreement (MKA) group before configuring MACsec security features for the group.

1. At the global configuration level, enter the **dot1x-mka-enable** command to enable MACsec on the device.

   ```
   device# configure terminal
   device(config)# dot1x-mka-enable
   device(config-dot1x-mka)#
   ```

   MACsec is enabled, and the device is placed at the dot1x-mka configuration level.

   > **NOTE**
   > When MKA is disabled, all the ports are brought to a down state. You must manually enable the ports again to bring the ports back up.

2. Enter the **mka-cfg-group** command followed by a group name to create a group.

   ```
   device# configure terminal
   device(config)# dot1x-mka
   device(config-dot1x-mka)# mka-cfg-group test1
   device(config-dot1x-mka-group-test1)#
   ```

   The group is created, and the device is placed at the group configuration level.

At the group configuration level, set key-server priority, and define MACsec security features to be applied to interfaces once they are assigned to the group.

Configuring MACsec key-server priority

MACsec uses a key-server to generate and distribute encryption parameters and secure key information to members of a MACsec connectivity association.

The key-server is elected by comparing key-server priority values during MACsec Key Agreement (MKA) message exchange between peer devices. The elected key-server is the peer with the lowest configured key-server priority, or with the lowest Secure Channel Identifier (SCI) in case of a tie. Key-server priority may be set to a value from 0 through 255. When no priority is configured, the device defaults to a priority of 16, which is not displayed in MACsec configuration details.

> **NOTE**
> If the key-server priority is set to 255, the device will not become the key-server.

Refer to Configuring MACsec on page 297 for an overview of enabling and configuring MACsec features.

At the dot1x-mka group configuration level, enter the **key-server-priority** command, and specify a value from 0 through 255 to define the key-server priority.

```
device# configure terminal
device(config)# dot1x-mka-enable
device(config-dot1x-mka)# mka-cfg-group test1
device(config-dot1x-mka-group-test1)# key-server-priority 20
```

In this example, the key-server priority is set to 20 for the MKA group test1.

Configure MACsec integrity and encryption for the group.

Configuring MACsec integrity and encryption

To ensure point-to-point integrity, MACsec computes an Integrity Check Value (ICV) on the entire Ethernet frame using the designated cipher suite. The designated cipher suite is also used for encryption.

MACsec adds the ICV to the frame before transmission. The receiving device recalculates the ICV and checks it against the computed value that has been added to the frame. Because the ICV is computed on the entire Ethernet frame, any modifications to the frame can be easily recognized.

By default, both encryption and integrity protection are enabled.

MACsec encrypts traffic between devices at the MAC layer and decrypts frames within participating networked devices. MACsec uses the Galois/Counter Mode Advanced Encryption Standard 128 (GCM-AES-128) cipher suite to encrypt data and to compute the ICV for each transmitted and received MACsec frame.

MACsec also encrypts the VLAN tag and the original Ethertype field in the Layer 2 header of the secured data. When initial bytes in a secure data packet must be transparent, a confidentiality offset of 30 or 50 bytes can be applied.

> **NOTE**
> Refer to Configuring MACsec on page 297 for an overview of enabling and configuring MACsec
> features.

1. At the dot1x-mka group configuration level, enter the **macsec cipher-suite** command with one of the available options:

   - gcm-aes-128: Enables encryption and integrity checking using the GCM-AES-128 cipher suite.
   - gcm-aes-128 integrity-only: Enables integrity checking without encryption.

   In the following example, MACsec encryption has been configured as a group test1 setting. By default, ICV integrity check is also enabled.

   ```
   device# configure terminal
   device(config)# dot1x-mka
   device(config-dot1x-mka)# mka-cfg-group test1
   device(config-dot1x-mka-group-test1)# macsec cipher-suite gcm-aes-128
   ```

   In the following example, MACsec has been configured for integrity protection only, without encryption.

   ```
   device# configure terminal
   device(config)# dot1x-mka
   device(config-dot1x-mka)# mka-cfg-group test1
   device(config-dot1x-mka-group-test1)# macsec cipher-suite gcm-aes-128 integrity-only
   ```

   > **NOTE**
   > The **no** form of the **macsec cipher-suite** command disables both encryption and integrity checking.

2. Enter the **macsec confidentiality-offset** command if an encryption offset is required:

    -   30: Encryption begins at byte 31 of the data packet.
    -   50: Encryption begins at byte 51 of the data packet.

        NOTE
        The default offset for MACsec encryption is zero bytes. Use the **no macsec confidentiality-offset** command to return the offset to zero bytes.

In the following example, the encryption offset is defined as 30 bytes. The first 30 bytes of each data packet carried within the MACsec frame are transmitted without encryption.

```
device# configure terminal
device(config)# dot1x-mka
device(config-dot1x-mka)# mka-cfg-group test1
device(config-dot1x-mka-group-test1)# macsec confidentiality-offset 30
```

Configuring MACsec frame validation

You can specify whether incoming frames are checked for MACsec (secTAG) headers and how invalid frames are handled.

NOTE
Refer to Configuring MACsec on page 297 for an overview of enabling and configuring MACsec features.

At the MKA group configuration level, enter the **macsec frame-validation** command, and select an option:

-   disable: Received frames are not checked for a MACsec header.
-   check: If frame validation fails, counters are incremented, but packets are accepted.
-   strict: If frame validation fails, packets are dropped, and counters are incremented.

In the following example, group test1 is configured to validate frames and discard invalid ones.

```
device# configure terminal
device(config)# dot1x-mka
device(config-dot1x-mka)# mka-cfg-group test1
device(config-dot1x-mka-group-test1)# macsec frame-validation strict
```

Configuring replay protection

MACsec replay protection detects repeated or delayed packets and acts as a safeguard against man-in-the-middle attacks.

When replay protection is configured, MACsec uses a separate replay packet number (PN) counter and gives each Ethernet frame a packet number. As frames are received, packet numbers are monitored.

Two modes of replay protection are supported: strict and out-of-order. In strict mode (the default), packets must be received in the correct incremental sequence. In out-of-order mode, packets are allowed to arrive out of sequence within a defined window.

NOTE
Refer to Configuring MACsec on page 297 for an overview of enabling and configuring MACsec features.

At the dot1x-mka group configuration level, enter the **macsec replay-protection** command with one of the available modes:

- strict: Frames must be received in exact incremental sequence.
- out-of-order *window size*: Frames are accepted out of order within the designated window size. Maximum window size is 4294967295.

In the following example, replay protection is enabled for group test1. Frames must be received in exact order.

```
device# configure terminal
device(config)# dot1x-mka
device(config-dot1x-mka)# mka-cfg-group test1
device(config-dot1x-mka-group-test1)# macsec replay-protection strict
```

In the following example, replay protection is enabled for group test1. Frames are accepted out of order within the designated window size (100).

```
device# configure terminal
device(config)# dot1x-mka
device(config-dot1x-mka)# mka-cfg-group test1
device(config-dot1x-mka-group-test1)# macsec replay-protection out-of-order window-size 100
```

Once you have configured desired MKA group settings, these settings can be applied to specific interfaces.

Enabling and configuring group interfaces for MACsec

After MACsec is enabled for the device, each MACsec interface must be individually enabled, and a configured group of parameters must be applied.

1. To enable MACsec, at the dot1x-mka configuration level, enter the **enable-mka ethernet** command, and specify the interface as *slot/port*.

   NOTE
   To enable a range of interfaces, use this form of the command: **enable-mka ethernet slot/port to slot/port**.

   In the following example, Ethernet port 1 on slot 1 of the device is enabled for MACsec security. The second code example shows the error message received when a port cannot be enabled for MACsec because a license is not installed.

   ```
   device# configure terminal
   device(config)# dot1x-mka
   device(config-dot1x-mka)# enable-mka ethernet 1/1
   device(config-dot1x-mka-eth-1/1)#

   device(config-dot1x-mka)# enable-mka ethernet 1/1
   Error: No MACsec License available for the port 1/1. Cannot enable MACsec !!!
   Error: MKA cannot be enabled on port 1/1
   ```

2. At the dot1x-mka interface configuration level, enter the **mka-cfg-group** command to specify the MKA group configuration to apply to the interface.

   In the following example, MACsec options configured for group test1 are applied to the enabled interface.

   ```
   device(config-dot1x-mka)# enable-mka ethernet 1/1
   device(config-dot1x-mka-eth-1/1)# mka-cfg-group test1
   ```

Configuring the pre-shared key

MACsec security is based on a pre-shared key, the Connectivity Association Key (CAK), which you define and name. Only MACsec-enabled interfaces that are configured with the same key can communicate over secure MACsec channels.

> **NOTE**
> Refer to Configuring MACsec on page 297 for an overview of enabling and configuring MACsec features.

At the dot1x-mka-interface configuration level, enter the **pre-shared-key** command to define and name the pre-shared key.

- *Key id*: Define the key ID value using 32 hexadecimal characters.
- key-name *hex string*: Give the key a name using from 2 through 64 hexadecimal characters.

In the following example, the pre-shared key with the hex value beginning with "0102" and the key name beginning with "1122" are applied to interface 1/1.

```
device# configure terminal
device(config)# dot1x-mka
device(config-dot1x-mka)# enable-mka ethernet 1/1
device(config-dot1x-mka-eth-1/1)# pre-shared-key 01020304050607080907A0B0C0D0E0F10 key-name 11223344
```

> **NOTE**
> The MKA configuration group must be associated with the interface before the pre-shared-key is configured.

Enable and configure each MACsec interface. Configure the same pre-shared key (CAK) on the interfaces between which a secure channel can be established.

Sample MACsec configuration

Here is a complete example of how to enable MACsec, configure general parameters, enable and configure interfaces, and assign a key that is shared with peers.

```
device(config)# dot1x-mka
device(config)# dot1x-mka-enable
device(config-dot1x-mka)# mka-cfg-group test1
device(config-dot1x-mka-cfg-group-test1)# key-server-priority 5
device(config-dot1x-mka-cfg-group-test1)# macsec cipher-suite gcm-aes-128
device(config-dot1x-mka-cfg-group-test1)# macsec confidentiality-offset 30
device(config-dot1x-mka-cfg-group-test1)# macsec frame-validation strict
device(config-dot1x-mka-cfg-group-test1)# macsec replay-protection strict
device(config-dot1x-mka-cfg-group-test1)# exit
device(config-dot1x-mka)# enable-mka ethernet 1/1
device(config-dot1x-mka-1/1)# mka-group test1
device(config-dot1x-mka-1/1)# pre-shared-key 0102030405060708090A0B0C0D0E0F10 key-name 11223344
device(config-dot1x-mka-1/1)# exit
device(config)# exit
device#
```

Displaying MACsec information

Use MACsec **show** commands to display information on MACsec for a device, group, or individual interface. MACsec **show** commands can be used to display configuration information, to report on MACsec sessions that are currently active on a device, or to monitor MACsec statistics on a particular interface.

Displaying MACsec configuration details

You can display configuration information for all MACsec groups on a device, or you can display details for a particular group.

1. In EXEC, Privileged EXEC, global configuration, or dot1x-mka interface mode, use the **show dot1x-mka config** command to display MACsec configuration details for the device.

   In the following example, MACsec parameters are displayed for the device and all groups configured on it. Specific MACsec interfaces are displayed as well as the pre-shared key for each interface.

   ```
   device(config)# show dot1x-mka config
   dot1x-mka-enable
     mka-cfg-group group1
       key-server-priority 20
       macsec frame-validation check
       macsec confidentiality-offset 30
       macsec replay-protection out-of-order window-size 100
     mka-cfg-group group2

   enable-mka ethernet 1/1 to ethernet 1/9
     mka-cfg-group    group1
     pre-shared-key 0102030405060708090A0B0C0D0E0F10 key-name 11223344
   enable-mka ethernet 1/10
     mka-cfg-group    group1
     pre-shared-key 0102030405060708090A0B0C0D0E0F10 key-name 11223344
   ```

2. In EXEC, Privileged EXEC, global configuration, or dot1x-mka interface mode, enter the **show dot1x-mka group** command to display information for configured groups. Add a group name to the command to narrow the information displayed to one group.

   The following example displays information for MKA group test1.

   ```
   device(config)# show dot1x-mka group test1
     Group name group1
     Key Server Priority    : 16
     Cipher Suite           : gcm-aes-128
     Capability             : Integrity, Confidentiality with offset
     Confidentiality Offset : 0
     Frame Validation       : strict
     Replay Protection      : strict
   ```

   > **NOTE**
   > Group information does not include the pre-shared key or enabled connections. Use the **show dot1x-mka config** command to obtain that information.

Displaying information on current MACsec sessions

You can display MACsec session activity for an interface, including the pre-shared key name, the most recent SAI information, and a list of peers.

1. For a quick overview of current MACsec sessions, enter the **show dot1x-mka sessions brief** command in EXEC, Privileged EXEC, global configuration, or dot1x-mka interface mode.

```
device(config)# show dot1x-mka sessions brief

Port    Link-Status  Secured  Key-Server  Negotiated Capability
----    -----------  -------  ----------  ---------------------

4/2     Up           Yes      Yes         Integrity, Confidentiality with offset 0
4/3     Up           Yes      Yes         Integrity, Confidentiality with offset 0
4/4     Up           Yes      Yes         Integrity, Confidentiality with offset 0
4/7     Up           Yes      Yes         Integrity, Confidentiality with offset 0
4/11    Up           Yes      Yes         Integrity, Confidentiality with offset 0
4/12    Up           Yes      Yes         Integrity, Confidentiality with offset 0
4/17    Up           Yes      Yes         Integrity, Confidentiality with offset 0
4/18    Up           Yes      Yes         Integrity, Confidentiality with offset 0
```

2. To display full details on current MACsec sessions, in EXEC, Privileged EXEC, global configuration, or dot1x-mka interface mode, enter the **show dot1x-mka sessions ethernet** command followed by the interface identifier.

```
device(config)# show dot1x-mka sessions ethernet 4/1


Interface              : 4/1

  DOT1X-MKA Enabled     : Yes
  DOT1X-MKA Active      : Yes

Configuration Status:
  Group Name            : 1
  Capability            : Integrity, Confidentiality with offset
  Confidentiality offset : 0

  Desired               : Yes
  Protection            : Yes
  Validation            : Strict
  Replay Protection     : None
  Replay Protection Size : 0
  Cipher Suite          : GCM-AES-128

  Authenticator         : No
  Key Server Priority   : 16
  Algorithm Agility     : 80C201

  CAK NAME              : 11223344

SCI Information:
  Actor SCI             : 0024388f6b900001
  Actor Priority        : 16
  Key Server SCI        : 0024388f6b900001
  Key Server Priority   : 16

MKA Status:
  Enabled               : Yes
  Authenticated         : No
  Secured               : Yes
  Failed                : No

Latest KI, KN and AN Information:
  Latest KI             : 42b4d71d520263cad8727d9100000001
  Tx Key Number         : 1
  Rx Key Number         : 0
  Tx Association Number  : 0
  Rx Association Number  : 0

Participant Information:
  SCI                   : 0024388f6b900001
  Key Identifier        : 1
  Member Identifier     : 42b4d71d520263cad8727d91
  Message Number        : 3491
  CKN Name              : 11223344
  Key Length(in bytes)  : 16

Secure Channel Information:
  No. of Peers (Live and Potential)  : 1
  Latest SAK Status             : Rx & TX
  Negotiated Capability         : Integrity, Confidentiality with offset 0

Peer Information(Live and Potential):
State Member Identifier       Message Number  SCI             Priority  Capability
----- -----------------       -----------     ----------------  --------  ----------
Live  66dfa9b5037a9c7aa8b5c71e 3490            0024389e2d300001 16         2
```

Displaying MKA protocol statistics for an interface

You can display a report on MKA protocol activity for a particular interface.

In EXEC, Privileged EXEC, global configuration, or dot1x-mka interface mode, enter the **show dot1x-mka statistics ethernet** command to display MKA protocol statistics for the designated interface.

```
device(config)# show dot1x-mka statistics ethernet 3/2

Interface                : 3/2

MKA in Pkts              : 89858
MKA in SAK Pkts          : 0
MKA in Bad Pkts          : 0
MKA in Bad ICV Pkts      : 0
MKA in Mismatch Pkts     : 0
MKA out Pkts             : 90225
MKA out SAK Pkts         : 192
```

Displaying MACsec secure channel activity for an interface

You can display currently enforced MACsec capabilities for a specific interface, along with secure channel statistics.

1. At the EXEC or Privileged EXEC level, enter the **clear macsec statistics ethernet** command for the designated interface. Results of the previous **show macsec ethernet** command are removed.

2. In EXEC, Privileged EXEC, global configuration, or dot1x-mka interface mode, enter the **show macsec statistics ethernet** command to display information on MACsec configuration and secure channel activity for a particular interface.

3. Enter the **show macsec statistics ethernet** command to display information on MACsec configuration and secure channel activity for a particular interface.

```
device(config)# clear macsec statistics ethernet 1/1
device(config)# show macsec statistics ethernet 1/1
Interface statistics
--------------------
rx Untagged Pkts        : 3              tx Untagged Pkts        : 0
rx Notagged Pkts        : 0              tx Too long Pkts        : 0
rx Bad Tag  Pkts        : 0
rx Unknown SCI Pkts     : 0
rx No SCI Pkts          : 0
rx Overrun Pkts         : 0

Transmit Secure Channels
------------------------

SC Statistics
  Protected Pkts        : 0              Protected Octets        : 0
  Encrypted Pkts        : 3              Encrypted Octets        : 144

SA[0] Statistics - In use
  Protected Pkts        : 3
  Encrypted Pkts        : 3


SA[1] Statistics
  Protected Pkts        : 0
  Encrypted Pkts        : 0


SA[2] Statistics
  Protected Pkts        : 0
  Encrypted Pkts        : 0


SA[3] Statistics
  Protected Pkts        : 0
  Encrypted Pkts        : 0


Receive Secure Channels
-----------------------

SC Statistics
  OK Pkts               : 0              Not Valid Pkts          : 0
  Unchecked Pkts        : 0              Not using SA Pkts       : 0
  Delayed Pkts          : 0              Unused SA Pkts          : 0
  Late Pkts             : 0              Validated Octets        : 0
  Invalid Pkts          : 0              Decrypted Octets        : 0

SA[0] Statistics - In use
  OK Pkts               : 0              Invalid Pkts            : 0
  Not using SA Pkts     : 0              Unused SA Pkts          : 0


SA[1] Statistics
  OK Pkts               : 0              Invalid Pkts            : 0
  Not using SA Pkts     : 0              Unused SA Pkts          : 0


SA[2] Statistics
  OK Pkts               : 0              Invalid Pkts            : 0
  Not using SA Pkts     : 0              Unused SA Pkts          : 0


SA[3] Statistics
  OK Pkts               : 0              Invalid Pkts            : 0
  Not using SA Pkts     : 0              Unused SA Pkts          : 0
```

Policy-Based Routing

Policy-based routing overview

Policy-based routing (PBR) is the process of altering a packet's path based on criteria other than the destination address. PBR allows you to use ACLs and route maps to selectively route an IP packet. The ACLs classify the traffic and the route maps that match on the ACLs set routing attributes for the traffic.

A PBR policy specifies the routing attributes for traffic that matches the policy. Using standard ACLs with route maps in PBR, an IP packet is routed based on their source IP address. With extended ACLs, you can route IP packets based on all of the clauses in the extended ACL.

You can configure the device to perform the following types of PBR based on a packet Layer 3 and Layer 4 information:

* Select the next-hop gateway.
* Send the packet to the null interface (null0).

To configure PBR, you define the policies using ACLs and route maps, then enable PBR on individual interfaces. The platform programs the ACLs on the interfaces, and routes traffic that matches the ACLs according to the instructions provided by the "set" statements in the route map entry.

The configuration of a set of match criteria and corresponding routing information (for example next hops) is referred to as a stanza. You can create multiple stanzas and assign an "Instance_ID" that controls the program positioning within the route map. Furthermore, when the route map is created, you specify a deny or permit construct. In addition, the ACL used for the "match" criteria also contains a deny or permit construct.

The deny or permit nomenclature has a different meaning within the context of the PBR operation than it does within the normal context of user-applied ACLs (where deny and permit are directly correlated to the forwarding actions of forward and drop). The following table lists the behavior between the permit and deny actions specified at the route-map level, in conjunction with the permit and deny actions specified at the ACL rule level.

TABLE 39 Behavior of Permit and Deny actions in different contexts

Route-map level permit and deny actions	ACL clause permit and deny actions	Resulting Ternary Content Addressable Memory (TCAM) action
Permit	Permit	The "set" statement of the route-map entry is applied.
Permit	Deny	The packet is "passed" and routed normally. The contents of the "set" command are not applied. A rule is programmed in the TCAM as a "permit" with no result actions preventing any further statements of the route-map ACL from being applied.
Deny	Permit	The packet is "passed" and routed normally. There should be no "set" commands following the "match" command of a deny route-map. A rule is programmed in the TCAM as a "permit" with no result actions preventing any further statements of the route-map ACL from being applied.
Deny	Deny	No TCAM entry is provisioned; no other route-map ACL entries will be compared against. If no subsequent matches are made, the packet is forwarded as normal.

LAG formation with PBR policy

PBR can be applied on LAG primary port ,only if the port is untagged. When a LAG is formed, the PBR policy on the primary port applies to all the secondary ports. If a different PBR policy exists on a secondary port at the time of a LAG formation, that policy is overridden by the PBR policy on the primary port. If the primary port does not have a PBR policy, then the secondary ports will not have a PBR policy.

When a LAG is removed, the PBR policy that was applied to the LAG virtual interface is unbound (removed) from former secondary ports. If global PBR is configured, the secondary ports adhere to the global PBR; otherwise, no PBR policy is bound to former secondary ports.

When a LAG is formed, all ports must have the same PBR configuration before deployment, during deployment the configuration on the primary port is replicated to all ports and on undeployment each port inherits the same PBR configuration.

Route map

Route maps enable you to define routing policy for the traffic causing a packet to be forwarded to a predetermined next-hop interface, bypassing the path determined by normal routing.

Each entry in a route map statement contains a combination of match and set statements. A route map specifies the match criteria that correspond to ACLs, followed by a set statement that specifies the resulting action if all of the match clauses are met. You can define multiple match and next-hop specifications (set statement) on the same interface. When a PBR policy has multiple next hops to a destination, PBR selects the first live next hop specified in the policy that is up. If none of the policy's direct routes or next hops is available, the packets are forwarded as per the routing table.

Match statement

The IP standard or extended ACLs can be used to establish the match criteria. Using the standard ACLs with route maps in PBR, an IP packet is routed based on their source IP address. Using the extended ACLs, you can route IP packets based on all of the clauses in the extended ACL.

Set statement

Traffic that matches a match statement in the route map is forwarded as defined by set commands. Multiple set commands can be configured and when a match condition is met, the device works sequentially through the list of set commands until it finds the first "next hop" that is operational and uses it. If that "next hop" goes down, the next hop as defined in a set command is chosen and if all next hop interfaces in the list are down, the packet is routed as determined in the IP Route Table. If a next hop interface that was down comes back up, the next hop selection process begins again and restarts its selection process from the top of the list.

The set clauses are evaluated in the following order:

1. Set clauses where the next hop is specified.
2. Set interface NULL0.

The order in which you enter the **set ip next-hop** commands determines the order preference. If no next hops are reachable, the egress interface is selected based on the order of interface configuration. The set interface NULL0 clause — regardless of which position it was entered — is always placed as the last selection in the list.

> NOTE
> The "match" and "set" statements described in this chapter are the only route-map statements supported for PBR. Other route-map statements described in the documentation apply only to the protocols with which they are described.

NOTE
If none of the clauses of a PBR routemap definition contains both 'match' and 'set' statements together, PBR does not work and the packets are forwarded as per the routing table.

The following are the PBR next hops that can be specified in a route map for a matched traffic:

- IPv4 address
- IPv6 address
- Null interface
- Physical interface
- GRE tunnel
- GRE physical interface
- Label Switched Path (LSP)
- VLAN flooding
- TVF domain

NOTE
For Layer 2 PBR, only the next hop types configured as set interface null0 or physical interface, and flood VLAN are selected. The next hop types configured as IP, IPv6, tunnel, or LSP is ignored.

Rule name

Rule name can be defined in the route map to organize and extract information about PBR configurations using the **rule-name** command. That is, you can assign a name to a specific instance in a route map. The following are the configuration notes for rule name.

- The same rule names can be used in multiple route maps.
- For each route map instance, there can be one rule name configured.
- Within each route map, the rule name has to be unique.
- The maximum number of rule names that can be configured in a system is 4096.

Configuring a PBR policy

To configure a PBR, you must define the policies using ACLs and route map, then enable PBR globally or on individual interfaces. The device programs the ACLs into the session CAM on the interfaces and routes traffic that matches the ACLs according to the instructions in the route maps.

The following are the basic steps to configure a PBR policy:

1. Configure ACLs that specify all the conditions required to match the desired packets to be routed using PBR.

   For more information about configuring ACLs, refer to the *Security Configuration Guide* for your platform.
2. Configure a route map that matches on the ACLs and sets the route information.
3. Enable PBR by applying the route map globally or on an untagged interface or virtual interface.

Policy-based routing (IPv4)

This section provides configuration details of PBR policy that specifies to match and route IPv4 packets.

Configuration considerations and guidelines for PBR

- PBR is supported in the full Layer 3 code only.
- PBR is not supported together with ingress ACLs on the same port.
- Global PBR is not supported with per-port-per-VLAN ACLs.
- Global PBR is supported only on default VRF.
- A PBR policy on an interface takes precedence over a global PBR policy.
- You cannot apply PBR on a port if that port already has inbound ACLs, inbound ACL-based rate limiting, or TOS-based QoS.
- PBR is supported on OpenFlow hybrid interfaces along with OpenFlow flows. For more information, refer to *Extreme NetIron SDN Configuration Guide*.
- The number of route maps that you can define is limited by the system memory. When a route map is used in a PBR policy, the PBR policy uses up to 200 instances of a Layer 3 route map, up to 5 ACLs in a matching policy of each route map instance.

  The following two conditions can cause more than 200 Layer 3 route-map instances to be used.
  - If one or more of first 200 instances have deny clause.
  - If the access-list used in the first 200 instances is not configured.
- Denial of Service (DoS) protection is supported in PBR.
- IPv4 ACL accounting is supported in PBR.
- ACLs with the **log** option configured in the **access-list** command should not be used for PBR purposes.
- PBR ignores implicit **deny ip any any** ACL entries to ensure that the traffic is compared to all ACLS for route maps that use multiple ACLs. However, if an explicit **deny ip any any** is configured, traffic matching this clause will be routed normally using Layer 3 paths and will not be compared to any ACL clauses that follow this clause.
- PBR always selects the first next hop from the next-hop list that is up. If a PBR policy next hop goes down, the policy uses another next hop if available. If no next hops are available, the device routes the traffic in the normal way.
- PBR is not supported for fragmented packets. If the PBR ACL filters on Layer 4 information such as TCP/UDP ports, fragmented packed are routed normally.
- You can change route maps or ACL definitions dynamically and do not need to rebind the PBR policy to an interface.
- Any changes to route maps or ACL definitions will be effective immediately for the interfaces where the PBR routemap is applied. There is no need to rebind. However, rebinding is required if a change is made to an IPv6 ACL.
- If a PBR policy is applied globally, inbound ACLs, inbound ACL-based rate-limiting or TOS-based QoS cannot be applied to any port on the device.
- If an IPv4 option packet matches a **deny** ACL filter with the **option** keyword, the packet will be forwarded based on Layer-3 destination. If the **ignore-options** command is configured on the incoming physical port, the packet will be forwarded based on its Layer-3 destination in hardware, otherwise the packet will be sent to the CPU for software forwarding.
- If an IPv4 option packet matches a **permit** ACL filter with the option keyword, it is hardware-forwarded based on its PBR next-hop (if available). If no PBR next-hop is available, the packet is either software or hardware-forwarded (depending on whether **ignore-options** is configured), based on an IP forwarding decision.
- Policy Based Routing (PBR) currently does not support the IPv4 and IPv6 features for changing the MTU.
- PBR route-map cannot be applied on VPLS, VLL, or VLL-Local endpoints and vice-versa.
- PBR policies are not supported on Layer-3 VPNs.

- In a PBR route-map definition, if even one route-map instance contains a "set next-hop-flood-vlan" statement, all instances of that route-map will apply to both routed and switched traffic.

- Flooding traffic to a POS interface is not allowed. It can only be flooded to Ethernet ports on the VLAN, including the default VLAN.

- When an incoming port is POS then the SA of the outgoing flooded packets will be 0.

- Where the next hop is a GRE tunnel:

  - Packets that are larger than the tunnel's MTU are subject to IP fragmentation and PBR processing of the fragmented packets.
  - For route changes of the tunnel destination, the appropriate information is automatically propagated to the PBR feature. Depending on the configuration of the route map, a route change can change the active next hop of the PBR if it leads to the active next hop going down which triggers a new next hop selection process.
  - In the situation where IPv6 traffic is encapsulated in an IPv4 tunnel:

    - › The remote end of the tunnel must not be terminated on a BR-MLX-10Gx24-DM module.
    - › To ensure that the IPv6 traffic is forwarded properly, IPv6 selective routes download for the default VRF must be disabled. IPv6 selective routes download is disabled by default; when it has been previously enabled, use the **no ipv6 selective-routes-download** command in global configuration mode to disable it.

Configuring an IPv4 PBR with IPv4 address as the next hop

The following steps configure an IPv4 PBR by setting an IPv4 address as the next hop in the route map. This task uses Access Control Lists (ACLs) which are explained in greater detail in the Security configuration guide for your platform.

1. Enter the **configure terminal** command to enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Define the required IPv4 ACLs to be added to the route map.

   ```
   device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
   ```

3. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all the match clauses are met.

   ```
   device(config)# route-map test-route permit 99
   ```

4. Add IPv4 ACLs to match the IP address that is permitted by the ACL.

   ```
   device(config-routemap test-route)# match ip address 99
   ```

5. Set the IPv4 address of the next hop to which the traffic that matches a match statement in the route map must be routed.

   ```
   device(config-routemap test-route)# set ip next-hop 192.168.3.1
   ```

   > NOTE
   > If the IP address used in this command is the IP address of a configured GRE tunnel, the configuration will choose the GRE tunnel interface as the next-hop address for matching packets. However, the support to set the next hop to GRE physical interface is recommended only for use on the BR-MLX 24x10G-DM module. For all other modules, the configuration will still be accepted but the next-hop selection will never choose this next-hop so it will not become active. If you want to set the next hop using a GRE tunnel, use the **set next-hop-ip-tunnel** command.

   Optionally, you can configure the route map to retain the Layer 2 and Layer 3 information of the packets using the **preserve-vlan** option. For more information about the usage of **preserve-vlan** option, refer to the Policy based routing with the preserve VLAN option section.

6. Enter the **exit** command to return to the global configuration mode.

    ```
    device(config-routemap test-route)# exit
    ```

7. Enable PBR by applying the route map globally or on an untagged interface or virtual interface.

    *   Enable IPv4 PBR globally to apply the route map to all interfaces.

        ```
        device(config)# ip policy route-map test-route
        ```

    *   Enable IPv4 PBR locally by applying the route map on an interface.

        ```
        device(config)# interface ethernet 1/1
        device(config-if-e1000-1/1)# ip policy route-map test-route
        ```

    *   Enable IPv4 PBR locally by applying the route map on a virtual interface.

        ```
        device(config)# interface ve 1
        device(config-vif-1)# ip policy route-map test-route
        ```

The following example shows the configuration steps to configure an IPv4 PBR by setting an IPv4 address as the next hop in the route map.

```
device# configure terminal
device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ip address 99
device(config-routemap test-route)# set ip next-hop 192.168.3.1
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ip policy route-map test-route
device(config-if-e1000-1/1)# end
```

The following example shows the configuration steps to configure an IPv4 PBR in which the route map is configured to preserve VLAN IDs and forward the packet to specific egress port.

```
device# configure terminal
device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ip address 99
device(config-routemap test-route)# set ip next-hop 192.168.3.1 preserve-vlan
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ip policy route-map test-route
device(config-if-e1000-1/1)# end
```

Configuring an IPv4 PBR with NULL0 interface as the next hop

The following steps configure an IPv4 PBR by setting NULL0 interface as the next hop in the route map.

1. Enter the **configure terminal** command to enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Define the required IPv4 ACLs to be added to the route map.

    ```
    device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
    ```

3. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all the match clauses are met.

    ```
    device(config)# route-map test-route permit 99
    ```

4. Add IPv4 ACLs to match the IP address that is permitted by the ACL.

```
device(config-routemap test-route)# match ip address 99
```

5. Set the next hop as NULL0 interface to send the traffic to the null interface, thus dropping the packet instead of forwarding it.

```
device(config-routemap test-route)# set interface null0
```

6. Enter the **exit** command to return to the global configuration mode.

```
device(config-routemap test-route)# exit
```

7. Enter configuration mode on the interface where you want to enable PBR by applying the route map.

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

PBR can be enabled globally by which the route map is applied to all interfaces using the **ip policy route-map** command from the global configuration mode.

8. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-1/1)# ip policy route-map test-route
```

The following example shows the configuration steps to configure an IPv4 PBR to send all traffic from 10.157.23.0 0.0.0.255 to the null interface, thus dropping the traffic instead of forwarding it.

```
device# configure terminal
device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ip address 99
device(config-routemap test-route)# set set interface null0
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ip policy route-map test-route
device(config-if-e1000-1/1)# end
```

VLAN flooding with PBR

VLAN flooding using PBR allows you to flood the matched traffic on all ports of the VLAN except the incoming physical port.

Any PBR policy with a set statement to flood the VLAN applies to both routed and switched traffic. This means that if any instance in a PBR route map contains the **set next-hop-flood-vlan** *vlan-id* command, all instances of that route-map will be applied to both routed and switched traffic.

> NOTE
> Always use Transparent VLAN Flooding for the VLAN that is specified in the **set next-hop-flood-vlan** *vlan-id* command.

If the VLAN specified in the route map is not configured, the PBR route map set statement will fall through to the next configured set statement. If no valid next hop is available, the packet is forwarded as per Layer 2 or Layer 3 forwarding decision. If the specified VLAN has no valid outgoing ports, (such as when all ports in the VLAN are down or when the VLAN is empty) matching packets will be dropped. There is also no Layer 3 header processing. For example, in the case of an IPv6 packet header, the hop-limit will not be decremented.

In the case of traffic incoming on MPLS uplink, PBR to VLAN flooding is only supported for IP traffic, and not for MPLS traffic.

Optionally, you can set the device to send packets flooded to the ports on the VLAN with the destination MAC address using the **outgoing-da** option. If the destination MAC address is not set using the **outgoing-da** option, the destination address is set as described in the following table.

TABLE 40 Destination address on VLAN flooded packets

Incoming port	Outgoing port	Routed traffic	Switched traffic
Ethernet	Ethernet	Replaced Destination Address	Original Destination Address from Incoming Packet
POS	Ethernet	Replaced Destination Address	N/A

The behavioral differences when deployed on a CES 2000 Series or CER 2000 Series device are described in the following table.

TABLE 41 Behavioral differences of VLAN flooding per platform

Scenario	CES 2000 Series/ CER 2000 Series	MLX Series/ XMR Series
Support for Switched Traffic	Not Supported	Supported
If the outgoing-da option is not set	Packet floods with original SA/DA	Switched traffic: Packet floods with original SA/DA. Routed traffic (L3 CAM miss): Packet floods with original SA/DA. Routed traffic (L3 CAM hit): Packet floods with original SA and DA as next-hop MAC.
If the outgoing-da option is set	DA: Configured outgoing-da SA: Outgoing port MAC	DA: Configured outgoing-da SA: Original SA
If the specified VLAN is configured but no ports are added or all ports are down	Go to next set statement.	Drop matching packets.
Ingress VLAN and PBR Flood VLAN are the same.	Packets are flooded back to source port.	Packets are not flooded back to source port.

Policy based routing with the preserve VLAN option

The Layer 2 and Layer 3 information of the packets that matches PBR policy can be retained (for example, the VLAN information and the MAC address are retained) using the **preserve-vlan** option in the set statement of the route map.

The **preserve-vlan** option can be configured in the following route map set policy configuration:

- **set ip next-hop**
- **set next-hop-flood-vlan**
- **set interface ethernet** .

  NOTE
  Policy-based routing support for preserve VLAN is supported only on XMR Series, MLX Series routers, and on the BR-MLX-10GX24 module.

The interface, on which PBR with preserve VLAN is configured, should be part of the VLANs through which packets are expected. TTL is not decremented when the packet is sent to the configured next hop. A route map policy with set policies to preserve VLAN can be applied on a physical port or on a VE port. The route map with preserve-vlan set policies cannot be configured globally. A route map used for PBR with a preserve VLAN policy must have the **preserve-vlan** option configured for each set policy. Packets not matching the

PBR policy with preserve-vlan will be dropped. If none of the policy's direct routes or next hops is available, the packets are forwarded as per the routing table.

Previously, PBR transparent VLAN flooding (TVF) replaced the ingress traffic's VLAN ID with the egress TVF VLAN ID, while flooding the egress TVF VLAN. Policy-based routing support for preserve VLAN allows the ingress packet VLAN header (VLAN ID and priority) to be preserved, while simultaneously flooding the PBR TVF VLAN.

The PBR TVF VLAN egress ports can be in strict tagged VLAN mode or dual VLAN mode. When PBR TVF VLAN egress ports are in strict tagged VLAN mode, the ingress tagged packets flood as "tagged" with the original VLAN ID and priority preserved. The ingress untagged packets flood as "tagged" with the default VLAN ID. When the PBR TVF VLAN egress ports are in dual VLAN mode, the ingress tagged packets flood as "tagged" with the original VLAN ID and priority preserved. The ingress untagged packets flood as "untagged" if the egress port is the untagged member of the ingress port's default VLAN; otherwise, the ingress untagged packets flood as "tagged" with the original VLAN ID and priority preserved.

Consider the following when PBR is configured for preserve VLAN:

- To preserve the ingress VLAN priority value, the ingress VLAN and the port QoS feature should not be configured at the same time.
- IPv4 and IPv6 ACL VLAN ID matches are supported for both ingress and egress ACLs.
- An egress ACL is supported to filter traffic.
- Policy-based routing support for the preserve VLAN option does not affect the feature implementation of policy-based routing support for the preserve VLAN.

The **allow-all-vlan pbr** command configures a physical interface to accept all VLAN packets for the purpose of PBR. This command reduces configuration complexity since the physical interface does not have to be configured individually in multiple VLAN interfaces. The **allow-all-vlan pbr** command cannot be configured when the physical port is configured with an IPv4 address, IPv6 address, MPLS, VPLS, VLL, ICL, Layer 3 VPN; or when the port is part of other VLAN. The **allow-all-vlan pbr** command cannot be applied to a VE. If OpenFlow and **allow-all-vlan pbr** commands are configured on the same port, the OpenFlow properties will have higher priority.

Configuring an IPv4 PBR to flood traffic on all ports of a VLAN

A PBR policy can be configured to flood the matched traffic on all ports of the VLAN except the incoming physical port using the **set next-hop-flood-vlan** statement in the route map.

The following steps configure an IPv4 PBR to flood the matched traffic on all ports of a VLAN.

1. Enter the **configure terminal** command to enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Define the required IPv4 ACLs to be added to the route map.

   ```
   device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
   ```

   For more information about configuring ACLs, refer to the *Security Configuration Guide* for your platform.

3. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

   ```
   device(config)# route-map test-route permit 99
   ```

4. Add IPv4 ACLs to match the IP address that is permitted by the ACL.

   ```
   device(config-routemap test-route)# match ip address 99
   ```

5. Set the next hop to route the matched traffic to flood on all ports of a VLAN.

   ```
   device(config-routemap test-route)# set next-hop-flood-vlan 10
   ```

   • (Optional) Set the packets that are flooded on the ports of a VLAN to carry the destination MAC address using the **outgoing-da** option.

     ```
     device(config-routemap test-route)# set next-hop-flood-vlan 10 outgoing-da 0000.1234.5678
     ```

   • (Optional) Set the route map to retain the Layer 2 and Layer 3 information of the packets that are flooded to the VLAN using the **preserve-vlan** option.

     ```
     device(config-routemap test-route)# set next-hop-flood-vlan 10 preserve-vlan
     ```

6. Enter the **exit** command to return to the global configuration mode.

   ```
   device(config-routemap test-route)# exit
   ```

7. Enter configuration mode on the interface where you want to enable IPv4 PBR by applying the route map.

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

8. Enable policy-based routing on the interface and specify the route map to be used.

   ```
   device(config-if-e1000-1/1)# ip policy route-map test-route
   ```

The following example shows the configuration steps to forward matching traffic to an interface.

```
device# configure terminal
device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ip address 99
device(config-routemap test-route)# set next-hop-flood-vlan 10
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ip policy route-map test-route
```

Configuring an IPv4 PBR with an interface as the next hop

Within a route-map instance, there can be only one match statement but multiple interfaces can be set as next hops. If there are multiple interfaces configured in the **set** statement, the first one in the configured order will be used to forward traffic. If the actively used interface is down, the next interface in the configuration order will take over. If the previously down interface comes back up, the traffic will be reverted to the first interface.

The following steps configure an IPv4 PBR to forward matching traffic to an interface.

1. Enter the **configure terminal** command to enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Define the required IPv4 ACLs to be added to the route map.

   ```
   device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
   ```

3. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

   ```
   device(config)# route-map test-route permit 99
   ```

4. Add IPv4 ACLs to match the IP address that is permitted by the ACL.

    ```
    device(config-routemap test-route)# match ip address 99
    ```

5. Set an interface as the next hop for the route map and preserve VLAN and L2 information for the matching packets.

    ```
    device(config-routemap test-route)# set interface ethernet 3/1 preserve-vlan
    ```

6. Enter the **exit** command to return to the global configuration mode.

    ```
    device(config-routemap test-route)# exit
    ```

7. Enter configuration mode on the interface where you want to enable IPv4 PBR by applying the route map.

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

8. Enable policy-based routing on the interface and specify the route map to be used.

    ```
    device(config-if-e1000-1/1)# ip policy route-map test-route
    ```

The following example shows the configuration steps to forward matching traffic to an interface.

```
device# configure terminal
device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ip address 99
device(config-routemap test-route)# set interface ethernet 3/1 preserve-vlan
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ip policy route-map test-route
```

Configuring an IPv4 PBR with a tunnel as the next hop

Traffic can be configured to route over a GRE tunnel using PBR. The following steps configure a GRE tunnel interface as the next hop of a PBR route map.

You must configure GRE tunnel before configuring the traffic to route over a tunnel. For more information about GRE tunnel configuration, refer to the *Routing Configuration Guide*.

1. Enter the **configure terminal** command to enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Define the required IPv4 ACLs to be added to the route map.

    ```
    device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
    ```

3. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all the match clauses are met.

    ```
    device(config)# route-map test-route permit 99
    ```

4. Add IPv4 ACLs to match the IP address that is permitted by the ACL.

    ```
    device(config-routemap test-route)# match ip address 99
    ```

5. Set the configured tunnel as the next hop for a route map.

    ```
    device(config-routemap test-route)# set next-hop-ip-tunnel 1
    ```

6. Enter the **exit** command to return to the global configuration mode.

```
device(config-routemap test-route)# exit
```

7. Enter configuration mode on the interface where you want to enable IPv4 PBR by applying the route map.

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

The IPv4 PBR can be enabled globally by which the route map is applied to all interfaces using the **ip policy route-map** command from the global configuration mode.

8. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-1/1)# ip policy route-map test-route
```

The following example shows the configuration steps to route routing traffic over a GRE tunnel using PBR.

```
device(config)# interface tunnel 1
device(config-tnif-1)# tunnel source ethernet 1/1
device(config-tnif-1)# tunnel destination 10.2.2.1
device(config-tnif-1)# tunnel mode gre ip
device(config-tnif-1)# ip address 10.4.4.4/24
device(config-tnif-1)# exit
device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ip address 99
device(config-routemap test-route)# set next-hop-ip-tunnel 1
device(config-routemap test-route)# end
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1/3)# ip policy route-map test-route
device(config-if-e1000-1/1/3)# end
```

Configuring an IPv4 PBR with TVF domain as the next hop for TVF with LAG load balancing

The transparent VLAN flooding (TVF) domain provides an infrastructure to increase the overall egress traffic streams. By setting TVF domain as a PBR next hop, traffic can be flooded to the TVF domain that supports 2016 TVF instances with LAG load balancing and, together with 4090 TVF instances without LAG load balancing, scales the overall egress traffic flow support to 6106.

You must configure the transparent VLAN flooding (TVF) domain to be set as PBR next hop. For more information about transparent VLAN flooding domain, refer to the *Switching Configuration Guide* for your platform.

The following steps configure the TVF domain as the next hop for a route map to support transparent VLAN flooding (TVF) with LAG load balancing.

1. Enter the **configure terminal** command to enter global configuration mode.

```
device# configure terminal
```

2. Configure system maximum requirements to support LAGs and TVF LAG load balancing.

```
device(config)# system-max trunk-num 32
device(config)# system-max tvf-lag-lb-fid-pool 4096
device(config)# system-max tvf-lag-lb-fid-group 4
```

3. Configure the TVF domain and assign the ports.

```
device(config)# tvf-domain 1
device(config-tvf-domain-1)# port ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-tvf-domain-1)# exit
```

4. Define the required IPv4 ACLs to be added to the route map.

```
device(config)# ip access-list extended v4_brc_Test_p000_bi_moof
device(config-ext-v4_brc_Test_p000_bi_moof)# permit vlan 1001 ip any any
device(config-ext-v4_brc_Test_p000_bi_moof)# exit
```

5. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

```
device(config)# route-map Mall permit 1001
```

6. Add IPv4 ACLs to match the IP address that is permitted by the ACL.

```
device(config-routemap Mall)# match ip address v4_brc_Test_p000_bi_moof
```

7. Enter the **set next-hop-tvf-domain** command to configure a TVF domain as the next hop for a route map.

```
device(config-routemap Mall)# set next-hop-tvf-domain 1
```

8. Enter the **exit** command to return to the global configuration mode.

```
device(config-routemap Mall)# exit
```

9. Configure a VLAN and add the interfaces required for the TVF LAG and LAG load balancing.

```
device(config)# vlan 1 name DEFAULT-VLAN
device(config-vlan-1)# no untagged ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-vlan-1)# no spanning-tree
```

10. Enter the **exit** command to return to the global configuration mode.

```
device(config-vlan-1)# exit
```

11. Create a LAG and add the interfaces.

```
device(config)# lag blue static
device(config-lag-blue)# ports ethernet 5/1 ethernet 6/1
device(config-lag-blue)# primary-port 5/1
device(config-lag-blue)# deploy
```

12. Enter the **exit** command to return to the global configuration mode.

```
device(config-lag-blue)# exit
```

13. Enter configuration mode on the interface where you want to enable IPv4 PBR by applying the route map.

```
device(config)# interface ethernet 11/1
```

The IPv4 PBR can be enabled globally by which the route map is applied to all interfaces using the **ip policy route-map** command from the global configuration mode.

14. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-11/1)# ip policy route-map Mall
```

The following example shows the configuration steps to configure the TVF domain as the next hop for a route map to support transparent VLAN flooding (TVF) with LAG load balancing.

```
device# configure terminal
device(config)# system-max trunk-num 32
device(config)# system-max tvf-lag-lb-fid-pool 4096
device(config)# system-max tvf-lag-lb-fid-group 4
device(config)# tvf-domain 1
device(config-tvf-domain-1)# port ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-tvf-domain-1)# exit
device(config)# ip access-list extended v4_brc_Test_p000_bi_moof
device(config-ext-v4_brc_Test_p000_bi_moof)# permit vlan 1001 ip any any
device(config-ext-v4_brc_Test_p000_bi_moof)# exit
device(config)# route-map Mall permit 1001
device(config-routemap Mall)# match ip address v4_brc_Test_p000_bi_moof
device(config-routemap Mall)# set next-hop-tvf-domain 1
device(config-routemap Mall)# exit
device(config)# vlan 1 name DEFAULT-VLAN
device(config-vlan-1)# no untagged ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-vlan-1)# no spanning-tree
device(config-vlan-1)# exit
device(config)# lag blue static
device(config-lag-blue)# ports ethernet 5/1 ethernet 6/1
device(config-lag-blue)# primary-port 5/1
device(config-lag-blue)# deploy
device(config-lag-blue)# exit
device(config)# interface ethernet 11/1
device(config-if-e1000-11/1)# ip policy route-map Mall
```

Configuring an IPv4 PBR with an LSP as the next hop

The following steps configure an IPv4 PBR to forward matching traffic to an RSVP-signaled LSP.

An LSP must be configured before configuring the IPv4 PBR to forward matching traffic to an RSVP-signaled LSP. For more information about LSP configuration, refer to the *Extreme NetIron MPLS Configuration Guide* for your platform.

1. Enter the **configure terminal** command to enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Define the required IPv4 ACLs to be added to the route map.

    ```
    device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
    ```

3. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

    ```
    device(config)# route-map test-route permit 99
    ```

4. Add IPv4 ACLs to match the IP address that is permitted by the ACL.

    ```
    device(config-routemap test-route)# match ip address 99
    ```

5. Set the configured LSP as the next hop for a route map.

    ```
    device(config-routemap test-route)# set next-hop-lsp t3
    ```

    For more information about LSP configuration, refer to the *Extreme NetIron MPLS Configuration Guide*.

6. Enter the **exit** command to return to the global configuration mode.

    ```
    device(config-routemap test-route)# exit
    ```

7. Enter configuration mode on the interface where you want to enable IPv4 PBR by applying the route map.

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

The IPv4 PBR can be enabled globally by which the route map is applied to all interfaces using the **ip policy route-map** command from the global configuration mode.

8. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-1/1)# ip policy route-map test-route
```

The following example shows the configuration steps to forward matching traffic to an RSVP-signalled LSP.

```
device(config)# router mpls
device(config-mpls)# mpls-interface ethernet 6/1
device(config-mpls)# lsp t3
device(config-mpls-lsp-t3)# to 10.1.1.1
device(config-mpls-lsp-t3)# enable
device(config-mpls-lsp-t3)# exit
device(config)# access-list 101 permit tcp any any
device(config)# access-list 101 deny ip any any
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ip address 101
device(config-routemap test-route)# set next-hop-lsp t3
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ip policy route-map test-route
```

Displaying PBR route map information

Various show commands can be used to display PBR configuration details.

Enter the **show route-map** command to display the route map information.

```
device# show route-map
Interface Ethernet 1/6
 ip policy route-map route1
```

Enter the **show route-map** *route-map name* command to display the details of the configured routing attributes.

```
device# show route-map route1
Interface Ethernet 1/6
ip policy route-map route1 permit 1 (Active)
    match ip address acl test1
    set ip next-hop 6.0.0.1 (selected)
    Policy routing matches: 1443 packets
```

Enter the **show route-map interface ethernet** *slot/port* command to display the route-map configuration details on a specific interface.

```
device# show route-map interface ethernet 1/6
Interface Ethernet 1/6
ip policy route-map route1 permit 1 (Active)
    match ip address acl test1
    set ip next-hop 6.0.0.1 (selected)
    Policy routing matches: 1543 packets
```

Policy-based routing (IPv6)

IPv6 PBR allows you to manually configure how IPv6 packets that match certain criteria can be forwarded instead of following the IPv6 Routing Table Manager (RTM) routes. This section provides configuration details of PBR policy that specifies to match and route IPv6 packets.

Configuration considerations

The configuration considerations are as follows:

- IPv6 PBR cannot be applied globally. IPv6 PBR can only be applied at the interface level.
- IPv6 PBR only supports default VRF. Multi-VRF is not supported.
- IPv6 PBR currently does not support the changing IPv4/IPv6 MTU.
- IPv6 PBR policies are not supported on Layer 3 VPNs.
- IPv6 PBR is applied to routed traffic only by default, except when the flood VLAN option is enabled.
- IPv6 PBR can only be configured on physical ports, Link Aggregation Groups (LAG) ports, and Virtual Ethernet (VEs).
- If IPv6 PBR is applied on a VE, it works only when the VE is enabled. When the VE is disabled, IPv6 PBR will not work and normal routing or switching takes place for the traffic received on the VE.
- The following combinations of IPv6 PBR and IPv4 PBR and IPv6 ACL and IPv4 ACL are allowed:
  - IPv6 PBR and IPv4 PBR can be applied to the same interface at the same time.
  - IPv6 PBR and IPv4 ACL can be applied to the same interface at the same time.
  - IPv4 PBR and IPv6 ACL can be applied to the same interface at the same time.
- IPv6 ACL-based rate limiting is not supported on MLX Series and XMR Series devices.
- You cannot apply IPv6 PBR on a port if that port already has inbound IPv6 ACLs.
- IPv6 PBR only supports IPv6 as the next hop and IPv6 PBR to VLAN flooding.
- The number of route maps that can be defined is limited by the system memory. When a route map is used in an IPv6 PBR policy, the IPv6 PBR policy uses up to 64 instances of a route map, up to 5 ACLs in a matching policy of each route map instance.

The following two conditions can cause more than 64 route map instances to be used.

1. If one or more of the first 64 instances has a deny clause.

2. If the access list used in the first 64 instances is not configured.

    •   ACLs with the **log** option configured should not be used for IPv6 PBR purposes.

    •   IPv6 PBR ignores implicit **deny ip any any** ACL entries to ensure that traffic is compared to all the ACLs for route maps that use multiple ACLs. However, if an explicit **deny ip any any** entry is configured, traffic matching this clause will be routed normally using Layer 3 paths and will not be compared to any ACL clauses that follow this clause.

    •   IPv6 PBR always selects the first next hop from the next hop list that is up. If an IPv6 PBR policy's next hop goes down, the policy uses another next hop if available. If no next hops are available, the device routes the traffic in the normal way. No IPv6 ECMP is supported.

    •   Any changes to route map definitions will be effective immediately for the interfaces where the IPv6 PBR route map is applied, without rebinding. However, rebinding is required for an ACL definition change to take effect.

    •   PBR policies are not supported on Layer-3 VPNs.

    •   In a PBR route-map definition, if even one route-map instance contains a "set next-hop-flood-vlan" statement, all instances of that route-map will apply to both routed and switched traffic.

    •   Flooding traffic to a POS interface is not allowed. It can only be flooded to Ethernet ports on the VLAN, including the default VLAN.

    •   When an incoming port is POS then the SA of the outgoing flooded packets will be 0.

Considerations specific to CES 2000 Series and CER 2000 Series

•   IPv6 PBR only supports routed traffic. Switched traffic is not supported.

•   IPv6 PBR on transit MPLS uplinks is not supported. IPv6 PBR on an Egress MPLS interface is supported. For example, an IPv6 PBR policy is applied on an MPLS interface.

    –   In a transit router, if a 6PE packet is received on an MPLS interface, after the MPLS Label is swapped/pushed, the underlying IPv6 packet will not be subjected to PBR and will be switched based on the MPLS Label.

    –   In an egress PE router, if a 6PE packet is received on an MPLS interface, after the MPLS Labels are removed, the underlying IPv6 packet will be subjected to PBR and will be forwarded based on PBR next hop.

•   IPv6 PBR to flood VLAN is not supported for switched traffic for the CES 2000 Series and CER 2000 Series.

Configuring an IPv6 PBR with IPv6 address as the next hop

The following steps configure an IPv6 PBR by setting an IPv6 address as the next hop in the route map. This task uses Access Control Lists (ACLs). For more information about configuring ACLs, refer to the *Security Configuration Guide* for your platform.

1. Enter the **configure terminal** command to enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Define the required IPv6 ACLs to be added to the route map.

    ```
    device(config)# ipv6 access-list v6acl
    device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
    ```

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

    ```
    device(config-ipv6-access-list v6acl)# exit
    ```

4. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all the match clauses are met.

    ```
    device(config)# route-map test-route permit 99
    ```

5. Add IPv6 ACLs to match the IP address that is permitted by the ACL.

```
device(config-routemap test-route)# match ipv6 address v6acl
```

6. Set the IPv6 address of the next hop to which the traffic that matches a match statement in the route map must be routed.

```
device(config-routemap test-route)# set ipv6 next-hop 2001:db8::1
```

7. Enter the **exit** command to return to the global configuration mode.

```
device(config-routemap test-route)# exit
```

8. Enter configuration mode on the interface where you want to enable PBR by applying the route map.

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

9. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
```

The following example shows the configuration steps to configure an IPv6 PBR by setting an IPv6 address as the next hop in the route map.

```
device# configure terminal
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
device(config)# exit
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ipv6 address v6acl
device(config-routemap test-route)# set ipv6 next-hop 2001:db8::1
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
device(config-if-e1000-1/1)# end
```

The following example shows the configuration steps to configure an IPv6 PBR in which the route map is configured to preserve VLAN IDs and forward the packet to specific egress port.

```
device# configure terminal
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
device(config)# exit
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ipv6 address v6acl
device(config-routemap test-route)# set ipv6 next-hop 2001:db8::1 preserve-vlan
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/2)# allow-all-vlan pbr
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
device(config-if-e1000-1/1)# end
```

Configuring an IPv6 PBR with NULL0 interface as the next hop

The following steps configure an IPv6 PBR by setting NULL0 interface as the next hop in the route map.

1. Enter the **configure terminal** command to enter global configuration mode.

```
device# configure terminal
```

2. Define the required IPv6 ACLs to be added to the route map.

```
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
```

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

```
device(config-ipv6-access-list v6acl)# exit
```

4. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all the match clauses are met.

```
device(config)# route-map test-route permit 99
```

5. Add IPv6 ACLs to match the IP address that is permitted by the ACL.

```
device(config-routemap test-route)# match ipv6 address v6acl
```

6. Set the next hop as NULL0 interface to send the traffic to the null interface, thus dropping the packet instead of forwarding it.

```
device(config-routemap test-route)# set interface null0
```

7. Enter the **exit** command to return to the global configuration mode.

```
device(config-routemap test-route)# exit
```

8. Enter configuration mode on the interface where you want to enable PBR by applying the route map.

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

9. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
```

The following example shows the configuration steps to configure an IPv6 PBR to send the IPv6 traffic to the null interface, thus dropping the traffic instead of forwarding it.

```
device# configure terminal
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
device(config)# exit
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ipv6 address v6acl
device(config-routemap test-route)# set interface null0
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
device(config-if-e1000-1/1)# end
```

Configuring an IPv6 PBR to flood traffic on all ports of a VLAN

A PBR policy can be configured to flood the matched traffic on all ports of the VLAN except the incoming physical port using the **set next-hop-flood-vlan** statement in the route map.

The following steps configure an IPv6 PBR to flood the matched traffic on all ports of a VLAN.

1. Enter the **configure terminal** command to enter global configuration mode.

```
device# configure terminal
```

2. Define the required IPv6 ACLs to be added to the route map.

```
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
```

3. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

```
device(config)# route-map test-route permit 99
```

4. Add IPv6 ACLs to match the IP address that is permitted by the ACL.

```
device(config-routemap test-route)# match ipv6 address v6acl
```

5. Set the next hop to route the matched traffic to flood on all ports of a VLAN.

```
device(config-routemap test-route)# set next-hop-flood-vlan 10
```

   - (Optional) Set the packets that are flooded on the ports of a VLAN to carry the destination MAC address using the **outgoing-da** option.

```
device(config-routemap test-route)# set next-hop-flood-vlan 10 outgoing-da 0000.1234.5678
```

   - (Optional) Set the route map to retain the Layer 2 and Layer 3 information of the packets that are flooded to the VLAN using the **preserve-vlan** option.

```
device(config-routemap test-route)# set next-hop-flood-vlan 10 preserve-vlan
```

6. Enter the **exit** command to return to the global configuration mode.

```
device(config-routemap test-route)# exit
```

7. Enter configuration mode on the interface where you want to enable IPv6 PBR by applying the route map.

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

8. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
```

The following example shows the configuration steps to forward matching traffic to an interface.

```
device# configure terminal
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ipv6 address v6acl
device(config-routemap test-route)# set next-hop-flood-vlan 10
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
```

Configuring an IPv6 PBR with TVF domain as the next hop for TVF with LAG load balancing

The transparent VLAN flooding (TVF) domain provides an infrastructure to increase the overall egress traffic streams. By setting TVF domain as a PBR next hop, traffic can be flooded to the TVF domain that supports 2016 TVF instances with LAG load balancing and, together with 4090 TVF instances without LAG load balancing, scales the overall egress traffic flow support to 6106.

You must configure the transparent VLAN flooding (TVF) domain to be set as PBR next hop. For more information about transparent VLAN flooding domain, refer to the *Switching Configuration Guide*.

The following steps configure the TVF domain as the next hop for a route map to support transparent VLAN flooding (TVF) with LAG load balancing.

1. Enter the **configure terminal** command to enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Configure system maximum requirements to support LAGs and TVF LAG load balancing.

   ```
   device(config)# system-max trunk-num 32
   device(config)# system-max tvf-lag-lb-fid-pool 4096
   device(config)# system-max tvf-lag-lb-fid-group 4
   ```

3. Configure the TVF domain and assign the ports.

   ```
   device(config)# tvf-domain 1
   device(config-tvf-domain-1)# port ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
   device(config-tvf-domain-1)# exit
   ```

4. Define the required IPv6 ACLs to be added to the route map.

   ```
   device(config)# ipv6 access-list v6_brc_Test_p000_bi_moof
   device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# permit vlan 1001 ip any any
   ```

5. Enter the **exit** command to return to global configuration mode.

   ```
   device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# exit
   ```

6. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

   ```
   device(config)# route-map Mall permit 1001
   ```

7. Add IPv6 ACLs to match the IPv6 address that is permitted by the ACL.

   ```
   device(config-routemap Mall)# match ipv6 address v6_brc_Test_p000_bi_moof
   ```

8. Enter the **set next-hop-tvf-domain** command to configure a TVF domain as the next hop for a route map.

   ```
   device(config-routemap Mall)# set next-hop-tvf-domain 1
   ```

9. Enter the **exit** command to return to the global configuration mode.

   ```
   device(config-routemap Mall)# exit
   ```

10. Configure a VLAN and add the interfaces required for the TVF LAG and LAG load balancing.

    ```
    device(config)# vlan 1 name DEFAULT-VLAN
    device(config-vlan-1)# no untagged ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
    device(config-vlan-1)# no spanning-tree
    ```

11. Enter the **exit** command to return to the global configuration mode.

    ```
    device(config-vlan-1)# exit
    ```

12. Create a LAG and add the interfaces.

    ```
    device(config)# lag blue static
    device(config-lag-blue)# ports ethernet 5/1 ethernet 6/1
    device(config-lag-blue)# primary-port 5/1
    device(config-lag-blue)# deploy
    ```

13. Enter the **exit** command to return to the global configuration mode.

```
device(config-lag-blue)# exit
```

14. Enter configuration mode on the interface where you want to enable IPv6 PBR by applying the route map.

```
device(config)# interface ethernet 11/1
```

15. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-11/1)# ipv6 policy route-map Mall
```

The following example shows the configuration steps to configure the TVF domain as the next hop for a route map to support transparent VLAN flooding (TVF) with LAG load balancing.

```
device# configure terminal
device(config)# system-max trunk-num 32
device(config)# system-max tvf-lag-lb-fid-pool 4096
device(config)# system-max tvf-lag-lb-fid-group 4
device(config)# tvf-domain 1
device(config-tvf-domain-1)# port ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-tvf-domain-1)# exit
device(config)# ipv6 access-list v6_brc_Test_p000_bi_moof
device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# permit vlan 1001 ip any any
device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# exit
device(config)# route-map Mall permit 1001
device(config-routemap Mall)# match ipv6 address v6_brc_Test_p000_bi_moof
device(config-routemap Mall)# set next-hop-tvf-domain 1
device(config-routemap Mall)# exit
device(config)# vlan 1 name DEFAULT-VLAN
device(config-vlan-1)# no untagged ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-vlan-1)# no spanning-tree
device(config-vlan-1)# exit
device(config)# lag blue static
device(config-lag-blue)# ports ethernet 5/1 ethernet 6/1
device(config-lag-blue)# primary-port 5/1
device(config-lag-blue)# deploy
device(config-lag-blue)# exit
device(config)# interface ethernet 11/1
device(config-if-e1000-11/1)# ipv6 policy route-map Mall
```

Configuring IPv6 PBR with next hop type as GRE tunnel

IPv6 traffic can be redirected over a GRE tunnel using IPv6 PBR. This enables communication between IPv6 end points without upgrading IPv4 network between them. Tunnels can be configured between end devices, however both tunnel endpoints must support IPv4 and IPv6 protocol stacks. The following steps configure a GRE tunnel interface as the next hop using PBR.

You must configure GRE tunnel before configuring the IPv6 traffic to be redirected over the tunnel using IPv6 PBR. For more information about GRE tunnel configuration, refer to the *Extreme NetIron Layer 3 Routing Configuration Guide*.

1. Enter the **configure terminal** command to enter global configuration mode.

```
device# configure terminal
```

2. Define the required IPv6 ACL that contains the matching criteria for the IPv6 traffic that needs to be routed using PBR.

```
device(config)# ipv6 access-list v6_brc_Test_p000_bi_moof
device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# permit ipv6 any any
```

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

```
device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# exit
```

4. Enter the **route-map** command to define the route map that matches on the ACLs and sets next hop as GRE tunnel.

```
device(config)# route-map Mall permit 1001
```

5. Add IPv6 ACLs to match the IPv6 address that is permitted by the ACL.

```
device(config-routemap Mall)# match ipv6 address v6_brc_Test_p000_bi_moof
```

6. Set the configured tunnel as the next hop for the route map.

```
device(config-routemap Mall)# set next-hop-ip-tunnel 1
```

7. Enter the **exit** command to return to the global configuration mode.

```
device(config-routemap Mall)# exit
```

8. Enter configuration mode on the interface where you want to enable IPv6 PBR by applying the route map.

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

9. Apply the route map to an interface.

```
device(config-if-e1000-1/1)# ipv6 policy route-map Mall
```

The following example shows the configuration steps to redirect IPv6 traffic over a GRE tunnel using PBR.

```
device(config)# interface tunnel 1
device(config-tnif-1)# tunnel source ethernet 1/1
device(config-tnif-1)# tunnel destination 10.2.2.1
device(config-tnif-1)# tunnel mode gre ip
device(config-tnif-1)# ipv6 address 2005:db8::1/64
device(config-tnif-1)# exit
device(config)# ipv6 access-list v6_brc_Test_p000_bi_moof
device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# permit ipv6 any any
device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# exit
device(config)# route-map Mall permit 1001
device(config-routemap Mall)# match ipv6 address v6_brc_Test_p000_bi_moof
device(config-routemap Mall)# set next-hop-ip-tunnel 1
device(config-routemap Mall)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ipv6 policy route-map Mall
```

Configuring IPv6 PBR with next hop type as LSP

IPv6 traffic can be redirected over a MPLS tunnel using IPv6 PBR. This enables communication between IPv6 end points without upgrading IPv4 network between them. Tunnels can be configured between end devices, however both tunnel endpoints must support IPv4 and IPv6 protocol stacks. The following steps must be configured to redirect IPv6 traffic with next hop type as LSP.

MPLS tunnel must be established before configuring next hop type as LSP. For more information about LSP configuration, refer to the *Extreme NetIron MPLS Configuration Guide*.

1. Enter the **configure terminal** command to enter global configuration mode.

```
device# configure terminal
```

2. Define the required IPv6 ACLs that contains the matching criteria for the IPv6 traffic that needs to be routed using PBR.

```
device(config)# ipv6 access-list v6_brc_Test_p000_bi_moof
device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# permit ipv6 any any
```

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

    ```
    device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# exit
    ```

4. Enter the **route-map** command to define the route map that matches on the ACLs and sets next hop as MPLS tunnel.

    ```
    device(config)# route-map Mall permit 1001
    ```

5. Add IPv6 ACLs to match the IPv6 address that is permitted by the ACL.

    ```
    device(config-routemap Mall)# match ipv6 address v6_brc_Test_p000_bi_moof
    ```

6. Set the configured LSP as the next hop for the route map.

    ```
    device(config-routemap Mall)# set next-hop-lsp t3
    ```

7. Enter the **exit** command to return to the global configuration mode.

    ```
    device(config-routemap Mall)# exit
    ```

8. Enter configuration mode on the interface where you want to enable IPv6 PBR by applying the route map.

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

9. Apply the route map to an interface.

    ```
    device(config-if-e1000-1/1)# ipv6 policy route-map Mall
    ```

The following example shows the configuration steps to redirect IPv6 traffic over a MPLS tunnel using PBR.

```
device(config)# router mpls
device(config-mpls)# mpls-interface ethernet 6/1
device(config-mpls)# lsp t3
device(config-mpls-lsp-t3)# to 10.1.1.1
device(config-mpls-lsp-t3)# enable
device(config-mpls-lsp-t3)# exit
device(config)# ipv6 access-list v6_brc_Test_p000_bi_moof
device(config-ipv6-access-list v6_brc_Test_p000_bi_moof)# permit ipv6 any any
device(config)# route-map Mall permit 1001
device(config-routemap Mall)# match ipv6 address v6_brc_Test_p000_bi_moof
device(config-routemap Mall)# set next-hop-lsp t3
device(config-routemap Mall)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ipv6 policy route-map Mall
```

Configuration examples for policy based routing

This section presents configuration examples for configuring and applying a PBR policy.

Basic example of IPv6 PBR

The following commands configure and apply an IPv6 PBR policy that routes HTTP traffic received on a virtual routing interface.

```
device# configure terminal
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
device(config)# exit
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ipv6 address v6acl
device(config-routemap test-route)# set ipv6 next-hop 2001:db8::1
```

```
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
device(config-if-e1000-1/1)# end
```

Combined example of IPv4 and IPv6 PBR

If both IPv4 and IPv6 configurations exist in a route map, the IPv4 portions of the route map will be ignored when the route map is used for IPv6 PBR, and the IPv6 portions of the route map will be ignored when the route map is used for IPv4 PBR. The following example configures commands configure and apply an IPv6 PBR policy that routes HTTP traffic received on a virtual routing interface.

In the following example the IPv4 traffic that matches the route map will be sent to 10.1.1.1 and IPv6 traffic that matches the route map will be sent to 2001:db8::1.

```
device# configure terminal
device(config)# access-list 99 permit 10.157.23.0 0.0.0.255
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
device(config)# exit
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match ip address 99
device(config-routemap test-route)# match ipv6 address v6acl
device(config-routemap test-route)# set ip next-hop 10.1.1.1
device(config-routemap test-route)# set ipv6 next-hop 2001:db8::1
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ip policy route-map test-route
device(config-if-e1000-1/1)# ipv6 policy route-map test-route
device(config-if-e1000-1/1)# end
```

Selectively applying normal routing to packets

Certain situations demand selected traffic to undergo normal routing based on IPv6 RTM while other traffic is to be forwarded based on PBR. One such situation is provided in the following scenario. The example provides a recommended solution.

To communicate with a direct connected host, the IPv6 address has to be resolved. To resolve the IPv6 address, an ICMPv6 Neighbor Solicitation (NS) will be sent and an ICMPv6 Neighbor Advertisement (NA) is expected as a reply. If the match ACL rules matches the IPv6 address of ICMPv6 NA, as per the PBR route map, it will be redirected to the first reachable next hop and the IPv6 address is not resolved. This results in packets being dropped. To avoid this situation, use the **deny** ACL clause to selectively allow ICMP traffic and define it as a permit route map entry, as shown in the following example.

```
device# configure terminal
device(config)# ipv6 access-list v6acl
device(config-ipv6-access-list v6acl)# deny icmp 2001:db8::/64 any
device(config-ipv6-access-list v6acl)# permit ipv6 2001:db8::/64 any
device(config-ipv6-access-list v6acl)# exit
device(config)# route-map v6pbr permit 10
device(config-routemap v6pbr)# match ipv6 address v6acl
device(config-routemap v6pbr)# set next-hop-flood-vlan 200 outgoing-da 0000.0022.3333
device(config-routemap v6pbr)# exit
device(config-routemap v6pbr)# interface ethernet 2/2
device(config-if-e10000-2/2)# ipv6 address 2001:db8::1/64
device(config-if-e10000-2/2)# ipv6 enable
device(config-if-e10000-2/2)# ipv6 policy route-map v6pbr
```

To allow normal routing for certain traffic, a corresponding deny ACL filter can be added before any permit ACL filter.

Displaying IPv6 PBR route map information

Various show commands can be used to display IPv6 PBR configuration details.

Enter the **show route-map** command to display the route map information.

```
device# show route-map ipv6_pbr_map
route-map ipv6_pbr_map permit 10
    match ipv6 address for_pbr_match_src
    set ipv6 next-hop 2001:db8::1
```

Enter the **show route-map binding** command to display the route map binding information.

```
device# show route-map binding ipv6_pbr_map
    IPv6 Bindings of ipv6_pbr_map :
        1/1
```

Displaying IPv6 ACL and route map information on the CES 2000 Series and CER 2000 Series

Enter the **show ipv6 access-list** command to display the IPv6 ACL output on the CES 2000 Series and CER 2000 Series.

```
device# show ipv6 access-list ipv6acl
ipv6 access-list ipv6acl: 3 entries
10: permit enable-accounting ipv6 2001:db8:3::/64 2001:db8:4::/64
20: permit ipv6 2001:db8:3::/64 any
30: permit ipv6 any 2001:db8:4::/64
```

Enter the **show route-map** command to display the route map output on the CES 2000 Series and CER 2000 Series.

```
device# show route-map pbr-1
route-map pbr-1 permit 5
match ipv6 address ipv6acl
set ipv6 next-hop 2001:db8:5::2
set ipv6 next-hop 2001:db8:5::3
set ipv6 next-hop 2001:db8:4::2
```

Enter the **show pbr** command to display the PBR selected next hop information on the CES 2000 Series and CER 2000 Series.

```
device# show pbr interface ve 20
Interface VE 20
No PBR routemap
IPv6 PBR Routemap pbr-1
PPCR 1:1: Routemap instance 5
Valid 1 ACL ipv6acl First table index 10
Metro Specific information in PPCR: 1:1
Instance No: 5
SET information
Next-Hop To NH value PumaNextStage SW NextStage id HW NextStage id Resolved
IPv6 address 2001:db8:5::2 IPv6 NEXT HOP 0x02000020 32 Y
IPv6 address 2001:db8:5::3 IPv6 NEXT HOP 0xFFFFFFFF -1 N
IPv6 address 2001:db8:4::2 IPv6 NEXT HOP 0x02000021 33 Y
```

The SET information table provides the hop status list. The ACL uses the first resolved next hop.

Displaying IPv6 accounting information

IPv6 PBR accounting is supported through the use of IPv6 ACL accounting in XMR Series and MLX Series devices.

ACL counters must be enabled to view accounting data.

Enter the **ipv6 access-list accounting brief policy-based-routing** command to display IPv6 PBR accounting details.

```
device(config)# enable-acl-counter
device# show ipv6 access-list accounting brief policy-based-routing
IPv6 Policy Based Routing Accounting Summary: (ac = accumulated since accounting started)
    Int    In ACL              Total In Hit   Out ACL             Total Out Hit
```

```
      2/6v6ACL1-udp                    0(1s)
                                       0(1m)
                                       0(5m)
                                       0(ac)
      2/6v6ACL1-tcp                    0(1s)
                                       0(1m)
                                       0(5m)
                                       0(ac)
      2/6ipv6-src-d                    0(1s)
                                       0(1m)
                                       0(5m)
                                       0(ac)
   VE 40v6ACL1-udp                 11659(1s)
                                  700150(1m)
                                 3500549(5m)
                               160277093(ac)
   VE 40v6ACL1-tcp                127319(1s)
                                 7645985(1m)
                                38227679(5m)
                              1750303054(ac)
```

Displaying telemetry command

The **show telemetry** command displays the information related to the telemetry configuration.

```
device(config#) show telemetry rule-name
Paths with leading * are configured but disabled, entries with + is for IPv6
Route-map ACL Output Output
Rule Name Input Policy Match VLAN Port(s)
*test-traffic N/A Testing TestACL N/A N/A
*xGW_path1 N/A other_map xGW_Filter2 N/A N/A
xGW_path1 10/1 10/2 xGW_map xGW_Filter1 2 N/A
xGW_path1 +10/1 xGW_map xGW_Filter1 2 N/A
xGW_path2 10/1 10/2 xGW_map xGW_Filter2 None 4/1
xGW_path2 +10/1 xGW_map None 4/1
show telemetry rule-name test_traffic
Paths with leading * are configured but disabled, entries with + is for IPv6
Route-map ACL Output Output
Rule Name Input Policy Match VLAN Port(s)
*test-traffic N/A TestMap TestACL N/A N/A
```

The following command provides a detail mode output in which the list of ports will be fully expanded and displayed if the ports are LAG or VLAN ports. For the output VLAN and output ports information, the command will show the active ports or VLANs that are currently being used for traffic forwarding.

```
device(config#) show telemetry detail rule-name
Rule name: test-traffic (disabled)
Input: None
Route-map Policy: Testing
IPv4 ACL match: Test_filter1
Output: None
Rule Name: xGW_Path1
Input: IPv4 - 10/1 10/2 IPv6 - 10/1
Route-map Policy: xGW_map
IPv4 ACL match: xGW_Filter1
IPv6 ACL match: xGW_Filter1
Output: IPv4 - VLAN 2 IPv6 - VLAN 2
Rule Name: xGW_Path2
Input: IPv4 - 10/1, 10/2 IPv6 - 10/1
Route-map Policy: xGW_map
IPv4 ACL match: xGW_Filter2
Output: IPv4 - 2/2 4/1 IPv6 - 2/2 4/1
```

Layer 2 policy-based routing

Layer 2 PBR allows you to manually configure how IPv6 packets that match certain Layer 2 criteria can be forwarded as per the route map configuration. This section provides configuration details of Layer PBR policy that specifies to match and route packets based on the routing policy.

Layer 2 PBR configuration considerations

Pay attention to the following configuration considerations for Layer 2 PBR:

- Layer 2 PBR cannot be applied globally. Layer 2 PBR can be applied only at the physical interface level.
- If both Layer 2 PBR and Layer 3 PBR are applied on the same interface (or Layer 3 PBR is applied globally), Layer 2 PBR only filters non-IP packets.
- If both Layer 2 PBR and Layer 3 PBR are applied on the same interface on BR-MLX-10Gx24 and BR-MLX-40Gx4-X cards, only Layer 3 PBR will work. In this case, Layer 3 PBR filters only IP packets.
- If only Layer 2 PBR is applied, both IP and non-IP packets will be filtered.
- Layer 2 PBR cannot be applied on a VE interface.
- Layer 2 PBR can be configured only on physical ports and Link Aggregation Group (LAG) ports.
- Layer 2 PBR cannot be applied on an interface where Layer 2 ACL or Layer 3 ACL is already applied.
- A Layer 2 ACL with the **log** or **mirror** option configured should not be used in Layer 2 PBR. If used, the **log** or **mirror** configurations are ignored.
- Layer 2 PBR cannot be applied on an interface where ACL-based rate limiting is already applied.
- Any changes to route maps or ACL definitions will be effective immediately for the interfaces where Layer 2 PBR is applied. There is no need to rebind using the **ip rebind-acl** command to change the CAM programming.
- Deletion of a Layer 2 ACL used in an applied route map is not allowed, unless the **force-delete-bound-acl** command is configured.
- When a route map is used, the Layer 2 PBR policy uses up to five ACLs in a matching policy of each route map instance.
- If Layer 2 PBR is applied using the **l2 policy route-map** command, only the options to set the next hop to a null interface or physical interface and VLAN flooding are selected. Other **set** statements configured in the route map configuration are ignored.
- The number of route maps that you can define is limited by the system memory. When a route map is used in a Layer 2 PBR policy, the Layer 2 PBR policy uses up to 1000 instances of a Layer 2 route map and up to five ACLs in a matching policy of each route map instance.
- The following two conditions can cause more than 1000 Layer 2 route map instances to be used:
  - If one or more of first 1000 instances have a **deny** clause.
  - If the ACL used in the first 1000 instances is not configured.
- Layer 2 PBR ignores implicit **deny any any** ACL entries to ensure that, for route maps that use multiple ACLs, the traffic is compared to all the ACLs. However, if an explicit **deny any any** ACL entry is configured, traffic matching this clause is not compared to any ACL clauses that follow this clause.
- A Layer 2 PBR route map cannot be applied on VPLS, VLL, or VLL-Local endpoints and vice versa.
- Layer 2 PBR policies are not supported on Layer 3 VPNs.

Configuring a Layer 2 PBR with an interface as the next hop

Within a route-map instance, there can be only one match statement but multiple interfaces can be set as next hops. If there are multiple interfaces configured in the **set** statement, the first one in the configured order will be used to forward traffic. If the actively used interface is down, the next interface in the configuration order will take over. If the previously down interface comes back up, the traffic will be reverted to the first interface.

The following steps configure a Layer 2 PBR to forward matching traffic to an interface.

1. Enter the **configure terminal** command to enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Define the required Layer 2 ACLs to be added to the route map.

   ```
   device(config)# mac access-list abc
   device(config-mac-acl-abc)# permit any any any etype 8000
   device(config-mac-acl-abc)# exit
   ```

3. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

   ```
   device(config)# route-map test-route permit 99
   ```

4. Add Layer 2 ACLs that specify all the conditions required to match the desired packets.

   ```
   device(config-routemap test-route)# match l2acl abc
   ```

5. Set an interface as the next hop for the route map and preserve VLAN and Layer 2 information for the matching packets.

   ```
   device(config-routemap test-route)# set interface ethernet 3/1 preserve-vlan
   ```

6. Enter the **exit** command to return to the global configuration mode.

   ```
   device(config-routemap test-route)# exit
   ```

7. Enter configuration mode on the interface where you want to enable Layer 2 PBR by applying the route map.

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

8. Enable policy-based routing on the interface and specify the route map to be used.

   ```
   device(config-if-e1000-1/1)# l2 policy route-map test-route
   ```

The following example shows the configuration steps to forward matching traffic to an interface.

```
device# configure terminal
device(config)# mac access-list abc
device(config-mac-acl-abc)# permit any any any etype 8000
device(config-mac-acl-abc)# exit
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match l2acl abc
device(config-routemap test-route)# set interface ethernet 3/1 preserve-vlan
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# l2 policy route-map test-route
```

Configuring a Layer 2 PBR with NULL0 interface as the next hop

The following steps configure a Layer 2 PBR by setting NULL0 interface as the next hop in the route map.

1. Enter the **configure terminal** command to enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Define the required Layer 2 ACLs to be added to the route map.

    ```
    device(config)# mac access-list abc
    device(config-mac-acl-abc)# permit any any any etype 8000
    ```

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

    ```
    device(config-mac-acl-abc)# exit
    ```

4. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all the match clauses are met.

    ```
    device(config)# route-map test-route permit 99
    ```

5. Add Layer 2 ACLs that specify all the conditions required to match the desired packets.

    ```
    device(config-routemap test-route)# match l2acl abc
    ```

6. Set the next hop as NULL0 interface to send the traffic to the null interface, thus dropping the packet instead of forwarding it.

    ```
    device(config-routemap test-route)# set interface null0
    ```

7. Enter the **exit** command to return to the global configuration mode.

    ```
    device(config-routemap test-route)# exit
    ```

8. Enter configuration mode on the interface where you want to enable PBR by applying the route map.

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

9. Enable policy-based routing on the interface and specify the route map to be used.

    ```
    device(config-if-e1000-1/1)# l2 policy route-map test-route
    ```

The following example shows the configuration steps to configure a Layer 2 PBR to send the traffic to the null interface, thus dropping the traffic instead of forwarding it.

```
device# configure terminal
device(config)# mac access-list abc
device(config-mac-acl-abc)# permit any any any etype 8000
device(config-mac-acl-abc)# exit
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match l2acl abc
device(config-routemap test-route)# set interface null0
device(config-routemap test-route)# exit
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# l2 policy route-map test-route
```

VLAN flooding with PBR

VLAN flooding using PBR allows you to flood the matched traffic on all ports of the VLAN except the incoming physical port.

> **NOTE**
> Always use Transparent VLAN Flooding for the VLAN that is specified in the **set next-hop-flood-vlan** *vlan-id* command.

If the VLAN specified in the route map is not configured, the PBR route map set statement will fall through to the next configured set statement. If no valid next hop is available, the packet is forwarded as per Layer 2 or Layer 3 forwarding decision. If the specified VLAN has no valid outgoing ports, (such as when all ports in the VLAN are down or when the VLAN is empty) matching packets will be dropped. There is also no Layer 3 header processing. For example, in the case of an IPv6 packet header, the hop-limit will not be decremented.

Optionally, you can set the device to send packets flooded to the ports on the VLAN with the destination MAC address using the **outgoing-da** option. If the destination MAC address is not set using the **outgoing-da** option, the destination address is set as described in the following table.

TABLE 42 Destination address on VLAN flooded packets

Incoming port	Outgoing port	Routed traffic	Switched traffic
Ethernet	Ethernet	Replaced Destination Address	Original Destination Address from Incoming Packet

The behavioral differences when deployed on a CES 2000 Series or CER 2000 Series device are described in the following table.

TABLE 43 Behavioral differences of VLAN flooding per platform

Scenario	MLX Series/ XMR Series
If the outgoing-da option is not set	Packet floods with original SA/DA.
	Routed traffic (CAM miss):
	Packet floods with original SA/DA.
	Routed traffic (CAM hit):
	Packet floods with original SA and DA as next-hop MAC.
If the outgoing-da option is set	DA: Configured outgoing-da
	SA: Original SA
In the **set next-hop-flood-vlan** command, if the specified VLAN ID is configured but no ports are added or all ports are down	Drop matching packets.
Ingress VLAN and PBR Flood VLAN are the same.	Packets are not flooded back to source port.

Policy based routing with the preserve VLAN option

The Layer 2 and Layer 3 information of the packets that matches PBR policy can be retained (for example, the VLAN information and the MAC address are retained) using the **preserve-vlan** option in the set statement of the route map.

The **preserve-vlan** option can be configured in the following route map set policy configuration:

- **set next-hop-flood-vlan**
- **set interface ethernet** .

> **NOTE**
> Policy-based routing support for preserve VLAN is supported only on XMR Series, MLX Series routers, and on the BR-MLX-10GX24 module.

The interface, on which PBR with preserve VLAN is configured, should be part of the VLANs through which packets are expected. TTL is not decremented when the packet is sent to the configured next hop. A route map policy with set policies to preserve VLAN can be applied on a physical port or on a VE port. The route map with preserve-vlan set policies cannot be configured globally. A route map used for PBR with a preserve VLAN policy must have the **preserve-vlan** option configured for each set policy. Packets not matching the PBR policy with preserve-vlan will be dropped. If none of the policy's direct routes or next hops is available, the packets are forwarded as per the routing table.

Previously, PBR transparent VLAN flooding (TVF) replaced the ingress traffic's VLAN ID with the egress TVF VLAN ID, while flooding the egress TVF VLAN. Policy-based routing support for preserve VLAN allows the ingress packet VLAN header (VLAN ID and priority) to be preserved, while simultaneously flooding the PBR TVF VLAN.

The PBR TVF VLAN egress ports can be in strict tagged VLAN mode or dual VLAN mode. When PBR TVF VLAN egress ports are in strict tagged VLAN mode, the ingress tagged packets flood as "tagged" with the original VLAN ID and priority preserved. The ingress untagged packets flood as "tagged" with the default VLAN ID. When the PBR TVF VLAN egress ports are in dual VLAN mode, the ingress tagged packets flood as "tagged" with the original VLAN ID and priority preserved. The ingress untagged packets flood as "untagged" if the egress port is the untagged member of the ingress port's default VLAN; otherwise, the ingress untagged packets flood as "tagged" with the original VLAN ID and priority preserved.

The **allow-all-vlan pbr** command configures a physical interface to accept all VLAN packets for the purpose of PBR. This command reduces configuration complexity since the physical interface does not have to be configured individually in multiple VLAN interfaces. The **allow-all-vlan pbr** command cannot be configured when the physical port is configured with an IPv4 address, IPv6 address, MPLS, VPLS, VLL, ICL, Layer 3 VPN; or when the port is part of other VLAN. The **allow-all-vlan pbr** command cannot be applied to a VE.

Configuring a Layer 2 PBR to flood traffic on all ports of a VLAN

A Layer 2 PBR policy can be configured to flood the matched traffic on all ports of the VLAN except the incoming physical port using the **set next-hop-flood-vlan** statement in the route map.

The following steps configure a Layer 2 PBR to flood the matched traffic on all ports of a VLAN.

1. Enter the **configure terminal** command to enter global configuration mode.

   ```
   device# configure terminal
   ```

2. Define the required Layer 2 ACLs to be added to the route map.

   ```
   device(config)# mac access-list abc
   device(config-mac-acl-abc)# permit any any any etype 8000
   device(config-mac-acl-abc)# exit
   ```

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

   ```
   device(config-mac-acl-abc)# exit
   ```

4. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

   ```
   device(config)# route-map test-route permit 99
   ```

5. Add Layer 2 ACLs that specify all the conditions required to match the desired packets.

   ```
   device(config-routemap test-route)# match l2acl abc
   ```

6. Set the next hop to route the matched traffic to flood on all ports of a VLAN.

    ```
    device(config-routemap test-route)# set next-hop-flood-vlan 10
    ```

    •   (Optional) Set the packets that are flooded on the ports of a VLAN to carry the destination MAC address using the **outgoing-da** option.

        ```
        device(config-routemap test-route)# set next-hop-flood-vlan 10 outgoing-da 0000.1234.5678
        ```

    •   (Optional) Set the route map to retain the Layer 2 information and the ingress packet VLAN header (VLAN ID and priority) of the packets that are flooded to the VLAN using the **preserve-vlan** option.

        ```
        device(config-routemap test-route)# set next-hop-flood-vlan 10 preserve-vlan
        ```

7. Enter the **exit** command to return to the global configuration mode.

    ```
    device(config-routemap test-route)# exit
    ```

8. Enter configuration mode on the interface where you want to enable Layer 2 PBR by applying the route map.

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

9. Enable policy-based routing on the interface and specify the route map to be used.

    ```
    device(config-if-e1000-1/1)# l2 policy route-map test-route
    ```

The following example shows the configuration steps to forward matching traffic to an interface.

```
device# configure terminal
device(config)# mac access-list abc
device(config-mac-acl-abc)# permit any any any etype 8000
device(config-mac-acl-abc)# exit
device(config)# route-map test-route permit 99
device(config-routemap test-route)# match l2acl abc
device(config-routemap test-route)# set next-hop-flood-vlan 10
device(config-routemap test-route)# exit
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# l2 policy route-map test-route
```

Configuring a Layer 2 PBR with TVF domain as the next hop for the TVF LAG load balancing

The transparent VLAN flooding (TVF) domain provides an infrastructure to increase the overall egress traffic streams. By setting TVF domain as a PBR next hop, traffic can be flooded to the TVF domain that supports 2016 TVF instances with LAG load balancing and, together with 4090 TVF instances without LAG load balancing, scales the overall egress traffic flow support to 6106.

You must configure the transparent VLAN flooding (TVF) domain to be set as PBR next hop. For more information about transparent VLAN flooding domain, refer to the *Switching Configuration Guide* for your platform.

The following steps configure the TVF domain as the next hop for a route map to support transparent VLAN flooding (TVF) with LAG load balancing.

1. Enter the **configure terminal** command to enter global configuration mode.

    ```
    device# configure terminal
    ```

2. Configure system maximum requirements to support LAGs and TVF LAG load balancing.

    ```
    device(config)# system-max trunk-num 32
    device(config)# system-max tvf-lag-lb-fid-pool 4096
    device(config)# system-max tvf-lag-lb-fid-group 4
    ```

3. Configure the TVF domain and assign the ports.

```
device(config)# tvf-domain 1
device(config-tvf-domain-1)# port ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-tvf-domain-1)# exit
```

4. Define the required Layer 2 ACLs to be added to the route map.

```
device(config)# mac access-list abc
device(config-mac-acl-abc)# permit any any any etype 8000
device(config-mac-acl-abc)# exit
```

5. Enter the **exit** command to return to global configuration mode.

```
device(config-mac-acl-abc)# exit
```

6. Enter the **route-map** command to define the route and specify the match criteria and the resulting action if all of the match clauses are met.

```
device(config)# route-map Mall permit 1001
```

7. Add Layer 2 ACLs that specify all the conditions required to match the desired packets.

```
device(config-routemap Mall)# match l2acl abc
```

8. Enter the **set next-hop-tvf-domain** command to configure a TVF domain as the next hop for a route map.

```
device(config-routemap Mall)# set next-hop-tvf-domain 1
```

9. Enter the **exit** command to return to the global configuration mode.

```
device(config-routemap Mall)# exit
```

10. Configure a VLAN and add the interfaces required for the TVF LAG and LAG load balancing.

```
device(config)# vlan 1 name DEFAULT-VLAN
device(config-vlan-1)# no untagged ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-vlan-1)# no spanning-tree
```

11. Enter the **exit** command to return to the global configuration mode.

```
device(config-vlan-1)# exit
```

12. Create a LAG and add the interfaces.

```
device(config)# lag blue static
device(config-lag-blue)# ports ethernet 5/1 ethernet 6/1
device(config-lag-blue)# primary-port 5/1
device(config-lag-blue)# deploy
```

13. Enter the **exit** command to return to the global configuration mode.

```
device(config-lag-blue)# exit
```

14. Enter configuration mode on the interface where you want to enable Layer 2 PBR by applying the route map.

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

15. Enable policy-based routing on the interface and specify the route map to be used.

```
device(config-if-e1000-1/1)# l2 policy route-map test-route
```

The following example shows the configuration steps to configure the TVF domain as the next hop for a route map to support transparent VLAN flooding (TVF) with LAG load balancing.

```
device# configure terminal
device(config)# system-max trunk-num 32
device(config)# system-max tvf-lag-lb-fid-pool 4096
device(config)# system-max tvf-lag-lb-fid-group 4
device(config)# tvf-domain 1
device(config-tvf-domain-1)# port ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-tvf-domain-1)# exit
device(config)# mac access-list abc
device(config-mac-acl-abc)# permit any any any etype 8000
device(config-mac-acl-abc)# exit
device(config)# route-map Mall permit 1001
device(config-routemap Mall)# match l2acl abc
device(config-routemap Mall)# set next-hop-tvf-domain 1
device(config-routemap Mall)# exit
device(config)# vlan 1 name DEFAULT-VLAN
device(config-vlan-1)# no untagged ethernet 1/1 ethernet 2/1 ethernet 5/1 ethernet 6/1
device(config-vlan-1)# no spanning-tree
device(config-vlan-1)# exit
device(config)# lag blue static
device(config-lag-blue)# ports ethernet 5/1 ethernet 6/1
device(config-lag-blue)# primary-port 5/1
device(config-lag-blue)# deploy
device(config-lag-blue)# exit
device(config)# interface ethernet 11/1
device(config-if-e1000-11/1)# l2 policy route-map Mall
```

Examples for managing MPLS, MAC-in-MAC, ARP, and LACP traffic types using Layer 2 PBR

The following methods can be used to match different traffic types using Layer 2 PBR.

TABLE 44 Matching criteria for different traffic types and supported line cards

Traffic type - Matching criteria	BR-MLX-1GCx24-X BR-MLX-1GFx24-X BR-MLX-10Gx4-X NI-MLX-1Gx48-T NI-MLX-1Gx48-T-A (Gen-1.1 cards)	NI-MLX-10Gx8-M NI-MLX-10Gx8-D BR-MLX-10Gx8-X (Gen-2 cards)	BR-MLX-10Gx24	BR-MLX-40Gx4-X BR-MLX-10Gx20 BR-MLX-100Gx2-CFP2 BR-MLX-10Gx4-M-IPSEC (Gen 2+ cards)

MPLS UNICAST - Etype 0x8847	Not supported	Use numerical etype matching
		
		`device(config)# mac access-list x2`
		`device(config-mac-acl-x2)#`
		`permit any any any etype 8847`
			Not supported	Use numerical etype matching
				
				`device(config)# mac access-list x2`
				`device(config-mac-acl-x2)#`
				`permit any any any etype 8847`

MPLS MULTICAST - Etype 0x8848	Not supported	Use numerical etype matching
		
		`device(config)# mac access-list x2`
		`device(config-mac-acl-x2)#`
		`permit any any any etype 8848`
			Not supported	Use numerical etype matching
				
				`device(config)# mac access-list x2`
				`device(config-mac-acl-x2)#`
				`permit any any any etype 8848`

TABLE 44 Matching criteria for different traffic types and supported line cards (continued)

Traffic type - Matching criteria	BR-MLX-1GCx24-X BR-MLX-1GFx24-X BR-MLX-10Gx4-X NI-MLX-1Gx48-T NI-MLX-1Gx48-T-A (Gen-1.1 cards)	NI-MLX-10Gx8-M NI-MLX-10Gx8-D BR-MLX-10Gx8-X (Gen-2 cards)	BR-MLX-10Gx24	BR-MLX-40Gx4-X BR-MLX-10Gx20 BR-MLX-100Gx2-CFP2 BR-MLX-10Gx4-M-IPSEC (Gen 2+ cards)
MAC In MAC - Etype 0x88A8	Not supported	Use numerical etype matching ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any any any etype 88A8```	Not supported	Use numerical etype matching ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any any any etype 88A8```
ARP - Etype 0x0806	Use etype arp to match etype for arp. ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any any any etype arp```	Use any of the following methods: • Use numerical etype matching • Use etype arp to match etype for arp ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any any any etype arp```	Use etype arp to match etype for arp. ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any any any etype arp```	Use any of the following methods: • Use numerical etype matching • Use etype arp to match etype for arp ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any any any etype arp```
LACP - MAC DA (01-80-C2-00-00-02)	Use MAC DA as matching criteria Example: ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any 0180.c200.0002 ffff.ffff.ffff```	Use MAC DA as matching criteria Example: ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any 0180.c200.0002 ffff.ffff.ffff```	Use MAC DA as matching criteria Example: ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any 0180.c200.0002 ffff.ffff.ffff```	Use MAC DA as matching criteria Example: ```device(config)# mac access-list x2 device(config-mac-acl-x2)# permit any 0180.c200.0002 ffff.ffff.ffff```

Upgrade and downgrade considerations

NetIron R05.8.00b supports Layer 2 PBR configuration in which matching Layer 2 ACLs is supported. If the device is downgraded to an earlier version, the configuration to match Layer 2 ACLs in the route map is removed. Layer 2 PBR that is applied on any interface is also removed.

Configuration examples

The following commands configure and apply a PBR policy that routes Layer 2 traffic that is flooded on VLAN 2 on physical interface 1/1.

```
device(config)# mac access-list abc
device(config-mac-acl-abc)# permit any any any etype 8000
device(config-mac-acl-abc)# exit
device(config)# route-map xGW_map permit 1
device(config-routemap xGW_map)# match l2acl abc
```

```
device(config-routemap xGW_map)# set next-hop-flood-vlan 2
device(config-routemap xGW_map)# exit
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# l2 policy route-map xGW_map
```

The following commands configure and apply a PBR policy that routes Layer 2 traffic that is flooded on VLAN 2 on physical interface 1/1. The example shows that you can have IPv4, IPv6 and Layer 2 ACLS in the route map configuration with the IPv4 address, IPv6 address, and VLAN as the next hop. If you apply Layer 2 policy, the route map is used only for Layer 2 PBR and the IPv4 and IPv6 portions of the route map are ignored. If you apply Layer 2 policy, only the Layer 2 ACL is programmed in the Layer 2 CAM.

You can apply IPv4 PBR and IPv6 PBR on the same interface along with Layer 2 PBR policy. If you apply the route map as an IPv4 policy using the **ip policy route-map** command, only the IPv4 ACL (101, 102) will be programmed in the IPv4 CAM.

If you apply the route map as an IPv6 policy using the **ipv6 policy route-map** command, only the IPv6 ACL (xGW_Filter1) will be programmed in the IPv6 CAM.

```
device(config)# access-list 101 permit ip 11.1.1.11/32 any
device(config)# access-list 102 permit ip  10.1.1.11/32 any
device(config)# ipv6 access-list xGW_Filter1
device(config-ipv6-access-list xGW_Filter1)# permit ipv6 any any
device(config-ipv6-access-list xGW_Filter1)# exit
device(config)# mac access-list abc
device(config-mac-acl-abc)# permit any any any etype 8000
device(config-mac-acl-abc)# exit
device(config)# route-map xGW_map permit 1
device(config-routemap xGW_map)# match ip address 101
device(config-routemap xGW_map)# match ipv6 address xGW_Filter1
device(config-routemap xGW_map)# match l2acl abc
device(config-routemap xGW_map)# set ip next-hop 100.1.1.10
device(config-routemap xGW_map)# set ipv6 next-hop 100:1::1:10
device(config-routemap xGW_map)# set next-hop-flood-vlan 2
device(config-routemap xGW_map)# exit
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# l2 policy route-map xGW_map
device(config-if-e10000-1/1)# ip policy route-map xGW_map
device(config-if-e10000-1/1)# ipv6 policy route-map xGW_map
```

You can match LACP packets using destination MAC address as matching criteria in the Layer 2 ACL configuration and flood it to a VLAN or drop it. The following example shows the method to manage LACP traffic type using Layer 2 PBR.

```
device(config)# mac access-list x2
device(config-mac-acl-x2)# permit any 0180.c200.0002 ffff.ffff.ffff
device(config-mac-acl-abc)# exit
device(config)# route-map xGW_map permit 1
device(config-routemap xGW_map)# match l2acl x2
device(config-routemap xGW_map)# set next-hop-flood-vlan 2
device(config-routemap xGW_map)# exit
device(config)# interface ethernet 1/1
device(config-if-e10000-1/1)# l2 policy route-map xGW_map
```

Displaying Layer 2 PBR information

A number of **show** commands are available to display the Layer 2 ACL accounting information and the route map information used in the Layer 2 PBR configuration.

Displaying Layer 2 PBR accounting information

ACL counters must be enabled to view accounting data. The following example displays the Layer 2 PBR accounting information on a physical interface.

```
device(config)# enable-acl-counter
device(config)# show access-list accounting ethernet 1/2 in l2 policy-based-routing
L2 Policy based Routing Accounting Information:
```

```
Routemap l2pbr10
ACL x10
   0:  10:  permit any any any etype any
       Hit count: (1 sec)                  0  (1 min)                0
                  (5 min)                  0  (accum)                0
```

Syntax: show access-list accounting ethernet *slot/port* **{ in | out } l2 policy-based-routing [omit-zero]**

The following example displays the Layer 2 PBR accounting summary.

```
device# show access-list accounting brief l2 policy-based-routing
  1/1                        x10                0(1s)
                                                0(1m)
                                                0(5m)
                                                0(ac)

  4/2                        x10                0(1s)
                                                0(1m)
                                                0(5m)
                                                0(ac)
```

Syntax: show access-list accounting brief l2 policy-based-routing [omit-zero]

Displaying Layer 2 PBR route map information

A number of **show** commands are available to display the route map information used in the Layer 2 PBR configuration.

Displaying route map information

To view the route map information, use the **show route-map** command. The following example displays route maps that are configured to match Layer 2 ACLs, IPv4 ACLs, and IPv6 ACLs.

```
device(config)# show route-map
route-map pb1 permit 1
 match l2acl  dx2
 set interface null0
route-map pbr permit 1
 match l2acl  abc
 set next-hop-flood-vlan 2 outgoing-da 1111.2222.3333
route-map vlpb permit 1
 match l2acl  dx3
 set interface null0
route-map Test1 permit 1
 match ip address  10
 match ipv6 address  TestIpv61
 set interface ethernet 3/3 preserve-vlan
 set interface ethernet 4/2 preserve-vlan
 set interface ethernet 1/1 preserve-vlan
route-map Test2 permit 1
 match l2acl  400
 set interface ethernet 3/3 preserve-vlan
route-map Test3 permit 1
 match ip address  10
 set interface ethernet 3/3 preserve-vlan
route-map l2pbr permit 1
 match l2acl  x10
```

Syntax: show route-map [*map-name*]

The *map-name* variable specifies the name of the PBR route map you want to view.

To view the route map binding information, enter a command such as the following.

NOTE
The **binding** option is available only on line cards.

```
device# show route-map binding
IPv4 Bindings of p2 :
  4/4

L2 PBR Bindings of pbr :
  3/1
```

Syntax: show route-map [**binding**] [*map-name*]

The *map-name* variable is the name of the route map binding you want to view.

Clearing Layer 2 ACL accounting information

To clear Layer 2 ACL accounting information, enter a command such as the following.

```
device(config)# clear access-list ethernet 2/1 l2 policy-based-routing
```

Syntax: clear access-list ethernet *slot/port* **l2 policy-based-routing**

The **ethernet** *slot/port* parameter specifies the interface on which you want to remove the Layer 2 ACL accounting information.

The **l2** parameter removes the Layer 2 ACL accounting information.

The **policy-based-routing** parameter removes the ACL accounting information of Layer 2 PBR.

Displaying telemetry command

The **show telemetry** command displays the information related to the telemetry configuration.

```
device(config)# show telemetry rule-name
Paths with leading * are configured but disabled, entries with + are for IPv6,
entries with & for L2
entries with # are for UDA
                        Route-map     ACL        Output      Output
Name              Input Policy        Match      VLAN        Port(s)/IP
*&test            N/A   testpbr7      x5         N/A         N/A
*v4test           N/A   v4pbr         101        N/A         N/A
*default-rulename N/A   abc           102        N/A         N/A
*&default-rulenam N/A   abc           x10        N/A         N/A
default-rulename  4/1   deny_pbr      deny_all               NULL0:
+default-rulename 4/1   pbr
*&default-rulenam N/A   vlanpbr       x10        N/A         N/A
e
*&default-rulenam N/A   l2pbr15       x10        N/A         N/A
e
&default-rulename 1/1   l2pbr10       x10
*default-rulename N/A   TGFlood       v4_Permit_A N/A        N/A
ny
*+default-rulename /A   TGFlood       v6_Permit_A N/A N/A
e ny
*&default-rulenam N/A   vlpbr1        dx1        N/A         N/A
e
*&default-rulenam N/A   vlan12        vlan1      N/A         N/A
```

The following command provides a detail mode output in which the list of ports will be fully expanded and displayed if the ports are LAG or VLAN ports. For the output VLAN and output ports information, the command shows the active ports or VLANs that are currently being used for traffic forwarding.

```
device(config)# show telemetry detail rule-name
Rule name: test (disabled)
Input:
Route-map Policy: testpbr7
```

```
L2 ACL match: x5
Output:
Rule name: v4test (disabled)
Input:
Route-map Policy: v4pbr
IPv4 ACL match: 101
Output:
Rule name: default-rulename
Input:
Route-map Policy: abc
L2 ACL match: x10
IPv4 ACL match: 102
Output:
Input:
Route-map Policy: test
Output:
Input: IPv4 - 4/1
Route-map Policy: deny_pbr
IPv4 ACL match: deny_all
Output: IPv4 -
Input: IPv6 - 4/1
Route-map Policy: pbr
L2 ACL match: abc
Output:
Input:
Route-map Policy: vlanpbr
L2 ACL match: x10
Output:
Input:
Route-map Policy: l2pbr15
L2 ACL match: x10
Output:
Input: L2 - 1/1
Route-map Policy: l2pbr10
L2 ACL match: x10
Output:
Input:
Route-map Policy: TGFlood
IPv4 ACL match: v4_Permit_Any
IPv6 ACL match: v6_Permit_Any
```

Protecting against Denial of Service Attacks

Denial of service protection overview

In a Denial-of-Service (DoS) attack, a router is flooded with useless packets for the purpose of slowing down or stopping normal operation.

Devices include measures to defend against two types of DoS attacks: Smurf attacks and TCP SYN attacks.

- Smurf attacks—Attacker sends ICMP echo request (ping) to broadcast address on the network of an intermediary and spoofs the IP address of the victim.
- TCP SYN attacks—Attacker floods a host with TCP SYN packets that have random source IP addresses that fill up the connection queue and service can be denied to legitimate TCP connections.

Protecting against smurf attacks

A smurf attack is a kind of DoS attack where an attacker causes a victim to be flooded with ICMP echo (pPing) replies sent from another network. Figure 15 illustrates how a smurf attack works.

FIGURE 15 How a smurf attack floods a victim with ICMP replies

The attacker sends an ICMP echo request packet to the broadcast address of an intermediary network. The ICMP echo request packet contains the spoofed address of a victim network as its source. When the ICMP echo request reaches the intermediary network, it is converted to a Layer 2 broadcast and sent to the hosts on the intermediary network. The hosts on the intermediary network then send ICMP replies to the victim network.

For each ICMP echo request packet sent by the attacker, a number of ICMP replies equal to the number of hosts on the intermediary network are sent to the victim. If the attacker generates a large volume of ICMP echo request packets, and the intermediary network contains a large number of hosts, the victim can be overwhelmed with ICMP replies.

Avoiding being an intermediary in a smurf attack

A smurf attack relies on the intermediary to broadcast ICMP echo request packets to hosts on a target subnet. When the ICMP echo request packet arrives at the target subnet, it is converted to a Layer 2 broadcast and sent to the connected hosts. This conversion takes place only when directed broadcast forwarding is enabled on the device.

To avoid being an intermediary in a smurf attack, make sure forwarding of directed broadcasts is disabled on the device. Directed broadcast forwarding is disabled by default. To disable directed broadcast forwarding, enter this command.

```
device(config)# no ip directed-broadcast
```

Syntax: [no] ip directed-broadcast

Avoiding being a victim in a smurf attack

You can configure the device to drop ICMP packets when excessive numbers are encountered, as is the case when the device is the victim of a smurf attack. The following example sets threshold values for ICMP packets targeted at the router and drop them when the thresholds are exceeded.

```
device(config)# ip icmp burst-normal 5000 burst-max 10000 lockup 300
```

Syntax: ip icmp burst-normal *value* **burst-max** *value* **lockup** *seconds*

The **burst-normal** value can be from 1 – 100000.

The **burst-max** value can be from 1 – 100000.

The **lockup** value can be from 1 – 10000.

The number of incoming ICMP packets per second are measured and compared to the threshold values, as follows:

- If the number of ICMP packets exceeds the **burst-normal** value, the excess ICMP packets are dropped.
- If the number of ICMP packets exceeds the **burst-max** value, all ICMP packets are dropped for the number of seconds specified by the lockup value. When the lockup period expires, the packet counter is reset and measurement is restarted.

In this example, if the number of ICMP packets received per second exceeds 5,000, the excess packets are dropped. If the number of ICMP packets received per second exceeds 10,000, the device drops all ICMP packets for the next 300 seconds (five minutes).

When incoming ICMP packets exceed the **burst-max** value, the following message is logged.

```
SYSLOG: Jul 26 12:30:31:<13>Jul 26 12:30:31 AB-850 ICMP:Local ICMP exceeds 10 burst packets,
stopping for 15 seconds!!
```

IPv6 traffic not subject to DOS attack filtering

The following IPv6 traffic exceptions (per section 4.4 of RFC 4890) are not subject to DoS attack filtering.

Error messages that are essential to the establishment and maintenance of communications:

- Destination unreachable (Type 1) – All codes
- Packet Too Big (Type 2)
- Time Exceeded (Type 3) – Code 0 only
- Parameter Problem (Type 4) – Codes 1 and 2 only

Address configuration and router selection messages:

- Router Solicitation (Type 133)
- Router Advertisement (Type 134)
- Neighbor Solicitation (Type 135)
- Neighbor Advertisement (Type 136)
- Redirect (Type 137)
- Inverse Neighbor Discovery Solicitation (Type 141)
- Inverse Neighbor Discovery Advertisement (Type 142)

Link-local multicast receiver notification messages:

- Listener Query (Type 130)
- Listener Report (Type 131)
- Listener Done (Type 132)
- Listener Report v2 (Type 143)

Multicast Router Discovery messages:

- Multicast router advertisement (Type 151)
- Multicast router solicitation (Type 152)
- Multicast router termination (Type 153)

Section 4.4 of RFC 4890 also recommends that the following traffic types must not be dropped, however these traffic types will continue to be subject to DoS attack filtering:

- Echo request (Type 128)
- Echo response (Type 129)
- Certificate path solicitation (Type 148)
- Certificate path advertisement (Type 149)

Protecting against TCP SYN attacks

TCP SYN attacks disrupt normal traffic flow by exploiting the way TCP connections are established. When a TCP connection starts, the connecting host sends a TCP SYN packet to the destination host. The destination host responds with a SYN ACK packet, and the connecting host sends back an ACK packet. This process, known as a "TCP three-way handshake", establishes the TCP connection.

While waiting for the connecting host to send an ACK packet, the destination host keeps track of the as-yet incomplete TCP connection in a connection queue. When the ACK packet is received, information about the connection is removed from the connection queue. Usually there is not much time between the destination host sending a SYN ACK packet and the source host sending an ACK packet, so the connection queue clears quickly.

In a TCP SYN attack, an attacker floods a host with TCP SYN packets that have random source IP addresses. For each of these TCP SYN packets, the destination host responds with a SYN ACK packet and adds information to the connection queue. However, since the source host does not exist, no ACK packet is sent back to the destination host, and an entry remains in the connection queue until it ages out (after approximately one minute). If the attacker sends enough TCP SYN packets, the connection queue can fill up, and service can be denied to legitimate TCP connections.

To protect against TCP SYN attacks, you can configure devices to drop TCP SYN packets when excessive numbers are encountered. You can set threshold values for TCP SYN packets that are targeted at the device and drop them when the thresholds are exceeded, as shown in this example.

```
device(config)# ip tcp burst-normal 10 burst-max 100 lockup 300
```

Syntax: ip tcp burst-normal *value* burst-max *value* lockup *seconds*

The **burst-normal** value can be from 1 - 100000.

The **burst-max** value can be from 1 - 100000.

The **lockup** value can be from 1 - 10000.

The number of incoming TCP SYN packets per second is measured and compared to the threshold values as follows:

- If the number of TCP SYN packets exceeds the **burst-normal** value, the excess TCP SYN packets are dropped.
- If the number of TCP SYN packets exceeds the **burst-max** value, all TCP SYN packets are dropped for the number of seconds specified by the **lockup** value. When the lockup period expires, the packet counter is reset and measurement is restarted.

In this example, if the number of TCP SYN packets received per second exceeds 10, the excess packets are dropped. If the number of TCP SYN packets received per second exceeds 100, the device drops all TCP SYN packets for the next 300 seconds (five minutes).

When incoming TCP SYN packets exceed the burst-max value, the following message is logged.

< date > < time >:N:Local TCP exceeds *< burst-max >* burst packets, stopping for *< lockup >* seconds!!

TCP security enhancement

A TCP security enhancement improves the way TCP inbound segments are handled. This enhancement eliminates or minimizes the possibility of a TCP reset attack, in which a perpetrator attempts to prematurely terminate an active TCP session, and a data injection attack, where an attacker injects or manipulates data in a TCP connection.

In both cases, the attack is blind, meaning the perpetrator does not have visibility into the content of the data stream between two devices, but blindly injects traffic. The attacker also does not see the direct effect (the continuing communications between the devices and the impact of the injected packet) but may see the indirect impact of a terminated or corrupted session.

The TCP security enhancement prevents and protects against the following types of attacks:

- Blind TCP reset attack using the reset (RST) bit.
- Blind TCP reset attack using the synchronization (SYN) bit
- Blind TCP data injection attack

The TCP security enhancement is automatically enabled. If necessary, you can disable this feature. Refer to Disabling the TCP security enhancement on page 354.

Protecting against a blind TCP reset attack using the RST bit

In a blind TCP reset attack using the RST bit, a perpetrator attempts to guess the RST segments to prematurely terminate an active TCP session.

To prevent a user from using the RST bit to reset a TCP connection, the RST bit is subject to the following rules when receiving TCP segments:

- If the RST bit is set and the sequence number is outside the expected window, the device silently drops the segment.
- If the RST bit is exactly the next expected sequence number, the device resets the connection.
- If the RST bit is set and the sequence number does not exactly match the next expected sequence value, but is within the acceptable window, the device sends an acknowledgement (ACK).

The TCP security enhancement is enabled by default. To disable it, refer to Disabling the TCP security enhancement on page 354.

Protecting against a blind TCP reset attack using the SYN bit

For a blind TCP reset attack, the attacker tries to guess the SYN bits to terminate an active TCP session.To protect against this type of attack, the SYN bit is subject to the following rules during arrival of TCP segments:

- If the SYN bit is set and the sequence number is outside the expected window, the device sends an ACK to the peer.
- If the SYN bit is set and the sequence number is an exact match to the next expected sequence, the device sends an ACK segment to the peer. Before sending the ACK segment, the software subtracts a 1 from the value being acknowledged.
- If the SYN bit is set and the sequence number is acceptable, the device sends an ACK segment to the peer.

This TCP security enhancement is enabled by default. To disable it, refer to Disabling the TCP security enhancement on page 354.

Protecting against a blind injection attack

In a blind TCP injection attack, the attacker tries to inject or manipulate data in a TCP connection. To reduce the chances of a blind injection attack, an additional check is performed on all incoming TCP segments.

This TCP security enhancement is enabled by default. To disable it, refer to Disabling the TCP security enhancement on page 354.

Disabling the TCP security enhancement

The TCP security enhancement is enabled by default. If necessary, you can disable this feature. When you disable this feature, the device reverts to the original behavior.

To disable the TCP security enhancement, enter the following command at the Global CONFIG level of the CLI.

```
device(config)# no ip tcp tcp-security
```

To re-enable the TCP security enhancement after it has been disabled, enter the following command.

```
device(config)# ip tcp tcp-security
```

Syntax: [no] ip tcp tcp-security

Protecting against UDP attacks

To protect against UDP attacks, you can configure devices to drop UDP packets when excessive numbers are encountered. You can set threshold values for UDP packets that are targeted at the device and drop them when the thresholds are exceeded.

In this example, if the number of UDP packets received per second exceeds 5000, the excess packets are dropped. If the number of UDP packets received per second exceeds 10,000, the device drops all UDP packets for the next 300 seconds (five minutes).

```
device(config)# ip udp burst-normal 5000 burst-max 10000 lockup 300
```

Syntax: [no] ip udp burst-normal *value* burst-max *value* lockup *seconds*

The **burst-normal** value can be from 1 - 100000.

The **burst-max** value can be from 1 - 100000.

The **lockup** value can be from 1 - 10000.

The **no** option removes the configuration and UDP rate limiting is disabled.

The number of incoming UDP packets per second is measured and compared to the threshold values as follows, apply to the individual service

- If the number of UDP packets exceeds the **burst-normal** value, the excess UDP packets are dropped.
- If the number of UDP packets exceeds the **burst-max** value, all UDP packets are dropped for the number of seconds specified by the **lockup** value. When the lockup period expires, the packet counter is reset and measurement is restarted.

  NOTE
  During the lockup period, all UDP packets including the BFD protocol packets, targeted at the device are also dropped.

Enhanced DOS attack prevention for IPv6

IPv6 was introduced to increase the address space. When an IPv6 packet is received, the device broadcast their IPv6 addresses to help clients find and connect to an IPv6 subnet. This created the possibility of DoS attack involving flooding the network segment with random RAs, which consumes CPU resources.

To rate limit the IPv6 subnet packets so the CPU is not overloaded, enter a command such as the following,

```
device(config)# ipv6 rate-limit subnet policy-map
```

Syntax: [no] ipv6 rate-limit subnet policy-map policy-map

The **policy-map** parameter specifies the policy map named in the policy-map variable to be used to provide parameters for rate limiting the port and VLAN specified. This command is only used when configuring traffic policing to a port using a policy map as described in "Applying traffic policing parameters using a policy map" on page 547.

Displaying statistics from a DoS attack

You can display statistics about ICMP and TCP SYN packets that were dropped, passed, or blocked because burst thresholds were exceeded using the **show statistics dos-attack** command.

```
device# show statistics dos-attack
Collecting local DOS attack statistic for slot 1... Completed successfully.
Collecting local DOS attack statistic for slot 2... Completed successfully.
Collecting local DOS attack statistic for slot 3... Completed successfully.
-------------------------- Local Attack Statistics --------------------------
ICMP Drop Count    ICMP Block Count    SYN Drop Count    SYN Block Count
3736338               188              3318293               175
```

Syntax: **show statistics dos-attack** [| **begin** *expression* | **exclude** *expression* | **include** *expression*]

The table below describes this output.

This field...	Displays...
Packet drop count	Number of packets that are dropped when the port is in lockup mode.
Packet Pass Count	Number of packets that are forwarded when the port is in rate-limiting mode.
Packet BLock Count	Number of times the port was shut down for a traffic flow that matched the ACL.

Clear DoS attack statistics

To clear statistics about ICMP and TCP SYN packets, enter the **clear statistics dos-attack** command.

```
device(config)# clear statistics dos-attack
```

Syntax: **clear statistics dos-attack**

```
-------------------------- Local Attack Statistics --------------------------
ICMP Drop Count    ICMP Block Count    SYN Drop Count    SYN Block Count
---------------    ----------------    -------------    ---------------
           0                   0                  0                  0
```

MAC Port-Based Authenication

MAC port-based authentication overview

MAC port-based authentication allows you to configure the device to learn a limited number of secure MAC addresses on an interface. The interface will forward only packets with source MAC addresses that match these secure addresses.

The secure MAC addresses can be specified manually, or the device can learn them automatically. After the device reaches the limit for the number of secure MAC addresses it can learn on the interface, if the interface then receives a packet with a source MAC address that is different from any of the secure learned addresses, it is considered a security violation.

When a security violation occurs, a Syslog entry and an SNMP trap are generated. In addition, the device takes one of two actions: it either drops packets from the violating address (but allows packets from the secure addresses), or it disables the port for a specified amount of time. You specify which of these actions takes place.

The secure MAC addresses are not flushed when an interface is disabled and brought up again. The secure addresses can be kept secure permanently (the default), or can be configured to age out, at which time they are no longer secure. You can configure the device to automatically save the list of secure MAC addresses to the startup-config file at specified intervals, allowing addresses to be kept secure across system restarts.

The port security feature applies only to Ethernet interfaces.

Configuration Considerations

When using the MAC port security feature, the following should be considered.

* If there is no port security configuration at the interface level, global level port security configuration is inherited.

* If a port security attribute is configured at the interface level, interface level configuration for that attribute takes precedence over global level configuration for the same attribute.The rest of the port security attributes that are not configured at the interface level will be inherited from the global level configuration.

Local and global resources

The port security feature uses a concept of local and global "resources" to determine how many MAC addresses can be secured on each interface. In this context, a "resource" is the ability to store one secure MAC address entry. Each interface is allocated 64 local resources. When the port security feature is enabled, the interface can store up to 64 secure MAC addresses using local resources.

Besides the maximum of 64 local resources available to an interface, there are 4096 global resources available. When an interface has secured enough MAC addresses to reach its limit for local resources, it can secure additional MAC addresses by using global resources. Global resources are shared among all the interfaces on a first-come, first-served basis.

The maximum number of MAC addresses any single interface can secure is 64 (the maximum number of local resources available to the interface), plus the number of global resources not allocated to other interfaces.

Configuring the MAC port security feature

To configure the MAC port security feature, you perform the following tasks:

- Enable the MAC port security feature
- Set the maximum number of secure MAC addresses for an interface
- Set the port security age timer
- Specify secure MAC addresses
- Configure the device to automatically save secure MAC addresses to the startup-config file
- Specify the action taken when a security violation occurs
- Deny specific MAC addresses
- Port Security MAC Violation Limits

Enabling the MAC port security feature

By default, the MAC port security feature is disabled on all interfaces. You can enable or disable the feature globally on all interfaces or on an individual interface.

To enable the feature globally, first go to the level for global port security and then enter **enable** , as follows.

```
device(config)# global-port-security
device(config-global-port-security)# enable
```

To disable the feature on all interfaces at once, do the following.

```
device(config)# global-port-security
device(config-global-port-security)#disable
```

Syntax: global-port-security

This command is for global enable port security.

To enable port security on a specific interface, first go to the level of a specific interface and then security level.

```
device(config)# interface ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-port-security-e100-7/11)# enable
```

Syntax: enable

This command applies to a specific interface or global configuration. The interface level take precedence over the global configuration.

Syntax: disable

This command applies to a specific interface or global configuration. The interface level take precedence over the global configuration.

Setting the maximum number of secure MAC addresses for an interface

When the port security feature is enabled, the interface can store 1 secure MAC address. You can increase the number of MAC addresses that can be secured to a maximum of 64, plus the total number of global resources available.

For example, to configure interface 7/11 to have a maximum of 10 secure MAC addresses.

```
device(config)# interface ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-if-e100-7/11)# maximum 10
```

Syntax: maximum number-of-addresses

The *number-of-addresses* parameter can be set to a number from 0 - (64 + the total number of global resources available) The total number of global resources is 4096. Setting the parameter to 0 prevents any addresses from being learned. The default is 1.

Setting the port security age timer

By default, a learned MAC address stays secure indefinitely. You can configure the device to age out secure MAC addresses after a specified amount of time and can do so for all timers globally of for a specific interface.

To set the port security age timer to 10 minutes on all interfaces, first go to the level for global security.

```
device(config)# global-port-security
device(config-global-port-security)# age 10
```

Syntax: global-port-security

Syntax: [no] age minutes

The default is 0 (never age out secure MAC addresses).

To set the port security age timer to 10 minutes on a specific interface, go to the interface level and then the port security level for that interface.

```
device(config)# interface ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-port-security-e100-7/11)# age 10
```

Syntax: port security

Syntax: [no] age minutes

The default is 0 (never age out secure MAC addresses).

Specifying secure MAC addresses

To specify a secure MAC address on an interface, enter commands such as the following.

```
device(config)# interface ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-port-security-e100-7/11)# secure 0050.DA18.747C
```

Syntax: [no] secure mac-address

Autosaving secure MAC addresses to the startup-config file

The learned MAC addresses can automatically be saved to the startup-config file at specified intervals. You can specify the autosave interval at the global level. For example, to set a 20-minute autosave interval globally for learned secure MAC addresses on the router, enter the following commands.

```
device(config)# global-port-security
device(config-port-security)# autosave 20
```

Syntax: global-port-security

Syntax: [no] autosave minutes

The interval range is 15 - 1440 minutes. By default, secure MAC addresses are not autosaved to the startup-config file. To remove autosave intervals, use the **no** form of the **autosave** command.

Setting to delete a dynamically learned MAC address on a disabled interface

By default, a dynamically learned MAC address is not deleted even though the port goes down. You can configure the device to delete a dynamically learned secure MAC addresses when a port goes down, for example, disabled either manually by a user or through a security violation.

You can configure the **delete-dynamic-learn** command at the global level.

To enable the **delete-dynamic-learn** command, enter a command such as the following.

```
device(config)# global-port-security
device(config-port-security)# delete-dynamic-learn
```

Syntax: **global-port-security**

Syntax: **[no] delete-dynamic-learn**

By default, **delete-dynamic-learn** is disabled.

Specifying the action taken when a security violation occurs

A security violation can occur when a user tries to plug into a port where a MAC address is already locked, or the maximum number of secure MAC addresses has been exceeded. When a security violation occurs, an SNMP trap and Syslog message are generated.

In addition, you configure the device to take one of two actions when a security violation occurs: either drop packets from the violating address (and allow packets from secure addresses), or disable the port altogether for a specified amount of time.

To configure the device to drop packets from a violating address and allow packets from secure addresses.

```
device(config)# interface ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-port-security-e100-7/11)# violation restrict
```

Syntax: **violation restrict**

To shut down the port when a security violation occurs.

```
device(config)# interface ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-port-security-e100-7/11)# violation shutdown
```

Syntax: **violation shutdown**

To specific the mac-addresses that will be denied. All other mac-addresses no specified will be allowed.

```
device(config)# interface ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-port-security-e100-7/11)# deny-mac-address
```

Syntax: **deny-mac-address**

> NOTE
> When using this feature with a 24-port 10/100 module (part number B24E) only the **shutdown** option is supported. The **restrict** option is not supported on the B24E.

Specifying the number of MAC addresses to be denied

You can specify the number of MAC addresses that are to be denied before the NetIron device shuts the port down using the **restrict-max-deny** command.

```
device(config)# interface ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-if-e100-7/11)# violation restrict
device(config-port-security-e100-7/11)# restrict-max-deny 40
```

Syntax: restrict-max-deny*number*

The *number* parameter indicates the number of MAC addresses that are to be denied before the NetIron device shuts the port down. The *number* range is between 1 and 1024. The code example indicates that the device will be shut down after 40 MAC addresses are denied.

> NOTE
> The default deny list size is 128. When restrict-max-deny is being configured, if there are more MAC addresses present in the deny MAC list than the value being specified for restrict-max-deny, an ERROR will be thrown to inform the user that the value being set should be less than or equal to the current number of MAC addresses in the deny MAC list, as shown in the following example.

```
device(config-port-security-e10000-4/6)#restrict-max-deny 5
ERROR: restrict-max-deny value cannot be configured to a value lesser than current denied MAC count 8.
```

Denying specific MAC addresses

You can configure the *violation deny mode* . The violation deny mode allows you to deny MAC addresses on a global level or on a per port level.

Denying MAC addresses globally

To deny a specific MAC address globally, enable the violation deny mode, then specify the MAC address to be denied.

```
device(config)# global-port-security
device(config-port-security)# violation deny
device(config-port-security)# deny-mac-address 0000.0000.0001 2
```

Global denied secure MAC addresses are denied system-wide. These MAC entries are added to the MAC table as deny entries, when a flow is received and are the only MAC addresses that are denied. All other MAC addresses are allowed.

A maximum of 512 deny MAC addresses can be configured on a global level.

Denying MAC addresses on an interface

You can specify which MAC addresses can be denied on an interface.

```
device(config)# internet ethernet 7/11
device(config-if-e100-7/11)# port security
device(config-port-security-e100-7/11)# violation deny
device(config-port-security-e100-7/11)# deny-mac-addr 0000.1111.2222 4
```

Only the configured MAC addresses are denied on the specified interface. All other MAC addresses are allowed.

A maximum of 64 deny MAC addresses can be configured at an interface level.

Displaying MAC addresses that have been denied

Use the **show port security global-deny** command to display all the MAC addresses that have been denied globally. Use the **show port security denied-macs** command to display all the denied MAC addresses

Port security MAC violation limit

Use the **violation restrict** command to specify how many packets the system can receive in a one-second interval from denied MAC address before the system shuts the port down.

Configuring port security

To enable this new mode, enter a command such as the following.

```
device(config)# global-port-security
device(config-port-security)# violation restrict 12
```

Syntax: violation restrict [**#-denied-packets processed**]

Enter 0 - 64000. This parameter has no default.

> **NOTE**
> With the introduction of this command, packets from denied MAC addresses are now processed in software by the LP. They are no longer programmed in the hardware.

In addition to the new processing of packets from denied MAC addresses, these packets can now be logged in the Syslog. And to prevent the Syslog from being overwhelmed with messages for denied packets, you can specify how many messages will be logged per second, based on a packet's IP address.

```
device(config)# global-port-security
device(config-port-security)# violation restrict 12
device(config-port-security)# deny-log-rate <7>
```

Syntax: deny-log-rate [**#-logs**]

The #-logs parameter specifies the count per line card. Enter 1 - 10. There is no default.

The logged message contains the packet's IP address and the MAC address of the denied packet. For example, the following configuration shows that violation restrict is configured;

```
interface ethernet 14/1
 port security
  enable
  maximum 5
  violation restrict 1000
  secure-mac-address 0000.0022.2222 10
  secure-mac-address 0000.0022.2223 10
  secure-mac-address 0000.0022.2224 10
  secure-mac-address 0000.0022.2225 10
  secure-mac-address 0000.0022.2226 10
```

When packet from MAC address 000.0022.2227, an address that is not a secured MAC address, the following Syslog message is generated.

```
SYSLOG: Mar 10 17:36:12:<12>3-RW-Core-3, Interface e14/1 shutdn due to high rate of denied mac
0000.0022.2227, vlan 10
SYSLOG: Mar 10 17:36:12:<14>3-RW-Core-3, Interface ethernet14/1, state
down - disabled
```

However, when **deny-log-rate** is configured,

```
interface ethernet 14/1
 disable
 port security
  enable
  maximum 5
  violation restrict 1000
  deny-log-rate 4
  secure-mac-address 0000.0022.2222 10
  secure-mac-address 0000.0022.2223 10
  secure-mac-address 0000.0022.2224 10
  secure-mac-address 0000.0022.2225 10
  secure-mac-address 0000.0022.2226 10
```

The following Syslog messages are generated.

```
Mar 10 17:38:51:I:Port security denied pkt: 0000.0022.2224 -> 0000.0011.1111
10.19.1.2 -> 10.19.1.1 [Protocol:114]
Mar 10 17:38:51:I:Port security denied pkt: 0000.0022.2224 -> 0000.0011.1111
10.19.1.2 -> 10.19.1.1 [Protocol:114]
Mar 10 17:38:51:I:Port security denied pkt: 0000.0022.2224 -> 0000.0011.1111
10.19.1.2 -> 10.19.1.1 [Protocol:114]
Mar 10 17:38:51:I:Port security denied pkt: 0000.0022.2224 -> 0000.0011.1111
10.19.1.2 -> 10.19.1.1 [Protocol:114]
Mar 10 17:38:51:I:Port security denied pkt: 0000.0022.2224 -> 0000.0011.1111
10.19.1.2 -> 10.19.1.1 [Protocol:114]
```

Displaying port security information

You can display the following information about the port security feature:

- The secure MAC addresses that have been saved to the startup-config file by the autosave feature
- The port security settings for an individual port or for all the ports on a specified module
- The secure MAC addresses configured on the device
- Port security statistics for an interface or for a module

Displaying port security settings

You can display the port security settings for an individual port or for all the ports on a specified module. For example, to display the port security settings for port 7/11, enter the following command.

```
device# show port security e 1/1
Port  Security  MacAddrs        Violation  PortShutdn(minutes) SecureMac Learn
                Learnt/Max  Total/Count/Type   Status/Time/Remain   AgeTime
----- -------- ----------- ----------------- ------------------- --------- ------
 1/1  disabled   0/1          0/  0/shutdown   no/permanent          permanent  yes
```

Syntax: show port security module | portnum

This command displays the following information

This field...	Displays...
Port	The slot and port number of the interface.
Security	Whether the port security feature has been enabled on the interface.
MacAddrsLearnt or Max	Learnt - The number of secure MAC addresses that have been learned on the interface. Max - The maximum number of secure MAC addresses that can be learned on the interface.

This field...	Displays...
ViolationTotal or Count or Type	Total - The total number of violations that have occurred on the interface.
	Count - The count of the current violation on the interface.
	Type - The action to be undertaken when a security violation occurs, either "shutdown" or "restrict".
PortShutdn (minutes) Status or Time or Remain	Status - Whether the interface has been shut down due to a security violation.
	Time - The number of seconds a port is shut down following a security violation, if the port is set to "shutdown" when a violation occurs.
	Remain - The number of seconds before the port is enabled again.
SecureMacAgeTime	The amount of time, in minutes, MAC addresses learned on the port will remain secure.
Learn	Whether the port is able to learn MAC addresses.

Displaying the secure MAC addresses on the device

To list the secure MAC addresses configured on the device, enter the following command.

```
device(config)# show port security mac
Port  Num-Addr Secure-Src-Addr Resource Age-Left  Shutdown/Time-Left
----- -------- --------------- -------- --------- ------------------
 7/11        1  0050.da18.747c   Local        10        no
```

Syntax: show port security mac

This command displays the following information.

This field...	Displays...
Port	The slot and port number of the interface.
Num-Addr	The number of MAC addresses secured on this interface.
Secure-Src-Addr	The secure MAC address.
Resource	Whether the address was secured using a local or global resource. Refer to Local and global resources on page 357 for more information.
Age-Left	The number of minutes the MAC address will remain secure.
Shutdown or Time-Left	Whether the interface has been shut down due to a security violation and the number of seconds before it is enabled again.

Displaying port security statistics

You can display port security statistics for an interface or for a module.

For example, to display port security statistics for interface 7/11.

```
device# show port security statistics e 7/11
Port  Total-Addrs Maximum-Addrs Violation Shutdown/Time-Left
----- ----------- ------------- --------- ------------------
 7/11           1             1         0        no
```

Syntax: show port security statistics portnum

This field...	Displays...
Port	The slot and port number of the interface.
Total-Addrs	The total number of secure MAC addresses on the interface.

This field...	Displays...
Maximum-Addrs	The maximum number of secure MAC addresses on the interface.
Violation	The number of security violations on the port.
Shutdown or Time-Left	Whether the port has been shut down due to a security violation and the number of seconds before it is enabled again.

To display port security statistics for a module, enter the following command.

```
device# show port security statistics 7
Module 7:
  Total ports: 0
  Total MAC address(es): 0
  Total violations: 0
  Total shutdown ports 0
```

Syntax: show port security statistics module

This field...	Displays...
Total ports:	The number of ports on the module.
Total MAC address(es):	The total number of secure MAC addresses on the module.
Total violations:	The number of security violations encountered on the module.
Total shutdown ports:	The number of ports on the module shut down as a result of security violations.

802.1x Port-Based Authentication

802.1X authentication overview

The IEEE 802.1X standard is designed to govern the authentication of devices attached to LAN ports. The 802.1X protocol defines a port-based authentication algorithm involving network data communication between client-based supplicant software, an authentication database on a server, and the authenticator device. Using 802.1X authentication, you can configure a device to grant access to a port based on information supplied by a client to an authentication server.

When a user logs on to a network that uses 802.1X authentication, the device grants (or does not grant) access to network services after the user is authenticated by an authentication server. The user-based authentication in 802.1X authentication provides an alternative to granting network access based on a user's IP address, MAC address, or subnetwork.

The Extreme implementation of 802.1X authentication supports the following RFCs:

- RFC 2284: PPP Extensible Authentication Protocol (EAP)
- RFC 2865: Remote Authentication Dial In User Service (RADIUS)
- RFC 2869: RADIUS Extensions

    NOTE
    SNMP is not supported for 802.1X authentication.

IETF RFC support

The implementation of 802.1x port-based authentication supports the following RFCs:

- RFC 2284 PPP Extensible Authentication Protocol (EAP)
- RFC 2865 Remote Authentication Dial In User Service (RADIUS)
- RFC 2869 RADIUS Extensions

How 802.1x port security works

This section explains the basic concepts behind 802.1x port security, including device roles, how the devices communicate, and the procedure used for authenticating clients.

Device roles in an 802.1x configuration

The 802.1x standard defines the roles of **clientorSupplicant** , **Authenticator** , and **Authentication Server** in a network.

The client (known as a **Supplicant** in the 802.1x standard) provides username or password information to the Authenticator. The Authenticator sends this information to the Authentication Server. Based on the client's information, the Authentication Server determines whether the client can use services provided by the Authenticator. The Authentication Server passes this information to the Authenticator, which then provides services to the client, based on the authentication result.

Figure 16 illustrates these roles.

FIGURE 16 Authenticator, client or supplicant, and authentication server in an 802.1x configuration

Authenticator – The device that controls access to the network. In an 802.1x configuration, the device serves as the Authenticator. The Authenticator passes messages between the client and the Authentication Server. Based on the identity information supplied by the client, and the authentication information supplied by the Authentication Server, the Authenticator either grants or does not grant network access to the client.

client or supplicant – The device that seeks to gain access to the network. clients must be running software that supports the 802.1x standard (for example, the Windows XP operating system). clients can either be directly connected to a port on the Authenticator, or can be connected by way of a hub.

Authentication server – The device that validates the client and specifies whether or not the client may access services on the device. The device supports Authentication Servers running RADIUS.

Communication between the devices

For communication between the devices, 802.1x port security uses the Extensible Authentication Protocol (EAP), defined in RFC 2284. The 802.1x standard specifies a method for encapsulating EAP messages so that they can be carried over a LAN. This encapsulated form of EAP is known as EAP over LAN (EAPOL). The standard also specifies a means of transferring the EAPOL information between the client or Supplicant, Authenticator, and Authentication Server.

EAPOL messages are passed between the Port Access Entity (PAE) on the Supplicant and the Authenticator. Figure 17 shows the relationship between the Authenticator PAE and the Supplicant PAE.

FIGURE 17 Authenticator PAE and supplicant PAE

Authenticator PAE – The Authenticator PAE communicates with the Supplicant PAE, receiving identifying information from the Supplicant. Acting as a RADIUS client, the Authenticator PAE passes the Supplicant's information to the Authentication Server, which decides whether the Supplicant can gain access to the port. If the Supplicant passes authentication, the Authenticator PAE grants it access to the port.

Supplicant PAE – The Supplicant PAE supplies information about the client to the Authenticator PAE and responds to requests from the Authenticator PAE. The Supplicant PAE can also initiate the authentication procedure with the Authenticator PAE, as well as send logoff messages.

Controlled and uncontrolled ports

A physical port on the device used with 802.1x port security has two virtual access points: a controlled port and an uncontrolled port. The controlled port provides full access to the network. The uncontrolled port provides access only for EAPOL traffic between the client and the Authentication Server. When a client is successfully authenticated, the controlled port is opened to the client. Figure 18 illustrates this concept.

FIGURE 18 Controlled and uncontrolled ports before and after client authentication
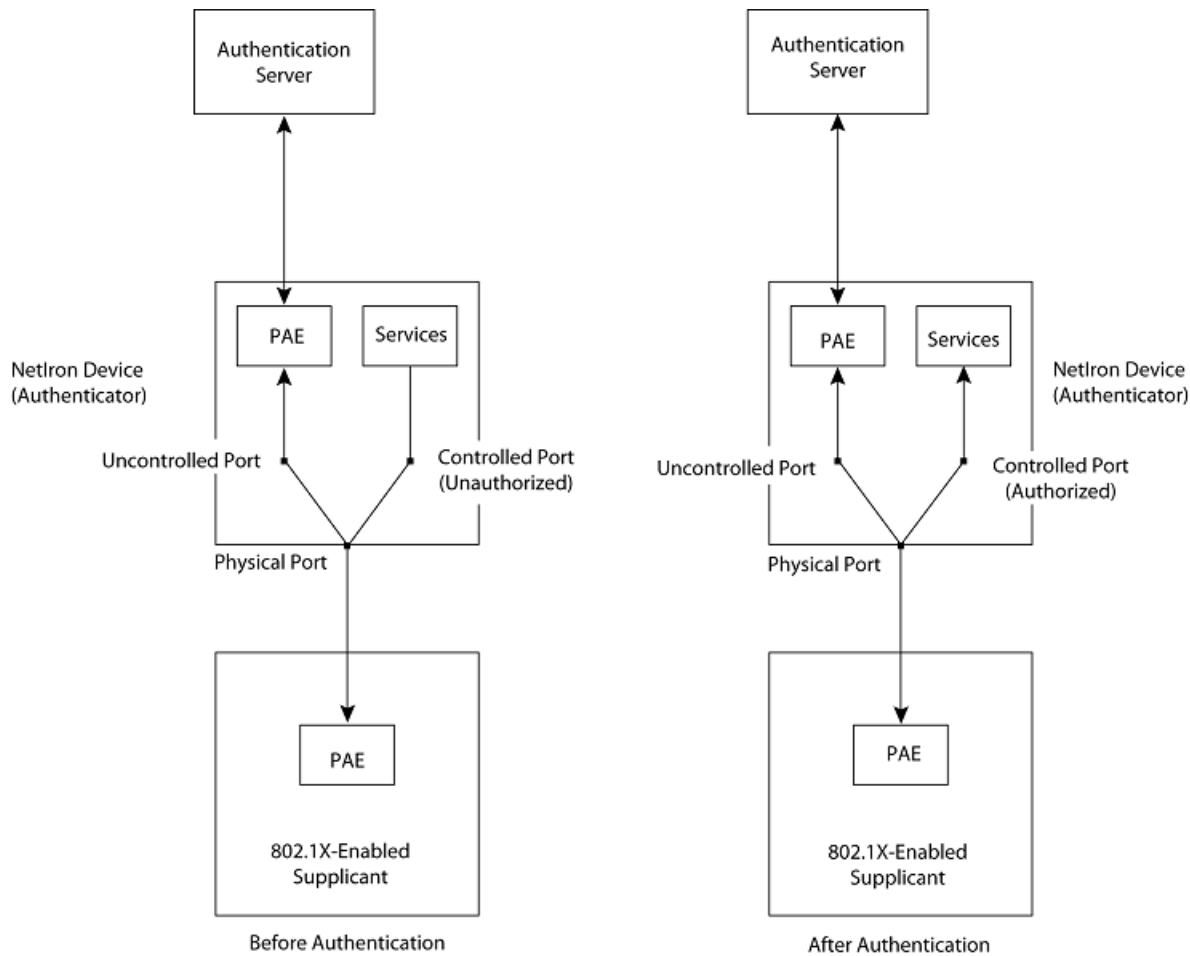
Before a client is authenticated, only the uncontrolled port on the Authenticator is open. The uncontrolled port allows only EAPOL frames to be exchanged between the client and the Authentication Server. The controlled port is in the unauthorized state and allows no traffic to pass through.

During authentication, EAPOL messages are exchanged between the Supplicant PAE and the Authenticator PAE, and RADIUS messages are exchanged between the Authenticator PAE and the Authentication Server. Refer to Message exchange during authentication on page 370 for an example of this process. If the client is successfully authenticated, the controlled port becomes authorized, and traffic from the client can flow through the port normally.

By default, all controlled ports on the device are placed in the authorized state, allowing all traffic. When authentication is activated on an 802.1x-enabled interface, the controlled port on the interface is placed initially in the unauthorized state. When a client connected to the port is successfully authenticated, the controlled port is then placed in the authorized state until the client logs off. Refer to Enabling 802.1x port security on page 383 for more information.

Message exchange during authentication

Figure 19 illustrates a sample exchange of messages between an 802.1x-enabled client, a device acting as Authenticator, and a RADIUS server acting as an Authentication Server.

FIGURE 19 Message exchange during authentication
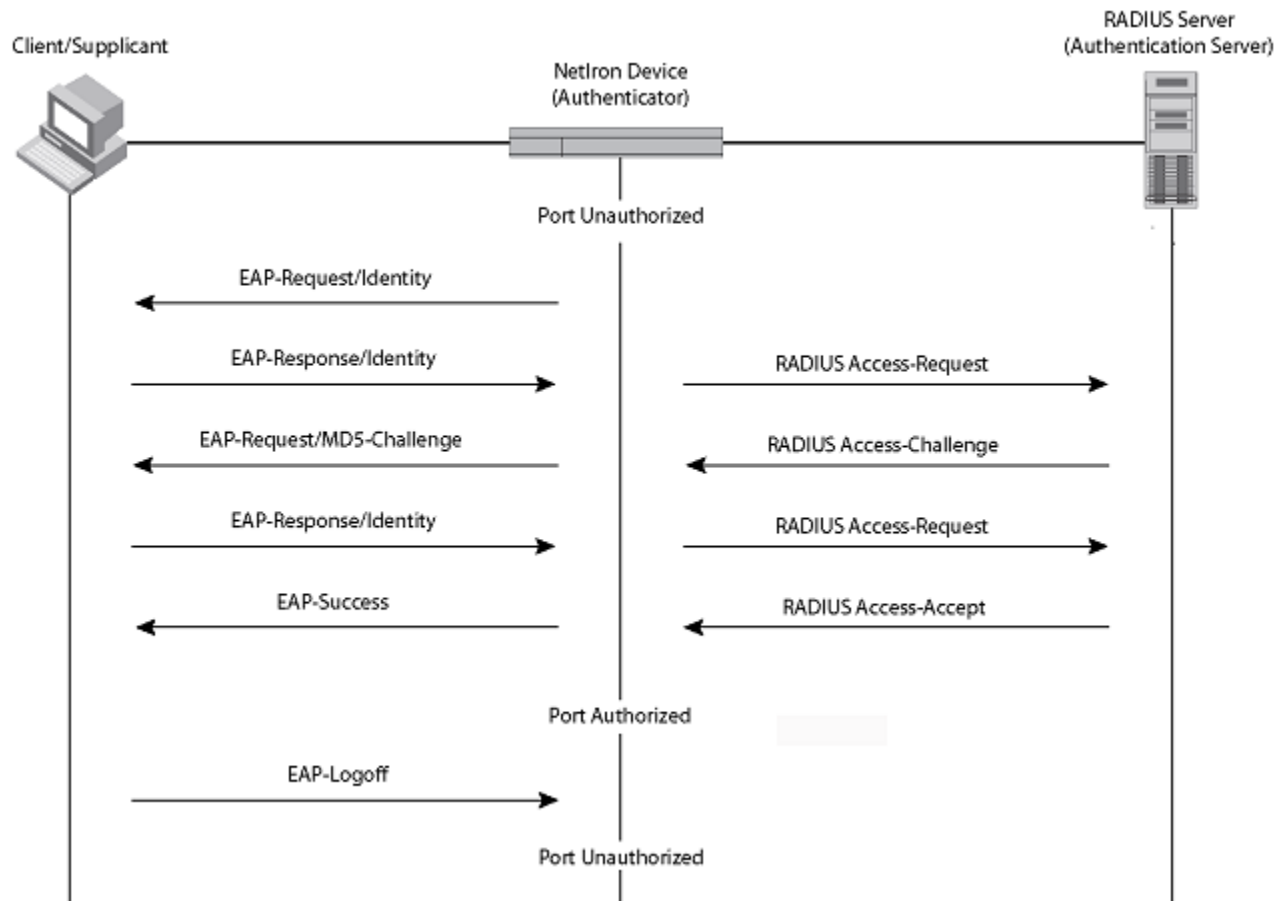
In this example, the Authenticator (the device) initiates communication with an 802.1x-enabled client. When the client responds, it is prompted for a user name (255 characters maximum) and password. The Authenticator passes this information to the Authentication Server, which determines whether the client can access services provided by the Authenticator. When the client is successfully authenticated by the RADIUS server, the port is authorized. When the client logs off, the port becomes unauthorized again.

Extreme's 802.1x implementation supports dynamic VLAN assignment. If one of the attributes in the Access-Accept message sent by the RADIUS server specifies a VLAN identifier, and this VLAN is available on the device, the client's port is moved from its default VLAN to the specified VLAN. When the client disconnects from the network, the port is placed back in its default VLAN. Refer to Dynamic VLAN assignment for 802.1X ports on page 379 for more information.

Extreme's 802.1x implementation supports dynamically applying an IP ACL or MAC address filter to a port, based on information received from the Authentication Server.

If a client does not support 802.1x, authentication cannot take place. The device sends EAP-Request or Identity frames to the client, but the client does not respond to them.

When a client that supports 802.1x attempts to gain access through a non–802.1x-enabled port, it sends an EAP start frame to the device. When the device does not respond, the client considers the port to be authorized, and starts sending normal traffic.

Devices support MD5-challenge TLS and any other EAP-encapsulated authentication types in EAP Request or Response messages. In other words, the devices are transparent to the authentication scheme used.

Authentication of multiple clients connected to the same port

Devices support 802.1X authentication for ports with more than one client connected to them. The following figure illustrates a sample configuration where multiple clients are connected to a single 802.1X port.

FIGURE 20 Multiple clients connected to a single 802.1X-enabled port

Clients/Supplicants running 802.1X-compliant client software

If there are multiple clients connected to a single 802.1X-enabled port, the device authenticates each of them individually. Each client's authentication status is independent of the others, so that if one authenticated client disconnects from the network, it has no effect on the authentication status of any of the other authenticated clients. The number of clients supported at the system level is 255.

By default, traffic from clients that cannot be authenticated by the RADIUS server is dropped.

You can optionally configure the device to assign the port to a "restricted" VLAN if authentication of the client is unsuccessful.

How 802.1x multiple client authentication works

When multiple clients are connected to a single 802.1x-enabled port on a router (as in Authentication of multiple clients connected to the same port on page 372), 802.1x authentication is performed in the following ways.

1. One of the 802.1x-enabled clients attempts to log into a network in which a device serves as an Authenticator.

2. The device creates an internal session (called a dot1x-mac-session) for the client. A dot1x-mac-session serves to associate a client's MAC address and username with its authentication status.

3. The device performs 802.1x authentication for the client. Messages are exchanged between the device and the client, and between the device and the Authentication Server (RADIUS server). The result of this process is that the client is either successfully authenticated or not authenticated, based on the username and password supplied by the client.

4. If the client is successfully authenticated, the client's dot1x-mac-session is set to "access-is-allowed". This means that traffic from the client can be forwarded normally.

5. If authentication for the client is unsuccessful the first time, multiple attempts to authenticate the client will be made as determined by the *attempts* variable in the **auth-fail-max-attempts** command.

6. If authentication for the client is unsuccessful more than the number of times specified by the *attempts* variable in the **auth-fail-max-attempts** command, an authentication-failure action is taken. The authentication-failure action can be either to drop traffic from the client, or to place the port in a "restricted" VLAN:
    - – If the authentication-failure action is to drop traffic from the client, then the client's dot1x-mac-session is set to "access-denied", causing traffic from the client to be dropped in hardware.
      – If the authentication-failure action is to place the port in a "restricted" VLAN, If the client's dot1x-mac-session is set to "access-restricted" then the port is moved to the specified restricted VLAN, and traffic from the client is forwarded normally.

7. When the client disconnects from the network, the device deletes the client's dot1x-mac-session. This does not affect the dot1x-mac-session or authentication status (if any) of the other clients connected on the port.
    - The client's dot1x-mac-session establishes a relationship between the username and MAC address used for authentication. If a user attempts to gain access from different clients (with different MAC addresses), he or she would need to be authenticated from each client.
    - If a client has been denied access to the network (that is, the client's dot1x-mac-session is set to "access-denied"), then you can cause the client to be re-authenticated by manually disconnecting the client from the network, or by using the **clear dot1x mac-session** command. Refer to the Clearing a dot1x-mac-session for a MAC_address task for information on using this command.
    - When a client has been denied access to the network, its dot1x-mac-session is aged out if no traffic is received from the client's MAC address over a fixed hardware aging period (70 seconds), plus a configurable software aging period. You can optionally change the software aging period for dot1x-mac-sessions or disable aging altogether. After the denied client's dot1x-mac-session is aged out, traffic from that client is no longer blocked, and the client can be re-authenticated.

802.1x port security and sFlow

sFlow is a system for observing traffic flow patterns and quantities within and among a set of the devices. sFlow works by taking periodic samples of network data and exporting this information to a collector.

When you enable sFlow forwarding on an 802.1x-enabled interface, the samples taken from the interface include the user name string at the inbound or outbound port, if that information is available.

For more information on sFlow, refer to the *Extreme NetIron Layer 2 Switching Configuration Guide*.

Configuring 802.1x port security

Configuring 802.1x port security on a device consists of the following tasks.

1. Configuring device interaction with the Authentication Server:
   - – Configuring an authentication method list for 802.1x on page 374
   - – Setting RADIUS parameters on page 374
   - – Dynamic VLAN assignment for 802.1X ports on page 379 (optional)

2. Configuring the device role as the Authenticator:
   - – Enabling 802.1x port security on page 383
   - – Initializing 802.1x on a port on page 386 (optional)

3. Configuring device interaction with clients:
   - – Periodic reauthentication on page 384 (optional)
   - – Manual reauthentication of a port on page 384 (optional)
   - – Quiet period for reauthentication on page 385 (optional)
   - – Retransmission interval for EAP-Request/Identity frames on page 385 (optional)
   - – Retransmission limit for EAP-Request/Identity frame on page 385 (optional)
   - – Retransmission timeout of EAP-Request frames to the client on page 385 (optional)
   - – Allowing multiple 802.1x clients to authenticate on page 386 (optional)

     NOTE
     Multi-Device Port Authentication and 802.1x authentication can both be enabled on a port; however only one of them can authenticate a MAC address or 802.1x client. Refer to "Support for multi-device port authentication and 802.1x on the same interface".

Configuring an authentication method list for 802.1x

To use 802.1x port security, you must specify an authentication method to be used to authenticate clients. The device supports RADIUS authentication with 802.1x port security. To use RADIUS authentication with 802.1x port security, you create an authentication method list for 802.1x and specify RADIUS as an authentication method, then configure communication between the device and RADIUS server.

```
device(config)# aaa authentication dot1x default radius
```

Syntax: [no] aaa authentication dot1x default *method-list*

For the *method-list*, enter at least one of the following authentication methods:

radius - Use the list of all RADIUS servers that support 802.1x for authentication.

none - Use no authentication. The client is automatically authenticated without the device using information supplied by the client.

> **NOTE**
> If you specify both **radius** and **none**, make sure **radius** comes before **none** in the method list.

Setting RADIUS parameters

To use a RADIUS server to authenticate access to a device, you must identify the server to the device.

```
device(config)#
 radius-server host 10.157.22.99 auth-port 1812 acct-port 1813 default key mirabeau dot1x
```

Syntax: **radius-server host** *ip-addr* | *server-name* [**auth-port** *number* **acct-port** *number* [**authentication-only** | **accounting-only** | **default** [**key 0** | **1** *string* [**dot1x**]]]]

The **host** *ip-addr* | *server-name* parameter is either an IP address or an ASCII text string.

The **auth-port** *number* parameter specifies what port to use for RADIUS authentication.

The **acct-port** *number* parameter specifies what port to use for RADIUS accounting.

The **dot1x** parameter indicates that this RADIUS server supports the 802.1x standard. A RADIUS server that supports the 802.1x standard can also be used to authenticate non-802.1x authentication requests.

You must configure RADIUS TLS server with support for the 802.1x standard using the following command example:

```
device(config)# radius-server host 10.25.105.44 ssl-auth-port 2083 dot1x
```

> **NOTE**
> To implement 802.1x port security, at least one of the RADIUS servers identified to the device must support the 802.1x standard.

Supported RADIUS attributes

Many IEEE 802.1x Authenticators will function as RADIUS clients. Some of the RADIUS attributes may be received as part of IEEE 802.1x authentication. The device supports the following RADIUS attributes for IEEE 802.1x authentication:

TABLE 45 Supported RADIUS attributes for authentication

Attribute ID	Attribute name	Data Type	Description
1	User-Name	String	Indicates the name of the user to be authenticated (RFC 2865).
4	NAS-IP-Address	Address	IP address of the network access server (NAS) which is requesting user authentication (RFC 2865).
5	NAS-Port	Integer	Physical port number of NAS that is authenticating the user. It is only used in Access-Request packets (RFC 2865). **NOTE** The physical port number refers to a physical connection on the NAS rather than a TCP or UDP port number. Either NAS-Port, NAS-Port-Type (61),or both should be present in the Access-Request packet, when the NAS differentiates between its ports.
6	Service-Type	Integer	Type of service the user has requested, or the type of service to be provided. It may be used in both Access-Request and Access-Accept packets. A NAS is not required to implement all service types, but must treat unknown or unsupported Service-Types as though an Access-Reject packet has been received (RFC 2865).
7	Framed-Protocol	Integer	Framing protocol to be used for framed access. It may be used in both Access-Request and Access-Accept packets (RFC 2865).
11	Filter-Id	Text	Name of the filter list for this user (RFC 2865).
12	Framed-MTU	Integer	Maximum transmission unit (MTU) to be configured for the user, when it is not negotiated by some other means (such as Point-to-Point protocol (PPP)). It may be used in Access-Accept packets or in an Access-

TABLE 45 Supported RADIUS attributes for authentication (continued)

Attribute ID	Attribute name	Data Type	Description
			Request packet as a hint by the NAS to the server that it would prefer that value; the server is not required to honor the hint (RFC 2865).
24	State	String	This attribute is available to be sent by the server to the client: • In an Access-Challenge and must be sent unmodified from the client to the server in the new Access-Request reply (if any) to that challenge. • In an Access-Accept that also includes a Termination-Action attribute with the value of RADIUS-Request. When the NAS performs the Termination-Action by sending a new Access-Request upon termination of the current session, it must include the State attribute unchanged in that Access-Request. In either preceding usage, the client must not interpret the attribute locally. A packet must have only one State attribute or none. Usage of the State attribute is implementation dependent (RFC 2865).
26	Vendor-Specific	String	This attribute allows vendors to support their own extended attributes (RFC 2865).
31	Calling-Station-Id	String	This attribute allows the NAS to send (in an Access-Request packet) the phone number that a call comes from. The phone number is identified using Automatic Number Identification (ANI) or similar technology. It is only used in Access-Request packets (RFC 2685).
32	NAS-Identifier	String	Network access server (NAS) identifier (the hostname of the device). It can be used instead of or in addition to NAS-IP-Address in the Access-Request packet that is sent to the authentication server. **NOTE** When the host name is not configured on the device, the NAS-Identifier attribute is not sent in the Access-Request packet (RFC 2865).
61	NAS-Port-Type	Integer	The physical port type that is authenticating the user. It can be used instead of or in addition to NAS-Port in the Access-Request packet that is sent to the authentication server. For network access control (NAC) authentication specify the NAS-Port-Type as "Ethernet". For device management authentication (when the user attempts to access the device by way of Telnet or SSH), specify the NAS-Port-Type as "Virtual". (RFC 2865).
64	Tunnel-Type	Integer	Tunnel protocol that is either in use by a tunnel terminator or to be used by a tunnel initiator (RFC 2868).
65	Tunnel-Medium-Type	Integer	Transport medium to use when creating a tunnel for protocols that can operate over multiple transport mediums (RFC 2868).

TABLE 45 Supported RADIUS attributes for authentication (continued)

Attribute ID	Attribute name	Data Type	Description
79	EAP-Message	String	An Extensible Authentication Protocol (EAP) packet (RFC 3579).
80	Message-Authenticator	String	A RADIUS Client receiving an Access-Accept, Access-Reject or Access-Challenge with a Message-Authenticator Attribute present must calculate the correct value of the Message-Authenticator and silently discard the packet if it does not match the value sent (RFC 2869).
81	Tunnel-Private-Group-ID	String	Group identifier for a particular tunneled session (RFC 2868).
95	NAS-IPv6-Address	Address	IPv6 address of the NAS that is requesting user authentication (RFC 3162).

The following table shows Microsoft vendor-specific RADIUS attributes that are supported on the Extreme device; these attributes are transmitted as sub-attributes within the Vendor-Specific attribute (26). For further information, refer to RFC 2548.

TABLE 46 Supported Microsoft vendor-specific RADIUS attributes for authentication

Attribute ID	Attribute Name	Data Type	Description
16	MS-MPPE-Send-Key	String	The MS-MPPE-Send-Key Attribute contains a session key for use by the Microsoft Point-to-Point Encryption Protocol (MPPE). This key is intended for encrypting packets sent from the NAS to the remote host. This attribute is only included in Access-Accept packets (RFC 2548).
17	MS-MPPE-Recv-Key	String	The MS-MPPE-Recv-Key Attribute contains a session key for use by the Microsoft Point-to-Point Encryption Protocol (MPPE). As the name implies, this key is intended for encrypting packets received by the NAS from the remote host. This attribute is only included in Access-Accept packets (RFC 2548).

The following table shows the RADIUS attributes that are supported for IEEE 802.1x accounting.

TABLE 47 Supported RADIUS attributes for accounting

Attribute ID	Attribute Name	Data Type	Description
4	NAS-IP-Address	address	IP address of the network access server which is requesting user authentication (RFC 2865).
5	NAS-Port	integer	Physical port number of the NAS that is authenticating the user. It is only used in Access-Request packets (RFC 2865). **NOTE** The physical port number refers to a physical connection on the NAS rather than a TCP or UDP port number. Either NAS-Port, NAS-Port-Type (61),or both should be present in the Access-Request packet, when the NAS differentiates between its ports.
6	Service-Type	Integer	Indicates the type of service the user has requested, or the type of service to be provided. It may be used in both Access-Request and Access-Accept packets. A

TABLE 47 Supported RADIUS attributes for accounting (continued)

Attribute ID	Attribute Name	Data Type	Description
			NAS is not required to implement all service types, but must treat unknown or unsupported Service-Types as though an Access-Reject packet has been received (RFC 2865).
12	Framed-MTU	Integer	Indicates the maximum transmission unit to be configured for the user, when it is not negotiated by some other means (such as PPP). It may be used in Access-Accept packets or in an Access-Request packet as a hint by the NAS to the server that it would prefer that value; the server is not required to honor the hint (RFC 2865).
31	Calling-Station-Id	String	This attribute allows the NAS to send (in an Access-Request packet) the phone number that a call comes from. The phone number is identified using Automatic Number Identification (ANI) or similar technology. It is only used in Access-Request packets (RFC 2685).
40	Acct-Status-Type	integer	Indicates whether this Accounting-Request marks the beginning (Start) or the end (Stop) of the user service. It may be used by the client to mark the start of accounting (for example, upon booting) by specifying Accounting-On and to mark the end of accounting (for example, just before a scheduled reboot) by specifying Accounting-Off (RFC 2866).
41	Acct-Delay-Time	integer	Indicates how many seconds the client has been trying to send this record for, and can be subtracted from the time of arrival on the server to find the approximate time of the event generating this Accounting-Request (network transit time is ignored). RFC 2866.
44	Acct-Session-Id	String	Unique Accounting ID used to match start and stop records in a log file. The start and stop records for a given session must have the same Acct-Session-Id. An Accounting-Request packet must have an Acct-Session-Id. An Access-Request packet may have an Acct-Session-Id; if it does, then the NAS must use the same Acct-Session-Id in the Accounting-Request packets for that session (RFC 2866).
45	Acct-Authentic	Integer	This attribute may be included in an Accounting-Request to indicate how the user was authenticated (by RADIUS, the NAS itself, or another remote authentication protocol). Users who are delivered service without being authenticated should not generate Accounting records (RFC 2866).
46	Acct-Session-Time	Integer	Indicates the number of seconds the user has received service for, and can only be present in Accounting-Request records where the Acct-Status-Type is set to Stop (RFC 2866).
49	Acct-Terminate-Cause	Integer	Indicates how the session was terminated, and can only be present in Accounting-Request records where the Acct-Status-Type is set to Stop (RFC 2866).
61	NAS-Port-Type	integer	The physical port type that is authenticating the user. It can be used instead of or in addition to NAS-Port in the Access-Request packet that is sent to the authentication server. For network access control (NAC) authentication specify the NAS-Port-Type as "Ethernet". For device

TABLE 47 Supported RADIUS attributes for accounting (continued)

Attribute ID	Attribute Name	Data Type	Description
			management authentication (when the user attempts to access the device by way of Telnet or SSH), specify the NAS-Port-Type as "Virtual". (RFC 2865).
95	NAS-IPv6-Address	address	IPv6 address of the NAS that is requesting user authentication (RFC 3162).

Dynamic VLAN assignment for 802.1X ports

The Extreme 802.1X implementation supports assigning a port to a VLAN dynamically, based on information received from an authentication server (RADIUS server).

When a client or supplicant successfully completes the EAP authentication process, the authentication server (RADIUS server) sends the authenticator (the device) a RADIUS Access-Accept message that grants the client access to the network. The RADIUS Access-Accept message contains attributes set for the user in the user's access profile on the RADIUS server.

If one of the attributes in the Access-Accept message specifies a VLAN identifier (ID), and this VLAN is available on the device, the client's port is moved from its default VLAN to the specified VLAN. When the client disconnects from the network, the port is placed back in its default VLAN.

> **NOTE**
> This feature is supported on port-based VLANs only. This feature cannot be used to place an 802.1X-enabled port into a Layer 3 protocol VLAN.

To enable 802.1X VLAN ID support on the device, you must add the following attributes to a user's profile on the RADIUS server.

TABLE 48 RADIUS attributes for dynamic VLAN assignment

Attribute name	Type	Value
Tunnel-Type	064	13 (decimal) - VLAN
Tunnel-Medium-Type	065	6 (decimal) - 802
Tunnel-Private-Group-ID	081	*vlan-name* (string) - Either the name or the number of a VLAN configured on the device.

The device reads the attributes as follows:

- If the Tunnel-Type or Tunnel-Medium-Type attributes in the Access-Accept message do not have the values specified above, the device ignores these Attribute-Value pairs. The client becomes authorized, but the client's port is not dynamically placed in a VLAN.

- If the Tunnel-Type or Tunnel-Medium-Type attributes (or both) are not present, then the client is moved to the unauthorized state displaying an error message on the device.

- If the Tunnel-Type or Tunnel-Medium-Type attributes in the Access-Accept message have the values specified in the table, but there is no value specified for the Tunnel-Private-Group-ID attribute, the client will not become authorized.

- When the device receives the value specified for the Tunnel-Private-Group-ID attribute, it checks whether the *vlan-name* string matches the name of a VLAN configured on the device. If there is a VLAN on the device with a name that matches the *vlan-name* string, then the client's port is placed in the VLAN with an ID that corresponds to the VLAN name.

- If the *vlan-name* string does not match the name of a VLAN, the device checks whether the string, when converted to a number, matches the ID of a VLAN configured on the device. If the ID matches, then the client's port is placed in the VLAN with that ID.

- If the *vlan-name* string does not match either the name or the ID of a VLAN configured on the device, then the client will not become authorized.

The **show interface** command displays the VLAN to which an 802.1X-enabled port has been dynamically assigned, as well as the port from which it was moved (that is, the port's default VLAN). Refer to the Displaying dynamically assigned VLAN information task for sample output indicating the port's dynamically assigned VLAN.

Considerations for dynamic VLAN assignment in an 802.1X multiple client configuration

The following considerations apply when a client in an 802.1X multiple client configuration is successfully authenticated, and the RADIUS Access-Accept message specifies a VLAN for the port:

- If the port is not already a member of a RADIUS-specified VLAN, and the RADIUS Access-Accept message specifies the name or ID of a valid VLAN on the device, then the port is placed in that VLAN.
- If the port is already a member of a RADIUS-specified VLAN, and the RADIUS Access-Accept message specifies the name or ID of a different VLAN, then it is considered an authentication failure. The port's VLAN membership is not changed.
- If the port is already a member of a RADIUS-specified VLAN, and the RADIUS Access-Accept message specifies the name or ID of that same VLAN, then traffic from the client is forwarded normally.
- If the RADIUS Access-Accept message specifies the name or ID of a VLAN that does not exist on the device, then it is considered an authentication failure.
- If the RADIUS Access-Accept message does not contain any VLAN information, the client's dot1x-mac-session is set to "access-is-allowed". If the port is already in a RADIUS-specified VLAN, it remains in that VLAN.

Disabling and enabling strict security mode for dynamic filter assignment

By default, 802.1x dynamic filter assignment operates in strict security mode. When strict security mode is enabled, 802.1x authentication for a port fails if the Filter-ID attribute contains invalid information, or if insufficient system resources are available to implement the per-user IP ACLs or MAC address filters specified in the Vendor-Specific attribute.

When strict security mode is enabled:

- If the Filter-ID attribute in the Access-Accept message contains a value that does not refer to an existing filter (that is, a MAC address filter or IP ACL configured on the device), then the client will not be authenticated, regardless of any other information in the message (for example, if the Tunnel-Private-Group-ID attribute specifies a VLAN to which to assign the port).
- If the Vendor-Specific attribute specifies the syntax for a filter, but there are insufficient system resources to implement the filter, then the port will not be authenticated.
- If the device does not have the system resources available to dynamically apply a filter to a port, then the port will not be authenticated.

  NOTE
  If the Access-Accept message contains values for both the Filter-ID and Vendor-Specific attributes, then the value in the Vendor-Specific attribute (the per-user filter) takes precedence.Also, if authentication for a port fails because the Filter-ID attribute referred to a non-existent filter, or there were insufficient system resources to implement the filter, then a Syslog message is generated.

When strict security mode is disabled:

- If the Filter-ID attribute in the Access-Accept message contains a value that does not refer to an existing filter (that is, a MAC address filter or IP ACL configured on the device), then the port is still authenticated, but no filter is dynamically applied to it.
- If the Vendor-Specific attribute specifies the syntax for a filter, but there are insufficient system resources to implement the filter, then the port is still authenticated, but the filter specified in the Vendor-Specific attribute is not applied to the port.

By default, strict security mode is enabled for all 802.1x-enabled interfaces, but you can manually disable or enable it, either globally or for specific interfaces.

To disable strict security mode globally, enter the following commands.

```
device(config)# dot1x-enable
device(config-dot1x)# no global-filter-strict-security
```

After you have globally disabled strict security mode on the device, you can re-enable it by entering the following command.

```
device(config-dot1x)# global-filter-strict-security
```

Syntax: [no] global-filter-strict-security

To disable strict security mode for a specific interface, enter commands such as the following.

```
device(config)# interface e 1
device(config-if-e10000-1)# no dot1x filter-strict-security
```

To re-enable strict security mode for an interface, enter the following command.

```
device(config-if-e10000-1)# dot1x filter-strict-security
```

Syntax: [no] dot1x filter-strict-security

The output of the **show dot1x** and **show dot1x config** commands has been enhanced to indicate whether strict security mode is enabled or disabled globally and on an interface.

Dynamically applying existing ACLs or MAC ACL

When a port is authenticated using 802.1X security, an IP ACL or MAC address filter that exists in the running configuration on the device can be dynamically applied to the port. To do this, you configure the Filter-Id (type 11) attribute on the RADIUS server. The Filter-Id attribute specifies the name or number of the IP ACL or MAC address filter.

The following table shows the syntax for configuring the Filter-Id attribute to refer to an IP ACL or MAC address filter.

TABLE 49 Syntax for Filter-Id attribute

Value	Description
ip.*number* .in	Applies the specified numbered ACL to the 802.1X authenticated port in the inbound direction.
ip.*name* .in	Applies the specified named ACL to the 802.1X authenticated port in the inbound direction.
ip.*number* .out	Applies the specified numbered ACL to the 802.1X authenticated port in the outbound direction.
ip.*name* .out	Applies the specified named ACL to the 802.1X authenticated port in the outbound direction.
mac.*number* .in	Applies the specified MAC address filter to the 802.1X authenticated port in the inbound direction.

The following table lists examples of values you can assign to the Filter-Id attribute on the RADIUS server to refer to IP ACLs and MAC address filters configured on a device.

TABLE 50 Example for Filter-Id attribute values

Possible values for the filter ID attribute on the RADIUS server	ACL or MAC address filter configured **on the** device
ip.2.in	access-list 2 permit host 10.48.0.3access-list 2 permit 10.0.0.0 10.255.255.255

TABLE 50 Example for Filter-Id attribute values (continued)

Possible values for the filter ID attribute on the RADIUS server	ACL or MAC address filter configured **on the** device
ip.102.in	access-list 102 permit ip 10.0.0.0 10.255.255.255 any
ip.fdry_filter.in	ip access-list standard fdry_filter permit host 10.48.0.3
mac.401.in	access-list 401 permit 3333.3333.3333 ffff.ffff.ffff any etype any
mac.402.in	access-list 402 permit 3333.3333.3333 ffff.ffff.ffff any etype any
mac.403.in	access-list 403 permit 2222.2222.2222 ffff.ffff.ffff any etype any

> **NOTE**
> - The *name* variable in the Filter-Id attribute is case-sensitive.
> - You can specify only numbered MAC address filters in the Filter-Id attribute. Named MAC address filters are not supported.
> - Dynamic IP ACL filters are supported for the inbound and outbound directions.
> - MAC address filters are supported only for the inbound direction. Outbound MAC address filters are not supported.
> - Dynamically assigned IP ACLs and MAC address filters are subject to the same configuration restrictions as non-dynamically assigned IP ACLs and MAC address filters.
> - Multiple IP ACLs and MAC address filters can be specified in the Filter-Id attribute, allowing multiple address filters to be simultaneously applied to an 802.1x authenticated port. Use commas, semicolons, or carriage returns to separate the address filters (for example, ip.3.in,mac.402.in).
> - If 802.1X is enabled on a VE port, ACLs, dynamic (802.1X assigned) or static (user configured), cannot be applied to the port.

Configuring per-user IP ACLs or MAC address filters

Per-user IP ACLs and MAC address filters make use of the Vendor-Specific (type 26) attribute to dynamically apply filters to ports. Defined in the Vendor-Specific attribute are Extreme ACL or MAC address filter statements. When the RADIUS server returns the Access-Accept message granting a client access to the network, the device reads the statements in the Vendor-Specific attribute and applies these IP ACLs or MAC address filters to the client's port. When the client disconnects from the network, the dynamically applied filters are no longer applied to the port. If any filters had been applied to the port previous to the client connecting, then those filters are reapplied to the port.

The following is the syntax for configuring the Extreme Vendor-Specific attribute with ACL or MAC address filter statements.

Value	Description
ipacl.e.in=*extended-acl-entries*	Applies the specified extended ACL entries to the 802.1x authenticated port in the inbound direction.
ipacl.e.out=*extended-acl-entries*	Applies the specified extended ACL entries to the 802.1x authenticated port in the outbound direction.
macfilter.in=*mac-access list-entries*	Applies the specified MAC address filter entries to the 802.1x authenticated port in the inbound direction.

The following table shows examples of IP ACLs and MAC address filters configured in the Extreme Vendor-Specific attribute on a RADIUS server. These IP ACLs and MAC address filters follow the same syntax as other Extreme ACLs and MAC address filters. Refer to the *Extreme NetIron Layer 2 Switching Configuration Guide* for information on syntax.

IP ACL or MAC address filter	Vendor-specific attribute on RADIUS server
Extended ACL with one entry (outbound direction)	ipacl.e.out=permit ip 10.0.0.0 10.255.255.255 any

IP ACL or MAC address filter	Vendor-specific attribute on RADIUS server
Mac address filter with one entry	macfilter.in= deny any any
Mac address filter with two entries	macfilter.in= permit 0000.0000.3333 ffff.ffff.0000 any, macfilter.in= permit 0000.0000.4444 ffff.ffff.0000 any

The RADIUS server allows one instance of the Vendor-Specific attribute to be sent in an Access-Accept message. However, the Vendor-Specific attribute can specify multiple IP ACLs or MAC address filters. You can use commas, semicolons, or carriage returns to separate the filters (for example: ipacl.e.in= permit ip any any,ipacl.e.in = deny ip any any).

Enabling 802.1x port security

By default, 802.1x port security is disabled on devices. To enable the feature on the device and enter the dot1x configuration level, enter the following command.

```
device(config)# dot1x-enable
device(config-dot1x)#
```

Syntax: [no] dot1x-enable

At the dot1x configuration level, you can enable 802.1x port security on all interfaces at once, on individual interfaces, or on a range of interfaces.

For example, to enable 802.1x port security on all interfaces on the device, enter the following command.

```
device(config-dot1x)# enable all
```

Syntax: [no] enable all

To enable 802.1x port security on interface 3/11, enter the following command.

```
device(config-dot1x)# enable ethernet 3/11
```

Syntax: [no] enable *portnum*

To enable 802.1x port security on interfaces 3/11 through 3/16, enter the following command.

```
device(config-dot1x)# enable ethernet 3/11 to 3/16
```

Syntax: [no] enable *portnum* to *portnum*

Setting the port control

To activate authentication on an 802.1x-enabled interface, you specify the kind of port control to be used on the interface. An interface used with 802.1x port security has two virtual access points:

- The controlled port can be either the authorized or unauthorized state. In the authorized state, it allows normal traffic to pass between the client and the authenticator. In the unauthorized state, it allows no traffic to pass through.
- The uncontrolled port allows only EAPOL traffic between the client and the authentication server.

Refer to Controlled and uncontrolled ports on page 369 for an illustration of this concept.

By default, all controlled ports on the device are in the authorized state, allowing all traffic. When you activate authentication on an 802.1x-enabled interface, its controlled port is placed in the unauthorized state. When a client connected to the interface is successfully authenticated, the controlled port is then placed in the authorized state for that client. The controlled port remains in the authorized state until the client logs off.

To activate authentication on an 802.1x-enabled interface, you configure the interface to place its controlled port in the authorized state when a client is authenticated by an authentication server. To do this, enter commands such as the following.

```
device(config)# interface e 3/1
device(config-if-e10000-3/1)# dot1x port-control auto
```

Syntax: [no] dot1x port-control [force-authorized | force-unauthorized | auto]

When an interface's control type is set to **auto** , its controlled port is initially set to unauthorized, but is changed to authorized when the connecting client is successfully authenticated by an Authentication Server.

The port control type can be one of the following:

force-authorized - The port's controlled port is placed unconditionally in the authorized state, allowing all traffic. This is the default state for ports on the device. Also, this parameter allows connection from multiple clients.

force-unauthorized - The controlled port is placed unconditionally in the unauthorized state.

auto - The controlled port is unauthorized until authentication takes place between the client and Authentication Server. Once the client passes authentication, the port becomes authorized. This has the effect of activating authentication on an 802.1x-enabled interface.

Notes

You cannot enable 802.1x port security on ports that have any of the following features enabled:

- Static MAC configurations
- Link aggregation
- Metro Ring Protocol (Brocade MRP)
- Tagged port
- Mirror port
- LAG port
- MAC port security
- Management Port
- VE members

Periodic reauthentication

The re-authentication interval is a global setting, applicable to all 802.1X-enabled interfaces. If you want to re-authenticate clients connected to a specific port manually, use the **dot1x re-authenticate** command. Refer to

Manual reauthentication of a port

When periodic reauthentication is enabled, the device reauthenticates clients connected to an 802.1X-enabled interface every 3,600 seconds (or the time specified by the **dot1x timeout re-authperiod** command) by default. You can also manually reauthenticate clients connected to a specific port.

Quiet period for reauthentication

If the device is unable to authenticate the client, the device waits for a specified amount of time before trying again. The amount of time the device remains idle between a failed authentication and a reauthentication attempt is specified with the **quiet-period** parameter. This timer also indicates how long a client that failed authentication will have its blocked entry programmed into the hardware.

Retransmission interval for EAP-Request/Identity frames

When the device sends a client an EAP-Request/Identity frame, it expects to receive an EAP-Response/Identity frame from the client. If the client does not send back an EAP-Response/Identity frame, the device waits a specified amount of time, and then retransmits the EAP-Request/Identity frame. You can specify the amount of time the device waits before retransmitting the EAP-Request/Identity frame to the client. This amount of time is specified using the **tx-period** parameter.

Retransmission limit for EAP-Request/Identity frame

If the device does not receive an EAP-Response/Identity frame from a client, the device waits 30 seconds (or the amount of time specified with the **timeout tx-period** command), and then retransmits the EAP-Request/Identity frame. By default, the device retransmits the EAP-Request/Identity frame a maximum of two times. If no EAP-Response/Identity frame is received from the client after two EAP-Request/Identity frame retransmissions, the device restarts the authentication process with the client.

Specifying a timeout for retransmission of messages to the Authentication Server

When performing authentication, the device receives EAPOL frames from the client and passes the messages on to the RADIUS server. The device expects a response from the RADIUS server within 30 seconds. If the RADIUS server does not send a response within 30 seconds, the device retransmits the message to the RADIUS server. The time constraint for retransmission of messages to the Authentication Server can be between 1 - 4294967295 seconds.

For the device, the possible values are: 1 - 4294967295.

For example, to configure the device to retransmit a message if the Authentication Server does not respond within 45 seconds, enter the following command.

```
device(config-dot1x)# servertimeout 45
```

Syntax: **servertimeout** *seconds*

Retransmission timeout of EAP-Request frames to the client

Acting as an intermediary between the RADIUS authentication server and the client, the device receives RADIUS messages from the RADIUS server, encapsulates them as EAPOL frames, and sends them to the client. When the device relays an EAP-Request frame from the RADIUS server to the client, it expects to receive a response from the client within 30 seconds. If the client does not respond within the allotted time, the device retransmits the EAP-Request frame to the client. The timeout value for retransmission of EAP-Request frames to the client can be from 1 through 4294967295 seconds.

Initializing 802.1x on a port

To initialize 802.1x port security on a port, or to flush all of its information on that port and start again, enter a command such as the following.

```
device# dot1x initialize e 3/1
```

Syntax: dot1x initialize *portnum*

Allowing multiple 802.1x clients to authenticate

If there are multiple clients connected to a single 802.1x-enabled port, the device authenticates each of them individually. When multiple clients are connected to the same 802.1x-enabled port, the functionality described in How 802.1x multiple client authentication works on page 373 is enabled by default. You can optionally do the following:

- Specify the authentication-failure action
- Specify the number of authentication attempts the device makes before dropping packets
- Disabling aging for dot1x-mac-sessions
- Configure aging time for blocked clients
- Clear the dot1x-mac-session for a MAC address

Specifying the authentication-failure action

In an 802.1x multiple client configuration, if RADIUS authentication for a client is unsuccessful, traffic from that client is either dropped in hardware (the default), or the client's port is placed in a "restricted" VLAN. You can specify which of these two authentication-failure actions is to be used. If the authentication-failure action is to place the port in a restricted VLAN, you can specify the ID of the restricted VLAN.

To specify that the authentication-failure action is to place the client's port in a restricted VLAN, enter the following command.

```
device(config)# dot1x-enable
device(config-dot1x)# auth-fail-action restricted-vlan
```

Syntax: [no] auth-fail-action restricted-vlan

To specify the ID of the restricted VLAN as VLAN 300, enter the following command.

```
device(config-dot1x)# auth-fail-vlanid 300
```

Syntax: [no] auth-fail-vlanid *vlan-id*

Specifying the number of authentication attempts the device makes before dropping packets

When the initial authentication attempt made by the device to authenticate the client is unsuccessful, the device immediately retries to authenticate the client. After three unsuccessful authentication attempts, the client's802.1x MAC authentication session is set to either "access-denied" or the port is moved to restricted VLAN.

You can optionally configure the number of authentication attempts the device makes. To do so, enter a command such as the following.

```
device(config-dot1x)# auth-fail-max-attempts 2
```

Syntax: [no] auth-fail-max-attempts *attempts*

By default, the device makes 3 attempts to authenticate a client. You can specify between 1 - 10 authentication attempts.

Display commands

The **show port security global-deny** command lists all the configured global deny MAC addresses.

The **show port security denied-macs** command lists all the denied MAC addresses in the system.

Clearing a dot1x-mac-session for a MAC address

You can clear the dot1x-mac-session for a specified MAC address, so that the client with that MAC address can be re-authenticated by the RADIUS server.

```
device# clear dot1x mac-session 00e0.1234.abd4
```

Syntax: clear dot1x mac-session *mac-address*

Displaying 802.1x information

Various show commands can be used to display the following 802.1x-related information:

- Information about the 802.1x configuration on the device and on individual ports
- Statistics about the EAPOL frames passing through the device
- Information about 802.1x-enabled ports dynamically assigned to a VLAN
- Information about the dynamically applied MAC and IP ACLs currently active on the device
- Information about the 802.1x multiple client configuration

Enter the **show dot1x** command to display the overall state of 802.1X authentication on the system.

```
device# show dot1x

802.1X Port-Based Authentication: Enabled
PAE Capability:                   Authenticator Only
Protocol Version:                 2
Auth Server:                      RADIUS
Readiness test timeout:           10
RADIUS Configuration
--------------------
Position:               1
Server Address:         10.24.65.6
Port:                   1812
Secret:                 xxxxxxxxx
Retry Interval:         5 seconds
```

Enter the **show dot1X all** command to display detailed 802.1X authentication information for all of the ports.

```
device# show dot1x all

802.1X Port-Based Authentication: Enabled
PAE Capability:                   Authenticator Only
Protocol Version:                 2
Auth Server:                      RADIUS
Readiness test timeout:           10

RADIUS Configuration
--------------------
Position:               1
Server Address:         10.20.106.144
Port:                   1812
Secret:                 testing123
Retry Interval:         4 seconds
Position:               2
Server Address:         10.20.106.189
```

```
Port:                    1812
Secret:                  testing123
Retry Interval:          4 seconds

802.1X info for interface Eth 1/31
----------------------------------
Port Control:            Auto
Protocol Version:        2
ReAuthentication:        Enabled
Auth Fail Max Attempts:  0
ReAuth Max:              2
Tx Period:               30 seconds
Quiet Period:            60 seconds
Supplicant Timeout:      30 seconds
Re-Auth Interval:        3600 seconds
Dynamic VLAN assigned:   50
Filter-strict-security:  Enabled
IP ACL assigned (IN|OUT): IPEXT-50 | IPEXT-OUT-50
MAC ACL assigned:        mac-ext
```

Enter the **show dot1x diagnostics interface** command to display all diagnostics information for the authenticator associated with a port.

```
device# show dot1x diagnostics interface ethernet 1/2

802.1X Diagnostics for interface Eth 1/2
----------------------------------------
authEnterConnecting:                1
authEaplogoffWhileConnecting:       0
authEnterAuthenticating:            1
authSuccessWhileAuthenticating:     1
authTimeoutWhileAuthenticating:     0
authFailWhileAuthenticating:        0
authEapstartWhileAuthenticating:    0
authEaplogoffWhileAuthenticating:   0
authReauthsWhileAuthenticated:      0
authEapstartWhileAuthenticated:     0
authEaplogoffWhileAuthenticated:    0
BackendResponses:                   11
BackendAccessChallenges:            10
BackendOtherrequestToSupplicant:    11
BackendAuthSuccess:                 1
BackendAuthFails:                   0
```

Enter the **show dot1x interface** command to display state of a specified interface.

```
show dot1x interface ethernet 1/31

802.1X info for interface Eth 1/31
----------------------------------
Port Control:            Auto
Protocol Version:        2
ReAuthentication:        Enabled
Auth Fail Max Attempts:  0
ReAuth Max:              2
Tx Period:               30 seconds
Quiet Period:            60 seconds
Supplicant Timeout:      30 seconds
Server Timeout:          30 seconds
Re-Auth Interval:        3600 seconds
Dynamic VLAN assigned:   50
Filter-strict-security:  Enabled
IP ACL assigned (IN|OUT): IPEXT-50 | IPEXT-OUT-50
MAC ACL assigned:        mac-ext
```

Enter the **show dot1x session-info interface** command to display information for all clients on the port .

```
device# show dot1x session-info interface ethernet 1/2

802.1X Session info for interface Eth 1/2
----------------------------------------
Mac Address: 0021.5ec6.15ce
```

```
----------------------------------------
User Name:              md5user2
Session Time:           2 secs
Terminate Cause:        Not terminated yet
Session Status:         Authorized
PAE State:              Authenticated
BE State:               Idle
VLAN:                   N/A
IP ACL (IN | OUT):      N/A | N/A
MAC ACL:                N/A
Current Id:             18
Id From Server:         17
```

Enter the **show dot1x statistics interface** command to display the statistics of a specified interface.

```
device# show dot1x statistics interface ethernet 1/2

802.1X statistics for interface Eth 1/2
----------------------------------------
EAPOL Frames Rx:                  12
EAPOL Frames Tx:                  43
EAPOL Start Frames Rx:            1
EAPOL Logoff Frames Rx:           0
EAP Rsp/Id Frames Rx:             1
EAP Response Frames Rx:           10
EAP Req/Id Frames Tx:             23
EAP Request Frames Tx:            10
Invalid EAPOL Frames Rx:          0
EAPOL Length Error Frames Rx:     0
EAPOL Last Frame Version Rx:      1
Invalid EAP Frames Rx:            0
EAP Length Error Frames Rx:       0
EAPOL Last Frame Src:             0021.5ec6.15ce
```

Displaying 802.1x configuration information

To display information about the 802.1x configuration on the device, enter the following command.

```
device# show dot1x
PAE Capability            : Authenticator Only
system-auth-control       : Enable
Number of ports enabled   : 25
re-authentication         : Disable
global-filter-strict-security: Enable
quiet-period              : 60 Seconds
tx-period                 : 30 Seconds
supptimeout               : 30 Seconds
servertimeout             : 30 Seconds
maxreq                    : 3
re-authperiod             : 3600 Seconds
Protocol Version          : 1
auth-fail-action          : Block Traffic
MAC Session Aging         : All
MAC Session Max Age       : 120 Seconds
Maximum Failed Attempts   : 3
```

Syntax: **show dot1x**

The following table describes the information displayed by the **show dot1x** command.

This field...	Displays...
PAE Capability	The Port Access Entity (PAE) role for the device. This is always "Authenticator Only".
system-auth-control	Whether system authentication control is enabled on the device. The **dot1x-enable** command enables system authentication control on the device.

This field...	Displays...
Number of ports enabled	Number of interfaces on the devices that have been enabled for 802.1x.
re-authentication	Whether periodic re-authentication is enabled on the device. Refer to Periodic reauthentication on page 384. When periodic re-authentication is enabled, the device automatically re-authenticates clients every 3,600 seconds by default.
global-filter-strict-security	Whether or not strict security mode is enabled globally.
quiet-period	When the device is unable to authenticate a client, the amount of time the device waits before trying again (default 60 seconds). Refer to Quiet period for reauthentication on page 385 for information on how to change this setting.
tx-period	When a client does not send back an EAP-response or identity frame, the amount of time the device waits before retransmitting the EAP-request or identity frame to a client (default 30 seconds). Refer to Retransmission interval for EAP-Request/Identity frames on page 385 for information on how to change this setting.
supp-timeout	When a client does not respond to an EAP-request frame, the amount of time before the device retransmits the frame. Refer to Retransmission timeout of EAP-Request frames to the client on page 385 for information on how to change this setting.
server-timeout	When the Authentication Server does not respond to a message sent from the client, the amount of time before the device retransmits the message. Refer to Specifying a timeout for retransmission of messages to the Authentication Server on page 385 for information on how to change this setting.
max-req	The number of times the device retransmits an EAP-request or identity frame if it does not receive an EAP-response or identity frame from a client (default 2 times). Refer to Retransmission limit for EAP-Request/Identity frame on page 385 for information on how to change this setting.
re-authperiod	How often the device automatically re-authenticates clients when periodic re-authentication is enabled (default 3,600 seconds). Refer to Periodic reauthentication on page 384 for information on how to change this setting.
security-hold-time	This field is not supported.
Protocol Version	The version of the 802.1x protocol in use on the device.
Auth-fail-action	The configured authentication-failure action. This can be Restricted VLAN or Block Traffic.
Mac Session Aging	Whether aging for dot1x-mac-sessions has been enabled or disabled for permitted or denied dot1x-mac-sessions.
Mac Session max-age	The configured software aging time for dot1x-mac-sessions.
Maximum Failed Attempts	The number of failed authentication attempts, if the authentication-failure action shows Restricted VLAN,

To display information about the 802.1x configuration on an individual port, enter a command such as the following.

```
device# show dot1x config e 1/3
Port 1/3 Configuration:
AuthControlledPortControl    : Auto
max-clients                  : 32
```

```
multiple-clients         : Enable
filter-strict-security   : Enable
```

Syntax: show dot1x config ethernet slot/port

The following additional information is displayed in the **show dot1x config** command for an interface.

This field...	Displays...
AuthControlledPortControl	The port control type configured for the interface. If set to auto, authentication is activated on the 802.1x-enabled interface.
multiple-hosts	Whether the port is configured to allow multiple Supplicants accessing the interface on the device through a hub.
	Refer to Allowing multiple 802.1x clients to authenticate on page 386 for information on how to change this setting.
max-clients	The maximum number of clients that can be authenticated on this interface.
multiple-clients	Shows if the interface is enabled or disabled for multiple client authentication.
filter-strict-security	Shows if the interface is enabled or disabled for strict security mode.

Displaying 802.1x statistics

To display 802.1x statistics for an individual port, enter a command such as the following.

```
device# show dot1x statistics e 3/3
Port 1/3 Statistics:
RX EAPOL Start:             0
RX EAPOL Logoff:            0
RX EAPOL Invalid:           0
RX EAPOL Total:             2
RX EAP Resp/Id:             1
RX EAP Resp other than Resp/Id: 1
RX EAP Length Error:        0
Last EAPOL Version:         1
Last EAPOL Source:          0050.da0b.8bef
TX EAPOL Total:             3
TX EAP Req/Id:              1
TX EAP Req other than Req/Id:  1
Num Sessions:               1
Num Restricted Sessions:    0
Num Authorized Sessions:    1
```

Syntax: show dot1x statistics [all | ethernet slot/port]

The following table describes the information displayed by the **show dot1x statistics** command for an interface.

This field...	Displays...
RX EAPOL Start	The number of EAPOL-Start frames received on the port.
RX EAPOL Logoff	The number of EAPOL-Logoff frames received on the port.
RX EAPOL Invalid	The number of invalid EAPOL frames received on the port.
RX EAPOL Total	The total number of EAPOL frames received on the port.
RX EAP Resp or Id	The number of EAP-Response or Identity frames received on the port
RX EAP Resp other than Resp or Id	The total number of EAPOL-Response frames received on the port that were not EAP-Response or Identity frames.
RX EAP Length Error	The number of EAPOL frames received on the port that have an invalid packet body length.

This field...	Displays...
Last EAPOL Version	The version number of the last EAPOL frame received on the port.
Last EAPOL Source	The source MAC address in the last EAPOL frame received on the port.
TX EAPOL Total	The total number of EAPOL frames transmitted on the port.
TX EAP Req or Id	The number of EAP-Request or Identity frames transmitted on the port.
TX EAP Req other than Req or Id	The number of EAP-Request frames transmitted on the port that were not EAP-Request or Identity frames.
Num sessions	Total number of dot1x sessions, which include authenticated, restricted, denied and sessions in the initial state.
Num Restricted Sessions	Number of current 802.1x sessions that failed authentication. The user configuration was moved into a restricted VLAN.
Num Authorized Sessions	Number of current 802.1x authenticated sessions that are authorized.

Clearing 802.1x statistics

You can clear the 802.1x statistics counters on all interfaces at once, on individual interfaces, or on a range of interfaces.

For example, to clear the 802.1x statistics counters on all interfaces on the device, enter the following command.

```
device# clear dot1x statistics all
```

Syntax: clear dot1x statistics all

To clear the 802.1x statistics counters on interface e 3/11, enter the following command.

```
device# clear dot1x statistics e 3/11
```

Syntax: clear dot1x statistics [mac-address | ethernet slot/portnum]

Displaying dynamically assigned VLAN information

The **show interface** command displays the VLAN to which an 802.1x-enabled port has been dynamically assigned, as well as the port from which it was moved (that is, the port's default VLAN).

The following is an example of the **show interface** command indicating the port's dynamically assigned VLAN. Information about the dynamically assigned VLAN is shown in bold type.

```
device# show interface e 1/3
GigabitEthernet1/3 is up, line protocol is up
STP Root Guard is disabled, STP BPDU Guard is disabled
Hardware is GigabitEthernet, address is 000c.dbe2.5800 (bia 000c.dbe2.5800)
Configured speed auto, actual 100Mbit, configured duplex fdx, actual fdx
Member of L2 VLAN ID 4094 (dot1x-RADIUS assigned), original L2 VLAN ID is 1,
port is untagged, port state is Forwarding
STP configured to ON, Priority is level0, flow control enabled
Priority force disabled, Drop precedence level 0, Drop precedence force disabled dhcp-snooping-trust
configured to OFF
mirror disabled, monitor disabled
Not member of any active trunks
Not member of any configured trunks
No port name
Internet address is 10.12.12.250/24, MTU 1522 bytes, encapsulation ethernet
300 second input rate: 810 bits/sec, 0 packets/sec, 0.00% utilization
300 second output rate: 1253 bits/sec, 1 packets/sec, 0.00% utilization
70178 packets input, 7148796 bytes, 0 no buffer
Received 0 broadcasts, 0 multicasts, 70178 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants, DMA received 70178 packets
NP received 555904 packets, Sent to TM 555900 packets
```

```
NP Ingress dropped 4 packets
91892 packets output, 10081165 bytes, 0 underruns
Transmitted 9853 broadcasts, 13330 multicasts, 68709 unicasts
0 output errors, 0 collisions, DMA transmitted 91892 packets
NP transmitted 4278949 packets, Received from TM 4768301 packets
```

In this example, the 802.1x-enabled port has been moved from VLAN 1 to VLAN 4094. When the client disconnects, the port will be moved back to VLAN 1.

Displaying information on MAC address filters and IP ACLs on an interface

You can display information about the user-defined and dynamically applied MAC address filters and IP ACLs currently active on an interface.

Displaying MAC address filters applied to an 802.1x-enabled port

Use the **show dot1x mac-address** command to display information about MAC filters applied to an interface. If the MAC address filter is dynamically assigned by 802.1x, the display shows the following.

```
device#show dot1x mac-address ethernet 1/1
Port 1/1 MAC Address Filter information:
  802.1x dynamic MAC Filter (user defined) :
    mac access-list 401 in
  Port default MAC Filter :
    mac access-list 400 in
```

The "Port default MAC Filter" appears if a default MAC filter has been configured on the port. This default MAC filter is the MAC filter that will be applied to the port once the dynamically assigned MAC filter is removed. If a default MAC filter has not been configured, the message "No Port default MAC is displayed.

When the dynamically assigned MAC address filter is removed, the display shows the following information.

```
device#show dot1x mac-address ethernet 1/1
Port 1/1 MAC Address Filter information:
  Port default MAC Filter :
    mac access-list 400 in
```

Syntax: show dot1x mac-address-filter [**all** | **ethernet** slot/port | | **begin** expression | **exclude** expression | **include** expression]

The **all** keyword displays all dynamically applied MAC address filters active on the device.

Use the **ethernet**slot/port parameter to display information for one port.

Displaying IP ACLs applied to an 802.1x-enabled port

Use the **show dot1x ip-acl** command to display the information about what IP ACLs have been applied to an 802.1x-enabled port. If the IP ACL was dynamically applied by 802.1x, the following information is displayed.

```
device# show dot1x ip-acl ethernet 1/1
Port 1/1 IP ACL information:
  802.1x dynamic IP ACL (user defined) in:
    ip access-list extended Port_1/1_E_IN in
  Port default IP ACL in:
    ip access-list 100 in
  No outbound ip access-list is set
```

The "Port default IP ACL" appears if a default IP ACL has been configured on the port. The default IP ACL is the IP ACL that will be applied to the port once the dynamically assigned IP ACL is removed. If a default IP ACL has not been configured, the message "No Port default IP ACL" is displayed.

When the dynamically assigned IP ACL is removed from the port, the display shows the following information.

```
device# show dot1x ip-acl ethernet 1/1
Port 1/1 IP ACL information:
  Port default IP ACL in:
    ip access-list 100 in
  No outbound ip access-list is set
```

Syntax: show dot1x ip-acl [all | ethernet slot/port | | begin expression | exclude expression | include expression]

The **all** keyword displays all dynamically applied IP ACLs active on the device.

Use the **ethernet**_slot/port_ parameter to display information for one port.

Displaying information about the dot1x-mac-sessions on each port

To display information about the dot1x-mac-sessions on each port on the device, enter the following command.

```
device# show dot1x mac-session
Port  MAC             Username                  VLAN Auth State ACL|MAC  Age
                                                               i|o|f
-------------------------------------------------------------------------
1/1   0050.da0b.8cd7  Mary M                       1    DENIED  n|n|n   0
1/2   0050.da0b.8cb3  adminmorn                 4094 PERMITTED  y|n|n   0
1/3   0050.da0b.8bef  reports                   4094 PERMITTED  y|n|n   0
1/4   0010.5a1f.6a63  testgroup                 4094 PERMITTED  y|n|n   0
1/5   0050.da1a.ff7e  admineve                  4094 PERMITTED  y|n|n   0
```

Syntax: show dot1x mac-session [brief | [begin expression | exclude expression | include expression]]

The table below describes the information displayed by the **show dot1x mac-session** command.

This field...	Displays...
Port	The port on which the dot1x-mac-session exists.
MAC	The MAC address of the client
Username	The username used for RADIUS authentication.
Vlan	The VLAN to which the port is currently assigned.
Auth-State	The authentication state of the dot1x-mac-session. This can be one of the following:
	permit – The client has been successfully authenticated, and traffic from the client is being forwarded normally.
	blocked – Authentication failed for the client, and traffic from the client is being dropped in hardware.
	restricted – Authentication failed for the client, but traffic from the client is allowed in the restricted VLAN only.
	init – The client is in is in the process of 802.1x authentication, or has not started the authentication process.
ACL	Whether or not an IP ACL is applied to incoming (i) and outgoing (o) traffic on the interface
MAC	Whether or not a MAC filter is applied to the port.
Age	The software age of the dot1x-mac-session.

Displaying information about the ports in an 802.1x multiple client configuration

To display information about the ports in an 802.1x multiple client configuration, enter the following command.

```
device# show dot1x mac-session brief
Port            Number of users        Dynamic Dynamic    Dynamic
         Restricted Authorized Total  VLAN    ACL (In/Out) MAC-Filt
---------+----------+----------+-----+-------+----------+--------
1/1             0          0     1 no          no/no   no
1/2             0          1     1 yes        yes/no   no
1/3             0          1     1 yes        yes/no   no
1/4             0          1     1 yes        yes/no   no
1/5             0          1     1 yes        yes/no   no
```

Syntax: show dot1x mac-session brief [| **begin expression** | **exclude expression** | **include expression**]

The following table describes the information displayed by the **show dot1x mac-session brief** command.

This field...	Displays...
Port	Information about the users connected to each port.
Number of users	The number of restricted and authorized (those that were successfully authenticated) users connected to the port.
Dynamic VLAN	Whether or not the port is a member of a RADIUS-specified VLAN.
Dynamic ACL	Whether or not a RADIUS-specified ACL has been applied to the port for incoming (in) and outgoing (out) traffic.
Dynamic MAC Filters	Whether or not a RADIUS-specified MAC Filter has been applied to the port.

Sample 802.1x configurations

This section illustrates a sample point-to-point configuration and a sample hub configuration that use 802.1x port security.

Point-to-point configuration

Figure 21 illustrates a sample 802.1x configuration with clients connected to three ports on the device. In a point-to-point configuration, only one 802.1x client can be connected to each port.

FIGURE 21 Sample point-to-point 802.1x configuration
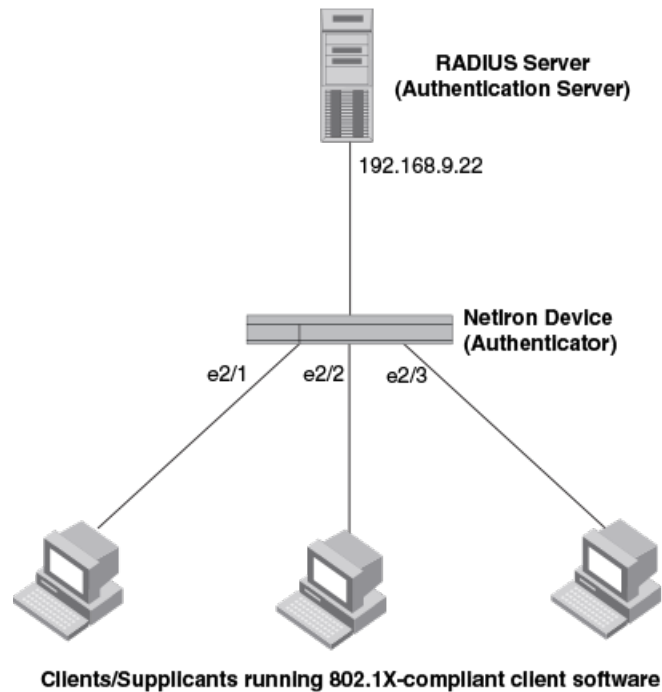
The following commands configure the device in Figure 21.

```
device(config)# aaa authentication dot1x default radius
device(config)# radius-server host 192.168.9.22 auth-port 1812 acct-port 1813 default key mirabeau dot1x
device(config)# dot1x-enable
device(config-dot1x)# enable e 2/1 to 2/3
device(config-dot1x)# re-authentication
device(config-dot1x)# timeout re-authperiod 2000
device(config-dot1x)# timeout quiet-period 30
device(config-dot1x)# timeout tx-period 60
device(config-dot1x)# max-req 6
device(config-dot1x)# exit
device(config)# interface e 2/1
device(config-if-e100-1)# dot1x port-control auto
device(config-if-e100-1)# exit
device(config)# interface e 2/2
device(config-if-e100-2)# dot1x port-control auto
device(config-if-e100-2)# exit
device(config)# interface e 2/3
device(config-if-e100-3)# dot1x port-control auto
device(config-if-e100-3)# exit
```

Hub configuration

Figure 22 illustrates a configuration where three 802.1x-enabled clients are connected to a hub, which is connected to a port on the device. The configuration is similar to that in Point-to-point configuration on page 395, except that 802.1x port security is enabled on only one port, and the **multiple-hosts** command is used to allow multiple clients on the port.

FIGURE 22 Sample 802.1x configuration using a hub

Clients/Supplicants running 802.1X-compliant client software

The following commands configure the device in .

```
device(config)# aaa authentication dot1x default radius
device(config)# radius-server host 192.168.9.22 auth-port 1812 acct-port 1813 default key mirabeau dot1x
device(config)# dot1x-enable e 2/1
device(config-dot1x)# re-authentication
device(config-dot1x)# timeout re-authperiod 2000
device(config-dot1x)# timeout quiet-period 30
device(config-dot1x)# timeout tx-period 60
device(config-dot1x)# max-req 6
device(config-dot1x)# exit
device(config)# interface e 2/1
device(config-if-e100-1)# dot1x port-control auto
device(config-if-e100-1)# exit
```

Configuring Multi-Device Port Authentication

Multi-device port authentication is a way to configure a device to forward or block traffic from a MAC address based on information received from a RADIUS server.

How multi-device port authentication works

The multi-device port authentication feature is a mechanism by which incoming traffic originating from a specific MAC address is switched or forwarded by the device only if the source MAC address is successfully authenticated by a RADIUS server. The MAC address itself is used as the username and password for RADIUS authentication; the user does not need to provide a specific username and password to gain access to the network. If RADIUS authentication for the MAC address is successful, traffic from the MAC address is forwarded in hardware.

If the RADIUS server cannot validate the user's MAC address, then it is considered an authentication failure, and a specified authentication-failure action can be taken. The default authentication-failure action is to drop traffic from the non-authenticated MAC address in hardware. You can also configure the device to move the port on which the non-authenticated MAC address was learned into a restricted or "guest" VLAN, which may have limited access to the network.

RADIUS authentication

The multi-device port authentication feature communicates with the RADIUS server to authenticate a newly found MAC address. The device supports multiple RADIUS servers; if communication with one of the RADIUS servers times out, the others are tried in sequential order. If a response from a RADIUS server is not received within a specified time (by default, 3 seconds) the RADIUS session times out, and the device retries the request up to three times. If no response is received, the next RADIUS server is chosen, and the request is sent for authentication.

The RADIUS server is configured with the usernames and passwords of authenticated users. For multi-device port authentication, the username and password is the MAC address itself; that is, the device uses the MAC address for both the username and the password in the request sent to the RADIUS server. For example, given a MAC address of 0007e90feaa1, the users file on the RADIUS server would be configured with a username and password both set to 0007e90feaa1. When traffic from this MAC address is encountered on a MAC-authentication-enabled interface, the device sends the RADIUS server an Access-Request message with 0007e90feaa1 as both the username and password. The format of the MAC address sent to the RADIUS server is configurable through the CLI.

The request for authentication from the RADIUS server is successful only if the username and password provided in the request matches an entry in the users database on the RADIUS server. When this happens, the RADIUS server returns an Access-Accept message back to the device. When the RADIUS server returns an Access-Accept message for a MAC address, that MAC address is considered authenticated, and traffic from the MAC address is forwarded normally by the device.

Authentication-failure actions

If the MAC address does not match the username and password of an entry in the users database on the RADIUS server, then the RADIUS server returns an Access-Reject message. When this happens, it is considered an authentication failure for the MAC address. When an authentication failure occurs, the device can either drop traffic from the MAC address in hardware (the default), or move the port on which the traffic was received to a restricted VLAN.

Extreme devices support multi-device port authentication on untagged ports only.

Supported RADIUS attributes

The Extreme devices support the following RADIUS attributes for multi-device port authentication and accounting.

TABLE 51 Supported RADIUS attributes for authentication

Attribute ID	Attribute name	Data Type	Description
1	User-Name	String	Indicates the name of the user to be authenticated (RFC 2865).
4	NAS-IP-Address	Address	IP address of the network access server (NAS) which is requesting user authentication (RFC 2865).
5	NAS-Port	Integer	Physical port number of NAS that is authenticating the user. It is only used in Access-Request packets (RFC 2865). **NOTE** The physical port number refers to a physical connection on the NAS rather than a TCP or UDP port number. Either NAS-Port, NAS-Port-Type (61),or both should be present in the Access-Request packet, when the NAS differentiates between its ports.
6	Service-Type	Integer	Type of service the user has requested, or the type of service to be provided. It may be used in both Access-Request and Access-Accept packets. A NAS is not required to implement all service types, but must treat unknown or unsupported Service-Types as though an Access-Reject packet has been received (RFC 2865).
7	Framed-Protocol	Integer	Framing protocol to be used for framed access. It may be used in both Access-Request and Access-Accept packets (RFC 2865).
11	Filter-Id	Text	Name of the filter list for this user (RFC 2865).
12	Framed-MTU	Integer	Maximum transmission unit (MTU) to be configured for the user, when it is not negotiated by some other means (such as Point-to-Point protocol (PPP)). It may be used in Access-Accept packets or in an Access-Request packet as a hint by the NAS to the server that it would prefer that value; the server is not required to honor the hint (RFC 2865).
24	State	String	This attribute is available to be sent by the server to the client: • In an Access-Challenge and must be sent unmodified from the client to the server in the new Access-Request reply (if any) to that challenge.

TABLE 51 Supported RADIUS attributes for authentication (continued)

Attribute ID	Attribute name	Data Type	Description
			• In an Access-Accept that also includes a Termination-Action attribute with the value of RADIUS-Request. When the NAS performs the Termination-Action by sending a new Access-Request upon termination of the current session, it must include the State attribute unchanged in that Access-Request.
			In either preceding usage, the client must not interpret the attribute locally.
			A packet must have only one State attribute or none. Usage of the State attribute is implementation dependent (RFC 2865).
26	Vendor-Specific	String	This attribute allows vendors to support their own extended attributes (RFC 2865).
31	Calling-Station-Id	String	This attribute allows the NAS to send (in an Access-Request packet) the phone number that a call comes from. The phone number is identified using Automatic Number Identification (ANI) or similar technology. It is only used in Access-Request packets (RFC 2685).
32	NAS-Identifier	String	Network access server (NAS) identifier (the hostname of the device). It can be used instead of or in addition to NAS-IP-Address in the Access-Request packet that is sent to the authentication server. **NOTE** When the host name is not configured on the device, the NAS-Identifier attribute is not sent in the Access-Request packet (RFC 2865).
61	NAS-Port-Type	Integer	The physical port type that is authenticating the user. It can be used instead of or in addition to NAS-Port in the Access-Request packet that is sent to the authentication server. For network access control (NAC) authentication specify the NAS-Port-Type as "Ethernet". For device management authentication (when the user attempts to access the device by way of Telnet or SSH), specify the NAS-Port-Type as "Virtual". (RFC 2865).
64	Tunnel-Type	Integer	Tunnel protocol that is either in use by a tunnel terminator or to be used by a tunnel initiator (RFC 2868).
65	Tunnel-Medium-Type	Integer	Transport medium to use when creating a tunnel for protocols that can operate over multiple transport mediums (RFC 2868).
79	EAP-Message	String	An Extensible Authentication Protocol (EAP) packet (RFC 3579).
80	Message-Authenticator	String	A RADIUS Client receiving an Access-Accept, Access-Reject or Access-Challenge with a Message-Authenticator Attribute present must calculate the correct value of the Message-Authenticator and silently discard the packet if it does not match the value sent (RFC 2869).

TABLE 51 Supported RADIUS attributes for authentication (continued)

Attribute ID	Attribute name	Data Type	Description
81	Tunnel-Private-Group-ID	String	Group identifier for a particular tunneled session (RFC 2868).
95	NAS-IPv6-Address	Address	IPv6 address of the NAS that is requesting user authentication (RFC 3162).

The following table shows Microsoft vendor-specific RADIUS attributes that are supported on the Extreme device; these attributes are transmitted as sub-attributes within the Vendor-Specific attribute (26). For further information, refer to RFC 2548.

TABLE 52 Supported Microsoft vendor-specific RADIUS attributes for authentication

Attribute ID	Attribute Name	Data Type	Description
16	MS-MPPE-Send-Key	String	The MS-MPPE-Send-Key Attribute contains a session key for use by the Microsoft Point-to-Point Encryption Protocol (MPPE). This key is intended for encrypting packets sent from the NAS to the remote host. This attribute is only included in Access-Accept packets (RFC 2548).
17	MS-MPPE-Recv-Key	String	The MS-MPPE-Recv-Key Attribute contains a session key for use by the Microsoft Point-to-Point Encryption Protocol (MPPE). As the name implies, this key is intended for encrypting packets received by the NAS from the remote host. This attribute is only included in Access-Accept packets (RFC 2548).

The following RADIUS attributes are supported for multi-device port accounting.

TABLE 53 Supported RADIUS attributes for accounting

Attribute ID	Attribute Name	Data Type	Description
4	NAS-IP-Address	address	IP address of the network access server which is requesting user authentication (RFC 2865).
5	NAS-Port	integer	Physical port number of the NAS that is authenticating the user. It is only used in Access-Request packets (RFC 2865). **NOTE** The physical port number refers to a physical connection on the NAS rather than a TCP or UDP port number. Either NAS-Port, NAS-Port-Type (61),or both should be present in the Access-Request packet, when the NAS differentiates between its ports.
6	Service-Type	Integer	Indicates the type of service the user has requested, or the type of service to be provided. It may be used in both Access-Request and Access-Accept packets. A NAS is not required to implement all service types, but must treat unknown or unsupported Service-Types as though an Access-Reject packet has been received (RFC 2865).
12	Framed-MTU	Integer	Indicates the maximum transmission unit to be configured for the user, when it is not negotiated by some other means (such as PPP). It may be used in Access-Accept packets or in an Access-Request

TABLE 53 Supported RADIUS attributes for accounting (continued)

Attribute ID	Attribute Name	Data Type	Description
			packet as a hint by the NAS to the server that it would prefer that value; the server is not required to honor the hint (RFC 2865).
31	Calling-Station-Id	String	This attribute allows the NAS to send (in an Access-Request packet) the phone number that a call comes from. The phone number is identified using Automatic Number Identification (ANI) or similar technology. It is only used in Access-Request packets (RFC 2685).
40	Acct-Status-Type	integer	Indicates whether this Accounting-Request marks the beginning (Start) or the end (Stop) of the user service. It may be used by the client to mark the start of accounting (for example, upon booting) by specifying Accounting-On and to mark the end of accounting (for example, just before a scheduled reboot) by specifying Accounting-Off (RFC 2866).
41	Acct-Delay-Time	integer	Indicates how many seconds the client has been trying to send this record for, and can be subtracted from the time of arrival on the server to find the approximate time of the event generating this Accounting-Request (network transit time is ignored). RFC 2866.
44	Acct-Session-Id	String	Unique Accounting ID used to match start and stop records in a log file. The start and stop records for a given session must have the same Acct-Session-Id. An Accounting-Request packet must have an Acct-Session-Id. An Access-Request packet may have an Acct-Session-Id; if it does, then the NAS must use the same Acct-Session-Id in the Accounting-Request packets for that session (RFC 2866).
45	Acct-Authentic	Integer	This attribute may be included in an Accounting-Request to indicate how the user was authenticated (by RADIUS, the NAS itself, or another remote authentication protocol). Users who are delivered service without being authenticated should not generate Accounting records (RFC 2866).
46	Acct-Session-Time	Integer	Indicates the number of seconds the user has received service for, and can only be present in Accounting-Request records where the Acct-Status-Type is set to Stop (RFC 2866).
49	Acct-Terminate-Cause	Integer	Indicates how the session was terminated, and can only be present in Accounting-Request records where the Acct-Status-Type is set to Stop (RFC 2866).
61	NAS-Port-Type	integer	The physical port type that is authenticating the user. It can be used instead of or in addition to NAS-Port in the Access-Request packet that is sent to the authentication server. For network access control (NAC) authentication specify the NAS-Port-Type as "Ethernet". For device management authentication (when the user attempts to access the device by way of Telnet or SSH), specify the NAS-Port-Type as "Virtual". (RFC 2865).
95	NAS-IPv6-Address	address	IPv6 address of the NAS that is requesting user authentication (RFC 3162).

Dynamic VLAN and ACL assignments

The multi-device port authentication feature supports dynamic VLAN assignment , where a port can be placed in a VLAN based on the MAC address learned on that interface. When a MAC address is successfully authenticated, the RADIUS server sends the device a RADIUS Access-Accept message that allows the device to forward traffic from that MAC address. The RADIUS Access-Accept message can also contain attributes set for the MAC address in its access profile on the RADIUS server.

If one of the attributes in the Access-Accept message specifies a VLAN identifier, and this VLAN is available on the device, the port is moved from its default VLAN to the specified VLAN.

To enable dynamic VLAN assignment for authenticated MAC addresses, you must add the following attributes to the profile for the MAC address on the RADIUS server. Dynamic VLAN assignment on multi-device port authentication-enabled interfaces is enabled by default.

Attribute name	Type	Value
Tunnel-Type	064	13 (decimal) - VLAN
Tunnel-Medium-Type	065	6 (decimal) - 802
Tunnel-Private-Group-ID	081	*vlan-name* (string) - either the name or the number of a VLAN configured on the device.

In addition to dynamic VLAN assignment, Extreme devices also support dynamic ACL assignment as is the case with 802.1x port security.

Support for authenticating multiple MAC addresses on an interface

The multi-device port authentication feature allows multiple MAC addresses to be authenticated or denied authentication on each interface. The maximum number of MAC addresses that can be authenticated on each interface is 256. The default is 32.

Support for multi-device port authentication and 802.1x on the same interface

On the Extreme devices, multi-device port authentication and 802.1x security can be enabled on the same port. However, only one of them can authenticate a MAC address or 802.1x client. If an 802.1x client responds, the software assumes that the MAC should be authenticated using 802.1x protocol mechanisms and multi-device port authentication for that MAC is aborted. Also, at any given time, a port can have either 802.1x clients or multi-device port authentication clients but not both.

Configuring multi-device port authentication

Configuring multi-device port authentication on the Extreme devices consists of the following tasks:

- Enabling multi-device port authentication globally and on individual interfaces
- Configuring an Authentication Method List for 802.1x
- Setting RADIUS Parameters
- Specifying the format of the MAC addresses sent to the RADIUS server (optional)
- Specifying the authentication-failure action (optional)
- Defining MAC address filters (optional)
- Configuring dynamic VLAN assignment (optional)
- Specifying to which VLAN a port is moved after its RADIUS-specified VLAN assignment expires (optional)

- Saving dynamic VLAN assignments to the running configuration file (optional)
- Clearing authenticated MAC addresses (optional)
- Disabling aging for authenticated MAC addresses (optional)
- Specifying the aging time for blocked MAC addresses (optional)

Enabling multi-device port authentication

You globally enable multi-device port authentication on the router.

To globally enable multi-device port authentication on the device, enter the following command.

```
device(config)# mac-authentication enable
```

Syntax: [no] mac-authentication enable

Syntax: [no] mac-authentication enable slot/portnum | all

The all option enables the feature on all interfaces at once.

You can enable the feature on an interface at the interface CONFIG level.

Configuring an authentication method list for 802.1x

To use 802.1x port security, you must specify an authentication method to be used to authenticate Clients. The Extreme device supports RADIUS authentication with 802.1x port security. To use RADIUS authentication with 802.1x port security, you create an authentication method list for 802.1x and specify RADIUS as an authentication method, then configure communication between the device and the RADIUS server.

```
device(config)# aaa authentication dot1x default radius
```

Syntax: [no] aaa authentication dot1x default method-list

For the *method-list*, enter at least one of the following authentication methods:

radius - Use the list of all RADIUS servers that support 802.1x for authentication.

none - Use no authentication. The Client is automatically authenticated without the device using information supplied by the Client.

> NOTE
> If you specify both **radius** and **none**, make sure **radius** comes before **none** in the method list.

Setting RADIUS parameters

To use a RADIUS server to authenticate access to a device, you must identify the server to the device.

```
device(config)#
 radius-server host 10.157.22.99 auth-port 1812 acct-port 1813 default key mirabeau dot1x
```

Syntax: radius-server host ip-addr | server-name [auth-port number acct-port number [authentication-only | accounting-only | default [key 0 | 1 string [dot1x]]]]

The **host**ip-addr | *server-name* parameter is either an IP address or an ASCII text string.

The **auth-port**number parameter specifies what port to use for RADIUS authentication.

The **acct-port**number parameter specifies what port to use for RADIUS accounting.

The **dot1x** parameter indicates that this RADIUS server supports the 802.1x standard. A RADIUS server that supports the 802.1x standard can also be used to authenticate non-802.1x authentication requests.

> **NOTE**
> To implement 802.1x port security, at least one of the RADIUS servers identified to the device must support the 802.1x standard.

Specifying the format of the MAC addresses sent to the RADIUS server

When multi-device port authentication is configured, the device authenticates MAC addresses by sending username and password information to a RADIUS server. The username and password is the MAC address itself; that is, the device uses the MAC address for both the username and the password in the request sent to the RADIUS server.

By default, the MAC address is sent to the RADIUS server in the format *xxxxxxxxxxxx* . You can optionally configure the device to send the MAC address to the RADIUS server in the format *00-00-00-xx-xx-xx* , or the format *0000-00xx.xxxx* . To do this, enter a command such as the following.

```
device(config)# mac-authentication auth-passwd-format xxxx.xxxx.xxxx
```

Syntax: [no] mac-authentication auth-passwd-format 0000-00xx.xxxx | 00-00-00-xx-xx-xx | xxxxxxxxxxxx

Specifying the authentication-failure action

When RADIUS authentication for a MAC address fails, you can configure the device to perform one of two actions:
- Drop traffic from the MAC address in hardware (the default)
- Move the port on which the traffic was received to a restricted VLAN

To configure the device to move the port to a restricted VLAN when multi-device port authentication fails, enter commands such as the following.

```
device(config)# interface e 3/1
device(config-if-e100-3/1)# mac-authentication auth-fail-action restrict-vlan 100
```

Syntax: [no] mac-authentication auth-fail-action restrict-vlan [vlan-id]

If the ID for the restricted VLAN is not specified at the interface level, the global restricted VLAN ID applies for the interface.

To specify the VLAN ID of the restricted VLAN globally, enter the following command.

```
device(config)# mac-authentication auth-fail-vlan-id 200
```

Syntax: [no] mac-authentication auth-fail-vlan-id vlan-id

The command above applies globally to all MAC-authentication-enabled interfaces.

Note that the restricted VLAN must already exist on the device. You cannot configure the restricted VLAN to be a non-existent VLAN.

To configure the device to drop traffic from non-authenticated MAC addresses in hardware, enter commands such as the following.

```
device(config)# interface e 3/1
device(config-if-e100-3/1)# mac-authentication auth-fail-action block-traffic
```

Syntax: [no] mac-authentication auth-fail-action block-traffic

Dropping traffic from non-authenticated MAC addresses is the default behavior when multi-device port authentication is enabled.

Defining MAC address filters

You can specify MAC addresses that do not have to go through multi-device port authentication. These MAC addresses are considered pre-authenticated, and are not subject to RADIUS authentication. To do this, you can define MAC address filters that specify the MAC addresses to exclude from multi-device port authentication.

You should use a MAC address filter when the RADIUS server itself is connected to an interface where multi-device port authentication is enabled. If a MAC address filter is not defined for the MAC address of the RADIUS server and applied on the interface, the RADIUS authentication process would fail since the device would drop all packets from the RADIUS server itself.

For example, the following command defines a MAC address filter for address 0000.0058.aca4.

```
device(config)# mac-authentication mac-filter 1 permit 0000.0058.aca4
```

Syntax: [no] mac-authentication mac-filter filter

The following commands apply the MAC address filter on an interface so that address 0010.dc58.aca4 is excluded from multi-device port authentication.

```
device(config)# interface e 3/1
device(config-if-e100-3/1)# mac-authentication apply-mac-auth-filter 1
```

Syntax: [no] mac-authentication apply-mac-auth-filter filter-id

Configuring dynamic VLAN assignment

An interface can be dynamically assigned to a VLAN based on the MAC address learned on that interface. When a MAC address is successfully authenticated, the RADIUS server sends the device a RADIUS Access-Accept message that allows the device to forward traffic from that MAC address. The RADIUS Access-Accept message can also contain attributes set for the MAC address in its access profile on the RADIUS server.

If one of the attributes in the Access-Accept message specifies a VLAN identifier, and this VLAN is available on the device, the port is moved from its default VLAN to the specified VLAN.

To enable dynamic VLAN assignment for authenticated MAC addresses, you must add the following attributes to the profile for the MAC address on the RADIUS server (dynamic VLAN assignment on multi-device port authentication-enabled interfaces is enabled by default and can be disabled). Refer to Dynamic VLAN and ACL assignments on page 404 for a list of the attributes that must be set on the RADIUS server

Dynamic VLAN assignment on a multi-device port authentication-enabled interface is enabled by default. If it is disabled, enter commands such as the following command to enable it.

```
device(config)# interface e 3/1
device(config-if-e100-3/1)# mac-authentication enable-dynamic-vlan
```

Syntax: [no] mac-authentication enable-dynamic-vlan

If a previous authentication attempt for a MAC address failed, and as a result the port was placed in the restricted VLAN, but a subsequent authentication attempt was successful, the RADIUS Access-Accept message may specify a VLAN for the port. By default, the device moves the port out of the restricted VLAN and into the RADIUS-specified VLAN. You can optionally configure the device to ignore the RADIUS-specified VLAN in the RADIUS Access-Accept message, and leave the port in the restricted VLAN.

To do this, enter the following command.

```
device(config)#
mac-authentication no-override-restrict-vlan
```

Syntax: [no] mac-authentication no-override-restrict-vlan

NOTE

- For untagged ports, if the VLAN ID provided by the RADIUS server is valid, then the port is removed from its current VLAN and moved to the RADIUS-specified VLAN as an untagged port.
- If you configure dynamic VLAN assignment on a multi-device port authentication enabled interface, and the Access-Accept message returned by the RADIUS server does not contain a Tunnel-Private-Group-ID attribute, then it is considered an authentication failure, and the configured authentication failure action is performed for the MAC address.
- If the *vlan-name* string does not match either the name or the ID of a VLAN configured on the device, then it is considered an authentication failure, and the configured authentication failure action is performed for the MAC address.
- If an untagged port had previously been assigned to a VLAN though dynamic VLAN assignment, and then another MAC address is authenticated on the same port, but the RADIUS Access-Accept message for the second MAC address specifies a different VLAN, then it is considered an authentication failure for the second MAC address, and the configured authentication failure action is performed. Note that this applies only if the first MAC address has not yet aged out. If the first MAC address has aged out, then dynamic VLAN assignment would work as expected for the second MAC address.

Specifying the VLAN to which a port is moved after the RADIUS-specified VLAN assignment expires

When a port is dynamically assigned to a VLAN through the authentication of a MAC address, and the MAC session for that address is deleted on the device, then by default the port is removed from its RADIUS-assigned VLAN and placed back in the VLAN where it was originally assigned.

A port can be removed from its RADIUS-assigned VLAN when any of the following occur:

- The link goes down for the port
- The MAC session is manually deleted with the **mac-authentication clear-mac-session** command
- The MAC address that caused the port to be dynamically assigned to a VLAN ages out

For example, say port 1/1 is currently in VLAN 100, to which it was assigned when MAC address 0007.eaa1.e90f was authenticated by a RADIUS server. The port was originally configured to be in VLAN 111. If the MAC session for address 0007.eaa1.e90f is deleted, then port 1/1 is moved from VLAN 100 back into VLAN 111.

You can optionally specify an alternate VLAN to which to move the port when the MAC session for the address is deleted. For example, to place the port in the restricted VLAN, enter commands such as the following.

```
device(config)# interface e 3/1
device(config-if-e100-3/1)# mac-auth move-back-to-old-vlan port-restrict-vlan
```

Syntax: [no] mac-authentication move-back-to-old-vlan disable | port-configured-vlan | port-restrict-vlan | system-default-vlan

The **disable** keyword disables moving the port back to its original VLAN. The port would stay in its RADIUS-assigned VLAN.

The **port-configured-vlan** keyword removes the port from its RADIUS-assigned VLAN and places it back in the VLAN where it was originally assigned. This is the default.

The **port-restrict-vlan** keyword removes the port from its RADIUS-assigned VLAN and places it in the restricted VLAN.

The **system-default-vlan** keyword removes the port from its RADIUS-assigned VLAN and places it in the DEFAULT-VLAN.

Saving dynamic VLAN assignments to the running configuration file

You can configure the device to save the RADIUS-specified VLAN assignments to the device's running configuration file. To do this, enter the following command.

```
device(config)# mac-authentication save-dynamicvlan-to-config
```

Syntax: [no] mac-authentication save-dynamicvlan-to-config

By default, the dynamic VLAN assignments are not saved to the running configuration file. Entering the **show running-config** command does not display dynamic VLAN assignments, although they can be displayed with the **show vlan** and **show auth-mac-address detail** commands.

Clearing authenticated MAC addresses

The Extreme router maintains an internal table of the authenticated MAC addresses (viewable with the **show authenticated-mac-address** command). You can clear the contents of the authenticated MAC address table either entirely, or just for the entries learned on a specified interface. In addition, you can clear the MAC session for an address learned on a specific interface.

To clear the entire contents of the authenticated MAC address table, enter the following command.

```
device(config)# clear auth-mac-table
```

Syntax: clear auth-mac-table

To clear the authenticated MAC address table of entries learned on a specified interface, enter a command such as the following.

```
device(config)# clear auth-mac-table e 3/1
```

Syntax: clear auth-mac-table slot/portnum

To clear the MAC session for an address learned on a specific interface, enter commands such as the following.

```
device(config)# interface e 3/1
device(config-if-e100-3/1)# mac-authentication clear-mac-session 00e0.1234.abd4
```

Syntax: mac-authentication clear-mac-session mac-address

This command removes the Layer 2 CAM entry created for the specified MAC address. If the device receives traffic from the MAC address again, the MAC address is authenticated again.

Disabling aging for authenticated MAC addresses

MAC addresses that have been authenticated or denied by a RADIUS server are aged out if no traffic is received from the MAC address for a certain period of time:

- Authenticated MAC addresses or non-authenticated MAC addresses that have been placed in the restricted VLAN are aged out if no traffic is received from the MAC address over the device's normal MAC aging interval.
- Non-authenticated MAC addresses that are blocked by the device are aged out if no traffic is received from the address over a fixed hardware aging period (70 seconds), plus a configurable software aging period. (Refer to the next section for more information on configuring the software aging period).

You can optionally disable aging for MAC addresses subject to authentication, either for all MAC addresses or for those learned on a specified interface.

To disable aging for all MAC addresses subject to authentication on all interfaces where multi-device port authentication has been enabled, enter the following command.

```
device(config)# mac-authentication disable-aging
```

To disable aging for all MAC addresses subject to authentication on a specific interface where multi-device port authentication has been enabled, enter commands such as the following.

```
device(config)# interface e 3/1
device(config-if-e100-3/1)# mac-authentication disable-aging
```

Syntax: [no] mac-authentication disable-aging [denied-mac-only | permitted-mac-only]

denied-mac-only disables aging of denied sessions and enables aging of permitted sessions.

permitted-mac-only disables aging of permitted (authenticated and restricted) sessions and enables aging of denied sessions.

Specifying the aging time for blocked MAC addresses

When the device is configured to drop traffic from non-authenticated MAC addresses, traffic from the blocked MAC addresses is dropped in hardware, without being sent to the CPU. A Layer 2 CAM entry is created that drops traffic from the blocked MAC address in hardware. If no traffic is received from the blocked MAC address for a certain amount of time, this Layer 2 CAM entry is aged out. If traffic is subsequently received from the MAC address, then an attempt can be made to authenticate the MAC address again.

Aging of the Layer 2 CAM entry for a blocked MAC address occurs in two phases, known as hardware aging and software aging . The hardware aging period is fixed at 70 seconds and is non-configurable. The software aging time is configurable through the CLI.

Once the device stops receiving traffic from a blocked MAC address, the hardware aging begins and lasts for a fixed period of time. After the hardware aging period ends, the software aging period begins. The software aging period lasts for a configurable amount of time (by default 120 seconds). After the software aging period ends, the blocked MAC address ages out, and can be authenticated again if the device receives traffic from the MAC address.

To change the length of the software aging period for blocked MAC addresses, enter a command such as the following.

```
device(config)# mac-authentication max-age 180
```

Syntax: [no] mac-authentication max-age seconds

You can specify from 1 - 65535 seconds. The default is 120 seconds.

Displaying multi-device port authentication information

You can display the following information about the multi-device port authentication configuration:

- Information about authenticated MAC addresses
- Information about the multi-device port authentication configuration
- Authentication Information for a specific MAC address or port
- Multi-device port authentication settings and authenticated MAC addresses for each port where the multi-device port authentication feature is enabled
- The MAC addresses that have been successfully authenticated
- The MAC addresses for which authentication was not successful

Displaying authenticated MAC address information

To display information about authenticated MAC addresses on the ports where the multi-device port authentication feature is enabled, enter the following command.

```
device# show auth-mac-address
-------------------------------------------------------------------------
Port          Vlan  Accepted MACs    Rejected MACs    Attempted-MACs
-------------------------------------------------------------------------
1/18          100   1                100              0
1/20          40    0                0                0
1/22          100   0                0                0
4/5           30    0                0                0
```

Syntax: **show auth-mac-address**

The following table describes the information displayed by the **show auth-mac-address** command.

This field...	Displays...
Port	The port number where the multi-device port authentication feature is enabled.
Vlan	The VLAN to which the port has been assigned.
Accepted MACs	The number of MAC addresses that have been successfully authenticated
Rejected MACs	The number of MAC addresses for which authentication has failed.
Attempted-MACs	The rate at which authentication attempts are made for MAC addresses.

Displaying multi-device port authentication configuration information

To display a summary of multi-device port authentication that have been configured on the device, enter the following command.

```
device# show auth-mac configuration
Feature enabled                 : Yes
Global Fail-VLAN Id             : None
Username/Password format        : xxxx.xxxx.xxxx
Maximum Age                     : 120
Save dynamic VLAN configuration : No
Number of Ports enabled         : 25
-------------------------------------------------------------------------------
Port  Aging      Fail      Fail DynVLAN Override  Revert      MAC     DoS Protectn
                 Action    VLAN Support Restricted VLAN       Filter Enable Limit
-------------------------------------------------------------------------------
1/1   All        Blocked   N/A  Yes     Yes        Configured  No      No    512
1/2   Permitted  Blocked   101  No      Yes        Restricted  No      No    512
1/3   All        Blocked   N/A  Yes     Yes        Configured  No      No    512
1/4   Denied     Blocked   N/A  Yes     Yes        Configured  No      No    512
1/5   All        Blocked   N/A  Yes     Yes        Configured  No      No    512
1/6   None       Blocked   N/A  Yes     Yes        Sys.Default No      No    512
1/7   All        Blocked   N/A  Yes     Yes        Configured  No      No    512
1/8   All        Blocked   N/A  Yes     Yes        Configured  No      No    512
1/9   All        Blocked   N/A  Yes     Yes        Configured  No      No    512
1/10  All        Blocked   N/A  Yes     Yes        Configured  No      No    512
```

The following table describes the information displayed by the **show authenticated-mac-address configuration** command.

This field...	Displays...
Feature enabled	Whether the multi-device port authentication feature is enabled on the device.
Number of Ports enabled	The number of ports on which the multi-device port authentication feature is enabled.
Aging	Shows which MAC addresses are aged out.

This field...	Displays...
	Denied – Only denied MAC addresses are aged out
	Permitted – Only permitted MAC addresses are aged out
	All – Both denied and permitted MAC addresses are aged out
	None – None of the MAC addresses are aged out
Port	Information for each multi-device port authentication-enabled port.
Fail-Action	What happens to traffic from a MAC address for which RADIUS authentication has failed: either block the traffic or assign the MAC address to a restricted VLAN.
Fail VLAN	The restricted VLAN to which non-authenticated MAC addresses are assigned, if the Fail-Action is to assign the MAC address to a restricted VLAN.
DynVLAN Support	Whether RADIUS dynamic VLAN assignment is enabled for the port.
Override Restricted	Whether or not a port in a restricted VLAN (due to a failed authentication) is removed from the restricted VLAN on a subsequent successful authentication on the port.
Revert VLAN	The VLAN that the port reverts to when the RADIUS-assigned dynamic VLAN expires.
MAC-filter	Whether a MAC filter has been applied to this port to specify pre-authenticated MAC addresses.
DOS Enable	Denial of Service status. This column will always show "No" since DOS is not supported.
Protection Limit	This is not applicable to the device, but the output always show "512".

Syntax: show auth-mac-address configuration

To display detailed information about the multi-device port authentication configuration and authenticated MAC addresses for a port where the feature is enabled, enter the following command.

```
device# show auth-mac-address detail
Port 1/18
 Dynamic-Vlan Assignment        : Enabled
 RADIUS failure action          : Block Traffic
 Override-restrict-vlan         : Yes
 Port VLAN                      : 4090 (Configured)
 DOS attack protection          : Disabled
 Accepted Mac Addresses         : 0
 Rejected Mac Addresses         : 0
 Aging of MAC-sessions          : Enable-All
 Port move-back vlan            : Port-Configured
 MAC Filter applied             : No
                              1 : 0000.0010.2000
 MAC TABLE
 -------------------------------------------
 MAC Address     Port    VLAN Access     Age
 -------------------------------------------
 00A1.0010.2000 1/18    1    Allowed    0
 00A1.0010.2001 1/18    1    Blocked    120
 00A1.0010.2002 1/18    1    Init       0
```

The following table describes the information displayed by the **show authenticated-mac-address** command.

This field...	Displays...
Port	The port to which this information applies.
Dynamic-Vlan Assignment	Whether RADIUS dynamic VLAN assignment has been enabled for the port.

This field...	Displays...
RADIUS failure action	What happens to traffic from a MAC address for which RADIUS authentication has failed: either block the traffic or assign the MAC address to a restricted VLAN.
Override-restrict-vlan	Whether a port can be dynamically assigned to a VLAN specified by a RADIUS server, if the port had been previously placed in the restricted VLAN because a previous attempt at authenticating a MAC address on that port failed.
Port VLAN	The VLAN to which the port is assigned, and whether the port had been dynamically assigned to the VLAN by a RADIUS server.
DOS attack protection	Whether denial of service attack protection has been enabled for multi-device port authentication, limiting the rate of authentication attempts sent to the RADIUS server.
Accepted MAC Addresses	The number of MAC addresses that have been successfully authenticated.
Rejected MAC Addresses	The number of MAC addresses for which authentication has failed.
Aging of MAC-sessions	Whether software aging of MAC addresses is enabled.
Max-Age of MAC-sessions	The configured software aging period.
Port move-back VLAN	The VLAN that the port reverts to when the RADIUS-assigned dynamic VLAN expires.
MAC Filter applied	Whether a MAC filter has been applied to this port to specify pre-authenticated MAC addresses.
MAC Table	The MAC addresses learned on the port.

Syntax: **show auth-mac-address detail**

Displaying multi-device port authentication information for a specific MAC address or port

To display authentication information for a specific MAC address or port, enter a command such as the following.

```
device# show auth-mac-address 0007.e90f.eaa1
---------------------------------------------------------------------------
MAC/IP Address      Port      Vlan     Access      Age
---------------------------------------------------------------------------
00A1.0010.2000      1/18      1        Allowed     0
```

Syntax: **show auth-mac-address mac-address | ip-address | slot/portnum**

The *ip-address* parameter lists the MAC address associated with the specified IP address.

The *slot/portnum* parameter lists the MAC addresses on the specified port.

The following table describes the information displayed by the **show auth-mac-address** command for a specified MAC address or port.

This field...	Displays...
MAC or IP Address	The MAC address for which information is displayed. If the packet for which multi-device port authentication was performed also contained an IP address, then the IP address is displayed as well.
Port	The port on which the MAC address was learned.
VLAN	The VLAN to which the MAC address was assigned.
Access	Whether or not the MAC address was allowed or denied access into the network.
Age	The age of the MAC address entry in the authenticated MAC address list.

Displaying the authenticated MAC addresses

To display the MAC addresses that have been successfully authenticated, enter the following command.

```
device# show auth-mac-addresses authorized-mac
MAC TABLE
-------------------------------------------
MAC Address     Port     VLAN Access       Age
-------------------------------------------
00A1.0010.2000 1/18     1     Allowed        0
00A1.0010.2001 1/18     1     Allowed      120
00A1.0010.2002 1/18     1     Allowed        0
```

Syntax: **show auth-mac-addresses authorized-mac**

Displaying the non-authenticated MAC addresses

To display the MAC addresses for which authentication was not successful, enter the following command.

```
device# show auth-mac-addresses unauthorized-mac
MAC TABLE
-------------------------------------------
MAC Address     Port     VLAN Access       Age
-------------------------------------------
00A1.0010.2000 1/18     1     Blocked        0
00A1.0010.2001 1/18     1     Blocked      120
00A1.0010.2002 1/18     1     Blocked        0
```

Syntax: **show auth-mac-addresses unauthorized-mac**

Secure Shell

Secure Shell (SSH) server is a mechanism for allowing secure remote access to management functions on a device. The SSH server provides a function similar to Telnet. Users can log into and configure the device using a publicly or commercially available SSH client program, just as they can with Telnet. However, unlike Telnet, which provides no security, SSH server provides a secure, encrypted connection to the device.

SSHv2 is supported on an Extreme device. The SSHv2 implementation is compatible with all versions of the SSHv2 protocol. At the beginning of an SSH server session, the device negotiates the version of SSHv2 to be used. The highest version of SSHv2 supported by both the device and the client is the version that is used for the session. Once the SSHv2 Version is negotiated, the host key algorithm with highest security ranking is negotiated and then the MAC, Encryption Algorithms are negotiated.

The maximum of 16 in-bound SSH server sessions are allowed. One out-bound SSH client session can be established from the device. The outbound session ID is always 17.

Also, the device supports Secure Copy (SCP) for securely transferring files between an Extreme device and an SCP-enabled remote host. Refer to the Using Secure Copy section for more information.

> **NOTE**
> SSH server and SSH client functionality are disabled by default. To gain access to a device through SSH server, you must enable it as described in this chapter.

SSH server version 2 support

SSHv2 is a substantial revision of Secure Shell, comprising the following hybrid protocols and definitions:

*   SSH server Transport Layer Protocol
*   SSH server Authentication Protocol
*   SSH server Connection Protocol
*   SECSH Public Key File Format
*   SSH server Fingerprint Format
*   SSH server Protocol Assigned Numbers
*   SSH server Transport Layer Encryption Modes
*   SCP or SFTP or SSH server URI Format

If you are using redundant management modules, you can synchronize the DSA host key pair and RSA Host key pair between the active and standby modules by entering the **sync-standby** command at the Privileged EXEC level of the CLI. By default these keys are synced to standby. The user can do force sync using the **sync-standby** command.

Supported SSHv2 clients

*   The following SSH clients have been tested with SSHv2 in NetIron 6.0.00a release for diffie-hellman-group14-sha1 key exchange algorithm support:
    –   Open SSH client 6.6.1p1, 5.3p1, and 7.1p1

- – NetIron SSH client

    NOTE
    Supported SSH client public key sizes are 1024 bits for DSA keys, and 1024 or 2048 bits for RSA keys.

  – PuTTY version 0.63

    NOTE
    On the PuTTy client, under the options that control key re-exchange, it is recommended that the maximum minutes before rekey be set to 0 and the maximum data before rekey be set to 0.

- The following SSH clients have been tested with SSHv2 (but not tested for diffie-hellman-group14-sha1 key exchange algorithm support in NetIron 6.0.00a release):

  – SSH server Secure Shell 3.2.3
  – Van Dyke SecureCRT 4.0, 4.1, 5.1, 5.5, 6.1, and 6.5.2
  – F-Secure SSH Client 5.3, 6.0, 6.1, and 6.2 beta
  – PuTTY 0.54 and 0.56

    NOTE
    On the PuTTy client, under the options that control key re-exchange, it is recommended that the maximum minutes before rekey be set to 0 and the maximum data before rekey be set to 0.

  – Open SSH server 3.5_p1, 3.6.1p2, 4.3p1, 5.3p1, 5.8p1, and 5.9p2
  – Solaris Sun-SSH-1.0, version 2.4

Supported features

SSHv2 provides an SSH server and an SSH client. The SSH server allows secure remote access management functions on a device.

SSHv2 support includes the following features:

- The following supported encryption cipher algorithms are listed in order of preference:

  – aes256-cbc: AES in CBC mode with 256-bit key
  – aes192-cbc: AES in CBC mode with 192-bit key
  – aes128-cbc: AES in CBC mode with 128-bit key

    NOTE
    CBC mode can be disabled using the **ip ssh encryption disable-aes-cbc** command.

  – 3des-cbc: Triple-DES

    NOTE
    3-DES can be disabled using the **ip ssh encryption aes-only** command.

  – aes256-ctr: AES in CTR mode with 256-bit key
  – aes192-ctr: AES in CTR mode with 192-bit key
  – aes128-ctr: AES in CTR mode with 128-bit key

- The following key exchange methods are listed in order of preference:

  – **diffie-hellman-group14-sha1**
  – **diffie-hellman-group1-sha1**

Beginning with NetIron 6.0.00a, the diffie-hellman-group14-sha1 key exchange algorithm can be used as the key exchange method to establish an SSH connection in addition to the currently supported diffie-hellman-group1-sha1 algorithm. Both key

exchange algorithms are applicable to client and server SSH roles on the device. The diffie-hellman-group14-sha1 key exchange algorithm is given higher priority than diffie-hellman-group1-sha1 because this algorithm provides enhanced encryption of shared secrets between two devices. When both options are offered, the clients or servers that support both key exchange methods use diffie-hellman-group14-sha1. However, you can configure the client to use either of the groups by providing respective key exchange algorithm on the command line for an inbound connection to the device from any Linux machine, as shown in the following examples:

```
ssh -o KexAlgorithms=diffie-hellman-group1-sha1 lab@10.24.12.69

ssh -o KexAlgorithms=diffie-hellman-group14-sha1 lab@10.24.12.69
```

For the outbound SSH session, diffie-hellman-group14-sha1 key exchange algorithm is used by default.

> **NOTE**
> High CPU usage is expected while establishing SSH sessions with the diffie-hellman-group14-sha1 than with using the diffie-hellman-group1-sha1 key exchange algorithm.

- Public key algorithm **ssh-dss**
- Public key algorithm **ssh-rsa**
- **hmac-sha1** algorithm that ensures data integrity.
- The **password** and **publickey** authentication methods.
- Sixteen inbound SSH server connections at one time
- Upto 16 inbound SSH clients
- One outbound SSH client
- SCP supports AES encryption.

Unsupported features

The following features are not supported:

- Compression
- TCP or IP port forwarding, X11 forwarding, and secure file transfer
- SSH server version 1

Configuring SSH server

The implementation of SSH server supports the following types of user authentication:

- DSA challenge-response authentication and RSA challenge-response authentication, where a collection of public keys is stored on the device. Only clients with a private key that corresponds to one of the stored public keys can gain access to the device using SSH server.
- Password authentication , where users attempting to gain access to the device using an SSH client are authenticated with passwords stored on the device or on a TACACS, TACACS+, or RADIUS server.

User authentication is enabled by default. You can configure the device to use any number of them.

To configure SSH on a device, perform the following tasks.

1. Generate a host DSA or RSA public and private key pair for the device.

   Refer to Table 54.

2.   Configure DSA or RSA challenge-response authentication.

       Refer to Configuring DSA public key authentication on page 424.

3. Set optional parameters.

    Refer to Setting optional parameters on page 425

    You can also view information about active SSH server connections on the device, as well as terminate the connections.

    To display SSH server configuration information, use the following command:

```
device# show ip ssh config
SSH server              : Enabled
SSH port                : tcp\22
Host Key                : DSA 1024
Encryption              : aes256-cbc, aes192-cbc, aes128-cbc, aes256-ctr, aes192-ctr, aes128-ctr,
3des-cbc
Permit empty password   : Yes
Authentication methods  : Password, Public-key, Interactive
Authentication retries  : 3
Login timeout (seconds) : 120
Idle timeout (minutes)  : 0
Strict management VRF   : Disabled
SCP                     : Enabled
SSH IPv4 clients        : All
SSH IPv6 clients        : All
SSH IPv4 access-group   :
SSH IPv6 access-group   :
SSH Client Keys         :
device#
```

Syntax: show ip ssh config

Table 54 shows the output information for the **show ip ssh config** command.

TABLE 54 show ip ssh config command output information.

Field	Description
SSH server	SSH server is enabled or disabled
SSH port	SSH port number

Encryption	The encryption used for the SSH connection. The following values are displayed when the Standard mode is enabled:
	
	• aes256-ctr, aes192-ctr, aes128-ctr, aes256-cbc, aes192-cbc, aes128-cbc, 3des-cbc indicate the different AES and CTR (counter mode) methods used for encryption.
	
	• 3-DES indicates 3-DES algorithm is used for encryption.
	
	    NOTE
	    CBC mode can be disabled using the **ip ssh encryption disable-aes-cbc** command.
	
	3-DES can be disabled using the **ip ssh encryption aes-only** command. The following values are displayed when the Standard mode with **ip ssh encryption aes-only** command is enabled:
	
	• aes256-ctr, aes192-ctr, aes128-ctr, aes256-cbc, aes192-cbc, aes128-cbc.
	
	The following values are displayed when the JITC mode is enabled using the **jitc enable** command: In this mode, the AES-CTR encryption mode is enabled, and the AES-CBC encryption for SSH is disabled.
	
	• aes256-ctr, aes192-ctr, aes128-ctr

Permit empty password	Empty password login is allowed or not allowed.
Authentication methods	The authentication methods used for SSH. The authentication can have one or more of the following values:

TABLE 54 show ip ssh config command output information. (continued)

Field	Description
	• Password - indicates that you are prompted for a password when attempting to log into the device. • Public-key - indicates that DSA or RSA challenge-response authentication is enabled. • Interactive - indicates the interactive authentication si enabled.
Authentication retries	The number of authentication retries. This number can be from 1 to 5.
Login timeout (seconds)	SSH login timeout value in seconds. This can be from 0 to 120.
Idle timeout (minutes)	SSH idle timeout value in minutes. This can be from 0 to 240.
Strict management VRF	Strict management VRF is enabled or disabled.
SCP	SCP is enabled or disabled.
SSH IPv4 clients	The list of IPv4 addresses to which SSH access is allowed. The default is "All".
SSH IPv6 clients	The list of IPv4 addresses to which SSh access is allowed. Default "All".
SSH IPv4 access-list	The IPv4 ACL used to permit or deny access using SSH.
SSH IPv6 access-list	The IPv6 ACL used to permit or deny access to device using SSH.

Generating a host key pair

When SSH server is configured, a public and private host DSA key pair is generated for the device. The SSH server on the device uses this host DSA key pair, along with a dynamically generated server DSA key pair , to negotiate a session key and encryption method with the client trying to connect to it.

The host DSA key pair is stored in the device's system-config file. Only the public key is readable. The public key should be added to a "known hosts" file (for example, $HOME/.ssh/known_hosts on OpenSSH Linux & UNIX systems) on the clients who want to access the device. Some SSH client programs add the public key to the known hosts file automatically; in other cases, you must manually create a known hosts file and place the device's public key in it. Refer to for an example of what to place in the known hosts file.

> **NOTE**
> This describes the OpenSSH (Linux) SSH client and server. Others are not the same procedure.

While the SSH server listener exists at all times, sessions can not be started from clients until a key is generated. Once a key is generated, clients can start sessions. The keys are not displayed in the configuration file by default. The default DSA is used when the DSA or RSA keyword not specified. To display the keys, use the **ssh show-host-keys** command in the Privileged EXEC mode. To generate a public and private DSA host key pair on a device, enter the following commands.

```
device(config)# crypto key generate
```

When a host key pair is generated, it is saved to the flash memory of all management modules.

To disable SSH server in SSHv2 on a device, enter the following commands.

```
device(config)# crypto key zeroize
```

When SSH server is disabled, it is deleted from the flash memory of all management modules.

> **NOTE**
> This command without the DSA or RSA keyword will delete both encryption key pairs (RSA and DSA).

Syntax: **crypto key generate | zeroize { dsa | rsa }**

The **generate** keyword places a host key pair in the flash memory and enables the SSH server on the device, if it is not already enabled.

The **zeroize** keyword deletes the host key pair from the flash memory and disables the SSH server if no other server host keys exist on the device.

The **dsa** keyword specifies a DSA host key pair. This keyword is optional. If you do not enter it, the **crypto key generate** command generates a DSA key pair by default, and the **crypto key zeroize** command works.

By default, public keys are hidden in the running configuration. You can optionally configure the device to display the DSA host key pair in the running configuration file entering the following command.

```
device# ssh show-host-keys
```

Syntax: **ssh show-host-keys**

To hide the public keys in the running configuration file, enter the following command.

```
device# ssh no-show-host-keys
```

Syntax: **ssh no-show-host-keys**

Enabling and disabling SSH server by generating and deleting host keys

To enable SSH server, you must generate a public and private DSA or RSA host key pair on the device. The SSH server on the ServerIron uses this host DSA or RSA key pair, along with a dynamically generated server DSA or RSA key pair, to negotiate a session key and encryption method with the client trying to connect to it.

While the SSH server listener exists at all times, sessions can not be started from the client until a host key is generated. After a host key is generated, clients can start sessions.

To disable SSH server, you delete all of the host keys from the device.

When a host key pair is generated, it is saved to the flash memory of all management modules. When a host key pair is deleted, it is deleted from the flash memory of all management modules.

The time range to initially generate SSH server keys varies. Refer to the section for initial SSH server key generation time ranges

Generating and deleting a DSA key pair

To generate a DSA key pair, enter the following command.

```
device(config)#crypto key generate dsa
```

To delete the DSA host key pair, enter the following command.

```
device(config)#crypto key zeroize dsa
```

Syntax: **crypto key generate | zeroize** *dsa*

The **generate** keyword places a host key pair in the flash memory and enables SSH server on the device, if it is not already enabled.

The **zeroize** keyword deletes the host key pair from the flash memory. This disables SSH server if no other server host keys exist on the device.

The **dsa** keyword specifies a DSA host key pair. This keyword is optional. If you do not enter it, the command **crypto key generate** generates a DSA key pair by default.

Generating and deleting an RSA key pair

To generate an RSA key pair, enter a command such as the following:

```
device(config)#crypto key generate rsa modulus 2048
```

To delete the RSA host key pair, enter the following command.

```
device(config)#crypto key zeroize rsa
```

Syntax: crypto key generate | zeroize rsa [modulus *modulus-size* **]**

The **generate** keyword places an RSA host key pair in the flash memory and enables SSH server on the device, if it is not already enabled.

The optional [**modulus**] parameter specifies the modulus size of the RSA key pair, in bits. The valid values for *modulus-size* are 1024 or 2048. The default value is 2048.

The **zeroize** keyword deletes the RSA host key pair from the flash memory. This disables SSH if no other authentication keys exist on the device.

The **rsa** keyword specifies an RSA host key pair.

Deleting DSA and RSA key pairs

To delete DSA and RSA key pairs from the flash memory, enter the following command:

```
device(config)#crypto key zeroize
```

Syntax: crypto key zeroize

The **zeroize** keyword deletes the host key pair from the flash memory. This disables SSH server.

Providing the public key to clients

The host DSA or RSA key pair is stored in the system-config file of the Extreme device. Only the public key is readable. Some SSH client programs add the public key to the known hosts file automatically; in other cases, you must manually create a known hosts file and place the public key of the Extreme device in it.

If you are using SSH server to connect to an Extreme device from a Linux or OpenSSH system, you may need to add the public key on the Extreme device to a "known hosts" file on the client OpenSSH system; for example, $HOME/.ssh/known_hosts. The following is an example of an entry in a known hosts file.

```
AAAAB3NzaC1kc3MAAACBAPY8ZOHY2yFSJA6XYC9HRwNHxaehvx5wOJ0rzZdzoSOXxbET W6ToHv8D1UJ/
z+zHo9Fiko5XybZnDIaBDHtblQ+Yp7StxyltHnXF1YLfKD1G4T6JYrdH YI14Om
1eg9e4NnCRleaqoZPF3UGfZia6bXrGTQf3gJq2e7Yisk/gF+1VAAAAFQDb8D5cv
wHWTZDPfX0D2s9Rd7NBvQAAAIEAlN92+Bb7D4KLYk3IwRbXblwXdkPggA4pfdtW9v
GfJ0/RHd+NjB4eo1D+0dix6tXwYGN7PKS5R/FXPNwxHPapcj9uL1Jn2AWQ2dsknf+i/FAA
vioUPkmdMc0zuWoSOEsSNhVDtX3WdvVcGcBq9cetzrtOKWOocJmJ80qadxTRHtUAAACB
AN7CY+KKv1gHpRzFwdQm7HK9bb1LAo2KwaoXnadFgeptNBQeSXG1vO+JsvphVMBJc9HS
n24VYtYtsMu74qXviYjziVucWKjjKEb11juqnF0GDlB3VVmxHLmxnAz643WK42Z7dLM5
sY29ouezv4Xz2PuMch5VGPP+CDqzCM4loWgV
```

Configuring DSA or RSA public key authentication

With DSA or RSA public key authentication, a collection of clients' public keys are stored on the Extreme device. Clients are authenticated using these stored public keys. Only clients that have a private key that corresponds to one of the stored public keys can gain access to the device using SSH server.

Setting up DSA or RSA private key authentication consists of the following steps.

1. Import authorized public keys into the Extreme device.
2. Enable DSA or RSA public key authentication.

Importing authorized public keys into the device

SSH clients that support DSA or RSA authentication normally provide a utility to generate a DSA or RSA key pair. The private key is usually stored in a password-protected file on the local host; the public key is stored in another file and is not protected. You must import the client public key for each client into the device.

Collect one public key of each key type (DSA and/or RSA) from each client to be granted access to the Extreme device and place all of these keys into one file. This public key file may contain up to 32 keys. The following is an example of a public key file containing one public key:

```
---- BEGIN SSH2 PUBLIC KEY ----
Comment: DSA Public Key  AAAAB3NzaC1kc3MAAACBAPY8ZOHY2yFSJA6XYC9HRwNHxaehvx5wOJ0rzZdzoSOXxbET W6ToHv8D1UJ/
z+zHo9Fiko5XybZnDIaBDHtblQ+Yp7StxyltHnXF1YLfKD1G4T6JYrdH YI14Om
1eg9e4NnCRleaqoZPF3UGfZia6bXrGTQf3gJq2e7Yisk/gF+1VAAAAFQDb8D5cv
wHWTZDPfX0D2s9Rd7NBvQAAAIEAlN92+Bb7D4KLYk3IwRbXblwXdkPggA4pfdtW9v
GfJ0/RHd+NjB4eo1D+0dix6tXwYGN7PKS5R/FXPNwxHPapcj9uL1Jn2AWQ2dsknf+i/FAA
vioUPkmdMc0zuWoSOEsSNhVDtX3WdvVcGcBq9cetzrtOKWOocJmJ80qadxTRHtUAAACB
AN7CY+KKv1gHpRzFwdQm7HK9bb1LAo2KwaoXnadFgeptNBQeSXG1vO+JsvphVMBJc9HS
n24VYtYtsMu74qXviYjziVucWKjjKEb11juqnF0GDlB3VVmxHLmxnAz643WK42Z7dLM5
sY29ouezv4Xz2PuMch5VGPP+CDqzCM4loWgV
---- END SSH2 PUBLIC KEY ----
```

> **NOTE**
> Each key in the public key file must begin and end with the first and last lines in this example. If your client does not include these lines in the public key, you must manually add them.

SSH server key generation time

The time range to initially generate SSH server keys varies. Refer to for initial SSH server key generation time ranges.

TABLE 55 Initial SSH server key generation time ranges (measured in seconds)

Device	Low	High	Average
MLX Series and XMR Series devices	8 seconds	141 seconds	44 seconds
CES 2000 Series and CER 2000 Series devices	4 seconds	77 seconds	25.2 seconds

Providing the public key to clients

If you are using SSH server to connect to a device from a Linux or OpenSSH system, you may need to add the device's public key to a "known hosts" file; for example, $HOME/.ssh/known_hosts. The following is an example of an entry in a known hosts file.

```
AAAAB3NzaC1kc3MAAACBAPY8ZOHY2yFSJA6XYC9HRwNHxaehvx5wOJ0rzZdzoSOXxbET
W6ToHv8D1UJ/z+zHo9Fiko5XybZnDIaBDHtblQ+Yp7StxyltHnXF1YLfKD1G4T6JYrdH YI14Om
1eg9e4NnCRleaqoZPF3UGfZia6bXrGTQf3gJq2e7Yisk/gF+1VAAAAFQDb8D5cv
wHWTZDPfX0D2s9Rd7NBvQAAAIEAlN92+Bb7D4KLYk3IwRbXblwXdkPggA4pfdtW9v
GfJ0/RHd+NjB4eo1D+0dix6tXwYGN7PKS5R/FXPNwxHPapcj9uL1Jn2AWQ2dsknf+i/FAA
vioUPkmdMc0zuWoSOEsSNhVDtX3WdvVcGcBq9cetzrtOKWOocJmJ80qadxTRHtUAAACB
AN7CY+KKv1gHpRzFwdQm7HK9bb1LAo2KwaoXnadFgeptNBQeSXG1vO+JsvphVMBJc9HS
n24VYtYtsMu74qXviYjziVucWKjjKEb11juqnF0GDlB3VVmxHLmxnAz643WK42Z7dLM5
sY29ouezv4Xz2PuMch5VGPP+CDqzCM4loWgV
```

Configuring DSA public key authentication

With DSA public key authentication, a collection of clients' public keys are stored on the device. Clients are authenticated using these stored public keys. Only clients that have a private key that corresponds to one of the stored public keys can gain access to the device using SSH server.

When DSA challenge-response authentication is enabled, the following events occur when a client attempts to gain access to the device using SSH server.

1. The client sends its public key to the device.
2. The device compares the client's public key to those stored in memory.
3. If there is a match, the device uses the public key to encrypt a random sequence of bytes.
4. The device sends these encrypted bytes to the client.
5. The client uses its private key to decrypt the bytes.
6. The client sends the decrypted bytes back to the device.
7. The device compares the decrypted bytes to the original bytes it sent to the client. If the two sets of bytes match, it means that the client's private key corresponds to an authorized public key, and the client is authenticated.

   Setting up DSA public key authentication consists of the following steps:

   1. Importing authorized public keys into the device.
   2. Enabling DSA public key authentication

Importing authorized public keys into the device

SSH clients that support DSA authentication normally provide a utility to generate an DSA key pair. The private key is usually stored in a password-protected file on the local host; the public key is stored in another file and is not protected. You should collect one public key from each client to be granted access to the device and place all of these keys into one file. This public key file is imported into the device.

The following is an example of a public key file containing one public keys.

```
---- BEGIN SSH2 PUBLIC KEY ----
Comment: DSA Public Key  AAAAB3NzaC1kc3MAAACBAPY8ZOHY2yFSJA6XYC9HRwNHxaehvx5wOJ0rzZdzoSOXxbET
W6ToHv8D1UJ/z+zHo9Fiko5XybZnDIaBDHtblQ+Yp7StxyltHnXF1YLfKD1G4T6JYrdH YI14Om
1eg9e4NnCRleaqoZPF3UGfZia6bXrGTQf3gJq2e7Yisk/gF+1VAAAAFQDb8D5cv
wHWTZDPfX0D2s9Rd7NBvQAAAIEAlN92+Bb7D4KLYk3IwRbXblwXdkPggA4pfdtW9v
GfJ0/RHd+NjB4eo1D+0dix6tXwYGN7PKS5R/FXPNwxHPapcj9uL1Jn2AWQ2dsknf+i/FAA
vioUPkmdMc0zuWoSOEsSNhVDtX3WdvVcGcBq9cetzrtOKWOocJmJ80qadxTRHtUAAACB
AN7CY+KKv1gHpRzFwdQm7HK9bb1LAo2KwaoXnadFgeptNBQeSXG1vO+JsvphVMBJc9HS
n24VYtYtsMu74qXviYjziVucWKjjKEb11juqnF0GDlB3VVmxHLmxnAz643WK42Z7dLM5
sY29ouezv4Xz2PuMch5VGPP+CDqzCM4loWgV
---- END SSH2 PUBLIC KEY ----
```

> **NOTE**
> Make sure the key ends with the complete phrase `"---- END SSH2 PUBLIC KEY ----"` before importing the public key. Otherwise, a warning is displayed whenever the device is reloaded.

You can import the authorized public keys into the active configuration by loading them from a file on a TFTP server and are saved on the EEPROM of the chassis.

> **NOTE**
> When one public-key file already exists, downloading a second public-key file will cause the second public-key file to overwrite the existing one. Downloading a public-key file when a public-key file already exists also erases currently loaded public-keys in the active configuration and loads only keys in the newly downloaded file.

To cause a public key file called pkeys.txt to be loaded from a TFTP server each time the device is booted, enter a command such as the following.

```
device(config)# ip ssh pub-key-file tftp 192.168.1.234 pkeys.txt
```

Syntax: ip ssh pub-key-file tftp ipv6 ipv6-addr | tftp-server-ip-addr filename [remove]

The *tftp-server-ip-addr* variable is the IP address of the tftp server that contains the public key file that you want to import into the device.

The *filename* variable is the name of the dsa public key file that you want to import into the device.

The **remove** parameter deletes the key from the system.

To display the currently loaded public keys, enter the following command.

```
device# show ip client-pub-key
---- BEGIN SSH2 PUBLIC KEY ----
Comment: DSA Public Key  AAAAB3NzaC1kc3MAAACBAPY8ZOHY2yFSJA6XYC9HRwNHxaehvx5wOJ0rzZdzoSOXxbET
W6ToHv8D1UJ/z+zHo9Fiko5XybZnDIaBDHtblQ+Yp7StxyltHnXF1YLfKD1G4T6JYrdH YI14Om
1eg9e4NnCRleaqoZPF3UGfZia6bXrGTQf3gJq2e7Yisk/gF+1VAAAAFQDb8D5cv
wHWTZDPfX0D2s9Rd7NBvQAAAIEAlN92+Bb7D4KLYk3IwRbXblwXdkPggA4pfdtW9v
GfJ0/RHd+NjB4eo1D+0dix6tXwYGN7PKS5R/FXPNwxHPapcj9uL1Jn2AWQ2dsknf+i/FAA
vioUPkmdMc0zuWoSOEsSNhVDtX3WdvVcGcBq9cetzrtOKWOocJmJ80qadxTRHtUAAACB
AN7CY+KKv1gHpRzFwdQm7HK9bb1LAo2KwaoXnadFgeptNBQeSXG1vO+JsvphVMBJc9HS
n24VYtYtsMu74qXviYjziVucWKjjKEb11juqnF0GDlB3VVmxHLmxnAz643WK42Z7dLM5
sY29ouezv4Xz2PuMch5VGPP+CDqzCM4loWgV
---- END SSH2 PUBLIC KEY ----
```

Syntax: show ip client-pub-key [| begin expression | exclude expression | include expression]

To clear the public keys from the buffers, enter the following command.

```
device# clear public-key
```

Syntax: clear public-key

Use the **ip ssh pub-key remove** command to delete the public key from the system.

Enabling DSA public key authentication

DSA public key authentication is enabled by default. You can disable or re-enable it manually.

To enable DSA public key authentication.

```
device(config)# ip ssh key-authentication yes
```

To disable DSA public key authentication.

```
device(config)# ip ssh key-authentication no
```

Syntax: ip ssh key-authentication yes | no

Setting optional parameters

You can adjust the following SSH server settings on the device:

- Number of SSH server authentication retries
- User authentication method the device uses for SSH server connections
- Whether or not the device allows users to log in without supplying a password
- Port number for SSH server connections
- SSH server login timeout value

- A specific interface to be used as the source for all SSH server traffic from the device
- Maximum idle time for SSH server sessions
- Disable 3-DES support

Setting the number of SSH server authentication retries

By default, the device attempts to negotiate a connection with the connecting host three times. The number of authentication retries can be changed to between 1 - 5.

For example, the following command changes the number of authentication retries to 5.

```
device(config)# ip ssh authentication-retries 5
```

Syntax: **ip ssh authentication-retries number**

> NOTE
> The **ip ssh authentication-retries** command is not applicable on an Extreme device which acts as an SSH client. When attempting to establish an SSH connection with wrong user credentials on an Extreme device acting as SSH client the session is not established and it is terminated, as the device does not check for SSH authentication retry configuration set using **ip ssh authentication-retries** command. The **ip ssh authentication-retries** command is applicable only to SSH clients like PUTTY, Secure CRT, and so on.

Deactivating user authentication

After the SSH server on the device negotiates a session key and encryption method with the connecting client, user authentication takes place. The implementation of SSH server supports DSA challenge-response authentication and password authentication.

With DSA challenge-response authentication, a collection of clients' public keys are stored on the device. Clients are authenticated using these stored public keys. Only clients that have a private key that corresponds to one of the stored public keys can gain access to the device using SSH server.

With password authentication, users are prompted for a password when they attempt to log into the device (provided empty password logins are not allowed; refer to Enabling empty password logins on page 427). If there is no user account that matches the user name and password supplied by the user, the user is not granted access.

You can deactivate one or both user authentication methods for SSH server. Note that deactivating both authentication methods essentially disables the SSH server entirely.

To disable DSA challenge-response authentication.

```
device(config)# ip ssh
key-authentication no
```

Syntax: **ip ssh key-authentication yes | no**

The default is "yes".

To deactivate password authentication.

```
device(config)# ip ssh password-authentication no
```

Syntax: **ip ssh password-authentication no | yes**

The default is "yes".

Enabling empty password logins

By default, empty password logins are not allowed. This means that users with an SSH client are always prompted for a password when they log into the device. To gain access to the device, each user must have a user name and password. Without a user name and password, a user is not granted access. Refer to "Setting up local user accounts" for information on setting up user names and passwords on the device.

If you enable empty password logins, users are not prompted for a password when they log in. Any user with an SSH client can log in without being prompted for a password.

To enable empty password logins.

```
device(config)# ip ssh permit-empty-passwd yes
```

Syntax: ip ssh permit-empty-passwd no | yes

Setting the SSH server port number

By default, SSH server traffic occurs on TCP port 22. You can change this port number. For example, the following command changes the SSH server port number to 2200.

```
device(config)# ip ssh port 2200
```

> **NOTE**
> If you change the default SSH server port number, you must configure SSH clients to connect to the new port. Also, you should be careful not to assign SSH server to a port that is used by another service. If you change the SSH server port number, it is recommended that you change it to a port number greater than 1024.

Syntax: ip ssh port number

Setting the SSH server login timeout value

When the SSH server attempts to negotiate a session key and encryption method with a connecting client, it waits a maximum of 120 seconds for a response from the client. If there is no response from the client after 120 seconds, the SSH server disconnects. You can change this timeout value to between 1 - 120 seconds. For example, to change the timeout value to 60 seconds.

```
device(config)# ip ssh timeout 60
```

Syntax: ip ssh timeout seconds

> **NOTE**
> The standard for the idle-timeout RADIUS attribute is for it to be implemented in seconds as opposed to the minutes that the device router uses. If this attribute is used for setting idle time instead of this configuration, the value from the idle-timeout RADIUS attribute will be converted to seconds and truncated to the nearest minute.

Designating an interface as the source for all SSH server packets

You can designate a loopback interface, virtual interface, or Ethernet port as the source for all SSH server packets from the device. The software uses the IP address with the numerically lowest value configured on the port or interface as the source IP address for SSH server packets originated by the device.

> **NOTE**
> When you specify a single SSH server source, you can use only that source address to establish SSH server management sessions with the device.

To specify the numerically lowest IP address configured on a loopback interface as the device's source for all SSH server packets, enter commands such as a the following.

```
device(config)# int loopback 2
device(config-lbif-2)# ip address 10.0.0.2/24
device(config-lbif-2)# exit
device(config)# ip ssh source-interface loopback 2
```

The commands in this example configure loopback interface 2, assign IP address 10.0.0.2/24 to the interface, then designate the interface as the source for all SSH server packets from the device.

Syntax: ip ssh source-interface ethernet slot/port | loopback num | ve num

The *num* parameter is a loopback interface or virtual interface number. The *slot/port* parameter specifies an ethernet port number.

```
device(config)# interface ethernet 1/4
device(config-if-e10000-1/4)# ip address 10.157.22.110/24
device(config-if-e10000-
1/4)# exit
device(config)# ip ssh source-interface ethernet 1/4
```

Configuring maximum idle time for SSH server sessions

By default, SSH server sessions do not time out. Optionally, you can set the amount of time an SSH server session can be inactive before the device closes it. For example, to set the maximum idle time for SSH server sessions to 30 minutes.

```
device(config)# ip ssh idle-time 30
```

Syntax: ip ssh idle-time minutes

If an established SSH server session has no activity for the specified number of minutes, the device closes it. An idle time of 0 minutes (the default value) means that SSH server sessions never time out. The maximum idle time for SSH server sessions is 240 minutes.

> **NOTE**
> The standard for the idle-timeout RADIUS attribute is for it to be implemented in seconds as opposed to the minutes that the device router uses. If this attribute is used for setting idle time instead of this configuration, the value from the idle-timeout RADIUS attribute will be converted from seconds to minutes and truncated to the nearest minute.

Filtering SSH server access using ACLs

You can permit or deny SSH server access to the device using ACLs. To configure an ACL that restricts SSH server access to the device, enter commands such as the following.

```
device(config)# access-list 12 deny host 10.157.22.98
device(config)# access-list 12 deny 10.157.23.0 10.0.0.255
device(config)# access-list 12 deny 10.157.24.0/24
device(config)# access-list 12 permit any
device(config)# ssh access-group 12
device(config)# write memory
```

Syntax: ssh access-group { num | name | ipv6 ipv6-acl-name }

Use the **ipv6** keyword if you are applying an IPv6 access list.

The *num* parameter specifies the number of a standard IPv4 ACL, 1 - 99.

The *name* parameter specifies a standard IPv4 access list name.

The *ipv6-acl-name* parameter specifies an IPv6 access list name.

These commands configure ACL 12, then apply the ACL as the access list for SSH server access. The device denies SSH server access from the IPv4 addresses listed in ACL 12 and permits SSH server access from all other IP addresses. Without the last ACL entry for permitting all packets, this ACL would deny SSH server access from all IP addresses.

> **NOTE**
> Access control lists are IP version specific. When both IPv4 and IPv6 ACLs are configured, the IPv4 ACL will be applied to sessions from IPv4 clients and the IPv6 ACL will be applied to sessions from IPv6 clients.

Refer to Filtering SSH server access using ACLs and Filtering SSH server access using ACLs for details on how to configure ACLs.

Disabling 3-DES

By default, both 3-DES and AES encryption algorithms are enabled on the device. You can disable 3-DES by entering the following command.

```
device(config)# ip ssh encryption aes-only
```

Syntax: [no] ip ssh encryption aes-only

Displaying SSH server connection information

A maximum of 16 SSH server connections can be active on the device at a given time. To display information about SSH server connections, enter the following command.

```
device# show ip ssh
    Session      Encryption     HostKey         Username        IP Address
    Inbound:
    Outbound:
     17          aes256-cbc    ssh-dss                          labuser         10.37.73.155
    SSH session type codes: N - Netconf, S - Scp

    SSH-v2.0 disabled
```

Syntax: show ip ssh [| begin expression | exclude expression | include expression]

This display shows the following information about the active SSH server connections.

This field...	Displays...
Session	The SSH server connection ID. This can be from 1 - 16.
Version	The SSH server version number. This should always be SSH-2.
Encryption	The encryption method used for the connection.
Username	The user name for the connection if password authentication is used. If public key authentication is used, username shows *none* . In the output, the username is truncated to 8 characters.

The **show who** command also displays information about SSH server connections.

```
device#show who
Console connections:
established, monitor enabled, in config mode
2 minutes 17 seconds in idle
Telnet connections (inbound):
1 closed
2 closed
3 closed
4 closed
5 closed
```

```
Telnet connection (outbound):
6 closed
SSH connections:
1 established, client ip address 10.43.2.4, user is hanuma
1 minutes 16 seconds in idle
2 established, client ip address 10.50.3.7, user is Mikaila
you are connecting to this session
18 seconds in idle
3 established, client ip address 10.47.8.20, user is Jenny
1 minutes 39 seconds in idle
4 established, client ip address 10.55.3.9, user is Mariah
41 seconds in idle
5 established, client ip address 10.9.4.11, user is Logan
23 seconds in idle
```

Syntax: show who [| begin expression | exclude expression | include expression]

Ending an SSH server connection

To terminate one of the active SSH server connections, enter the following command.

```
device# kill ssh 1
```

Syntax: kill ssh connection-id

Outbound SSHv2 client

SSH2 client allows you to connect from a device to an SSH2 server, including another device that is configured as an SSH2 server. You can start an outbound SSH2 client session while you are connected to the device by any connection method (SSH2, Telnet, console). Devices support one outbound SSH2 client session at a time.

The supported SSH2 client features are as follows:

- Encryption algorithms, in the order of preference:
    - aes256-ctr
    - aes192-ctr
    - aes128-ctr
    - aes256-cbc
    - aes192-cbc
    - aes128-cbc
    - 3des-cbc
- SSH2 client session authentication algorithms:
    - Password authentication
    - Public Key authentication
- Message Authentication Code (MAC) algorithm: hmac-sha1
- Key exchange algorithm: diffie-hellman-group1-sha1 and diffie-hellman-group14-sha1
- Compression algorithms are not supported.
- The client session can be established through either in-band or out-of-band management ports.
- The client session can be established through IPv4 or IPv6 protocol access.
- The client session can be established to a server listening on a non-default SSH server port.

Enabling SSHv2 client

When SSH2 server is enabled, you can use SSH client to connect to an SSH server using password authentication.

Configuring SSH2 client public key authentication

To use SSH client for public key authentication, you must generate SSH client authentication keys and export the public key to the SSH servers to which you want to connect.

The following sections describe how to configure SSH client public key authentication:

Generating and deleting a client DSA key pair

Client keys are independent of host keys. Both DSA and RSA client keys can co-exist in the system. The RSA client key will be used for outbound session when both exist. To generate a client DSA key pair, enter the following command.

```
device(config)#crypto key client generate dsa
```

To delete the DSA host key pair, enter the following command.

```
device(config)#crypto key client zeroize dsa
```

Syntax: crypto key client generate | zeroize dsa

The **generate** keyword places a host key pair in the flash memory.

The **zeroize** keyword deletes the host key pair from the flash memory.

The **dsa** keyword specifies a DSA host key pair.

Generating and deleting a client RSA key pair

Client keys are independent of host keys. Both DSA and RSA client keys can co-exist in the system. The RSA client key will be used for outbound session when both exist. To generate a client RSA key pair, enter a command such as the following:

```
device(config)#crypto key client generate rsa modulus 2048
```

To delete the RSA host key pair, enter the following command.

```
device(config)#crypto key client zeroize rsa
```

Syntax: crypto key client generate | zeroize rsa [modulus *modulus-size* **]**

The **generate** keyword places an RSA host key pair in the flash memory.

The **zeroize** keyword deletes the RSA host key pair from the flash memory.

The optional [**modulus**] parameter specifies the modulus size of the RSA key pair, in bits. The valid values for *modulus-size* are 1024 or 2048. It is used only with the **generate** parameter. The default value is 1024.

The **rsa** keyword specifies an RSA host key pair.

Exporting client public keys

Client public keys are stored in the following files in flash memory:

- A DSA key is stored in the file **$$sshdsapub.key** .
- An RSA key is stored in the file **$$sshrsapub.key** .

To copy key files to a TFTP server, you can use the **copy flash tftp** command.

To upload the client key to TFTP server, use a command such as the following.

```
device#copy flash tftp 10.37.73.154 client.key $$sshdsapub.key
```

Syntax: **copy flash tftp ip-addr client.key $$sshdsapub.key**

Importing client public keys

To download the client key to SSHv2 sever, use a command such as the following.

```
device(config)# ip ssh pub-key-file tftp 10.37.73.154 client.key
```

Syntax: **ip ssh pub-key-file tftp ip-addr client.key**

You must copy the public key to the SSH server. If the SSH server is a device supported by this guide, see the section Importing authorized public keys into the device on page 423.

Using an SSH2 client

The following sections describe how to configure SSH client:

- Initiating a SSH2 client on page 432
- Designating an interface as the outbound SSH session on page 433
- Ending an outbound SSH session on page 433

Initiating a SSH2 client

To start an SSH2 client connection to an SSH2 server using password authentication, enter a command such as the following:

```
device# ssh 10.10.10.2
```

To start an SSH2 client connection to an SSH2 server using public key authentication, enter a command such as the following:

```
device# ssh 10.10.10.2 public-key dsa
```

Syntax: **ssh** [*ipv6*] [**vrf** *vrf*] *ipv4-addr* | *ipv6-addr* | *host-name* [port][**outgoing-interface**{**ethernet**|**ve**}][**public-key**{**dsa**|**rsa**}]

To make IPv6 connections to SSH server, use parameter [ipv6] followed by IPv6 address.

SSH requests will be initiated only from the ports belonging to the specified *vrf*.The default value for *vrf* parameter is default-vrf.

The default value for port number is 22.

The parameter outgoing-interface {ethernet|ve} is applicable to IPv6 connections only.

To bring up public-key based client session, use the parameters [public-key {dsa|rsa}].

By default password based client session will be brought up.

Designating an interface as the outbound SSH session

You can designate a loopback interface, virtual interface, or Ethernet port as the outbound SSH session.

To specify an IP address as a loopback interface, enter commands such as a the following.

```
device(config)# int loopback 2
device(config-lbif-2)# ip address 10.0.0.2/24
device(config-lbif-2)# exit
device(config)# ip ssh source-interface loopback 2
```

To specify an IP address as an Ethernet port, enter commands such as a the following.

```
device(config)# interface ethernet 1/4
device(config-if-e10000-1/4)# ip address 10.157.22.110/24
device(config-if-e10000-
1/4)# exit
device(config)# ip ssh source-interface ethernet 1/4
```

Syntax: ip ssh source-interface ethernet slot/port | loopback num | ve num

The *slot/port* parameter specifies an ethernet port number.

The *num* parameter is a loopback interface or virtual interface number.

Ending an outbound SSH session

To clear an outbound SSH session, enter a command such as the following.

```
device# kill ssh 17
```

Syntax: kill *connection-id*

Displaying SSH2 client information

For information about displaying SSH2 client information, see the following sections:

- Displaying SSH server connection information on page 429
- Displaying SSH server connection information on page 429

Using Secure Copy

Secure Copy (SCP) uses security built into SSH server to transfer files between hosts on a network, providing a more secure file transfer method than Remote Copy (RCP), FTP, or TFTP. SCP automatically uses the authentication methods, encryption algorithm, and data compression level configured for SSH server. For example, if password authentication is enabled for SSH server, the user is prompted for a user name and password before SCP allows a file to be transferred. No additional configuration is required for SCP on top of SSH server.

You can use SCP to copy files on the device, including the startup configuration and running configuration files, to or from any other device running an SCP program (referred to here as an "SCP-enabled remote host").

SCP is enabled on the device by default and can be disabled. To disable SCP, enter the following command.

```
device(config)# ip ssh scp disable
```

Syntax: ip ssh scp disable | enable

NOTE
If you disable SSH server, SCP is also disabled.

NOTE
When using SCP to transfer files, you enter the **scp** commands on the SCP-enabled remote host, rather than the console on the device.

NOTE
Certain SCP client options, including -p and -r, are ignored by the SCP server on the device. If an option is ignored, the client is notified.

NOTE
```
User access privileges to enable SCP:
```

- The user should have Super User privilege level (allows complete read-and-write access to the system) to enable or use SCP.
- If password authentication is enabled for SSH server, the user is prompted for user password before the file transfer takes place.

NOTE
All SCP features are supported on the devices. However, some of the command options are unavailable on the CER 2000 Series and CES 2000 Series because they are not applicable to the CER 2000 Series and CES 2000 Series hardware (e.g., copying files to or from a Auxiliary flash card).

NOTE
The **scp** command will not display any **"success message"** on completion of data transfer. Instead use **showlog** command to validate scp image success.

NOTE
The SCP command, "runConfig-overwrite", is not supported on both MLX and CES/
CER.

Secure Copy feature for CES 2000 Series and CER 2000 Series

To copy a file to flash.

```
C:\> scp c:\<src-file> terry@192.168.1.50:flash:<dst-file>
```

To copy and append a configuration file (c:\cfg\brocadehp.cfg) to the running configuration file on a device at 192.168.1.50 and log in as user terry, enter the following command on the SCP-enabled client.

```
C:\> scp c:\cfg\brocadehp.cfg terry@192.168.1.50:runConfig
```

If you are copying the configuration file from the device to a PC or another machine (outbound), the command saves the running configuration file to the PC. If you are copying a configuration file from a PC to the device, (inbound) the command appends the source file to the running configuration file on the device.

If password authentication is enabled for SSH server, the user is prompted for user terry's password before the file transfer takes place.

When a configuration file is loaded using the Secure Copy feature, the following commands within the configuration file are supported.

- **isis metric command**
- **set-overload-bit command**
- **admin-group**
- **cspf-group**
- **bypass-lsp**

If you are copying a configuration file from a PC to the device, (inbound) the command replaces the source file on the device.

To copy the configuration file to the startup configuration file.

```
C:\> scp c:\cfg\brocadehp.cfg terry@192.168.1.50:startConfig
```

To copy the configuration file to a file called config1.cfg on the Auxiliary flash card in slot 1 on a management module.

```
C:\> scp c:\cfg\brocade.cfg terry@192.168.1.50:slot1:/config1.cfg
```

To copy the configuration file to a file called config1.cfg on the Auxiliary flash card in slot 2 on a management module.

```
C:\> scp c:\cfg\brocade.cfg terry@192.168.1.50:slot2:/config1.cfg
```

To copy the running configuration file on an device to a file called c:\cfg\fdryhprun.cfg on the SCP-enabled client.

```
C:\> scp terry@192.168.1.50:runConfig c:\cfg\fdryhprun.cfg
```

To copy the startup configuration file on a device to a file called c:\cfg\fdryhpstart.cfg on the SCP-enabled client.

```
C:\> scp terry@192.168.1.50:startConfig c:\cfg\fdryhpstart.cfg
```

To copy the software image (for example, xmr03300b228.bin) to the primary flash.

```
C:\> scp c:\xmr03300b228.bin local@192.168.1.50:flash:primary
```

To copy the software image (for example, xmr03300b228.bin) to the secondary flash.

```
C:\> scp c:\xmr03300b228.bin local@192.168.1.50:flash:secondary
```

Secure Copy Feature for XMR Series

The following encryption cipher algorithms are supported on the XMR Series. They are listed in order of preference:

- **aes256-cbc:** AES in CBC mode with 256-bit key
- **aes192-cbc:** AES in CBC mode with 192-bit key
- **aes128-cbc:** AES in CBC mode with 128-bit key
- **3des-cbc:** Triple-DES

Outbound commands:

The following is the list of outbound SCP command options supported (Upload from device to host).

The general syntax of the outbound SCP commands is as follows.

Syntax: scp user@IP Address:Source:src-name dst-file

src-name can be abbreviated

To copy the running configuration file on a device to a file on the SCP-enabled host.

```
C:\> scp <user>@<device-IpAddress>:runConfig <dst-file>
```

> **NOTE**
> If you are copying the running configuration file from the device to a PC or another machine (outbound), the command saves the running configuration file to the PC. If you are copying a configuration file from a PC to the device, (inbound) the command appends the source file to the running configuration file on the device.

To copy the startup configuration file on the device to a file on the SCP-enabled client.

```
C:> scp <user>@<device-IpAddress>:startConfig <dst-file>
```

To copy the MP primary image file from the device to a file on the SCP-enabled client.

```
C:> scp <user>@<device-IpAddress>:flash:primary <primary-image-file>
```

To copy the MP secondary image file from the device to a file on the SCP-enabled client.

```
C:> scp <user>@<device-IpAddress>:flash:secondary <secondary-image-file>
```

To copy a flash file from the device to a file on the SCP-enabled client.

```
C:> scp <user>@<device-IpAddress>:flash:<src-file> <dst-file>
```

To copy a file on Auxiliary flash slot from the device to a file on the SCP-enabled client (this command is not applicable to the CES 2000 Series or CER 2000 Series devices).

```
C:> scp <user>@<device-IpAddress>:slot1:/<src-file> <dst-file>
C:> scp <user>@<device-IpAddress>:slot2:/<src-file> <dst-file>
```

Inbound commands:

The following is the list of inbound SCP command options supported (for downloading from an SCP-enabled client to the device). The commands copy the files from the SCP-enabled client to the specified location on the device.

The general syntax of the Inbound SCP commands is as follows.

Syntax: scp file-name user@IP Address:Destination:file-name [:additional-options]

The last two tokens *file-name* and *additional-options* can be abbreviated. The others cannot be abbreviated.

Auxiliary flash command option

To download a file and store it in a Auxiliary flash (Slot 1 or Slot 2), enter the following command (not applicable to the CES 2000 Series or CER 2000 Series devices).

```
C:> scp <src-file> <user>@<device-IpAddress>:slot1:/<dst-file>
```

This commands transfers *src-file* to the device and saves it as /slot1/*dst-file*.

Flash MP command options

To download a file and store in MP Flash, enter the following command.

```
C:> scp <src-file> <user>@<device-IpAddress>:flash:<dst-file>
```

This command transfers *src-file* to the device and saves it as *dst-file* in flash

To download and replace MP Monitor image in Flash, enter the following command.

```
C:> scp <monitor-image-file> <user>@<device-IpAddress>:flash:monitor
```

This command transfers *monitor-image-file* to the device and replaces MP monitor image in flash.

To download and replace MP Primary image in Flash, enter the following command.

```
C:> scp <primary-image-file> <user>@<device-IpAddress>:flash:primary
```

This command transfers `primary-image-file` to the device and replaces MP Primary image in flash.

To download and replace MP secondary image in Flash, enter the following command.

```
C:> scp <secondary-image-file>
 <user>@<device-IpAddress>
:flash:secondary
```

This command transfers `secondary-image-file` to the device and replaces MP secondary image in flash.

To download and replace MP boot image in Flash, enter the following command.

```
C:> scp <boot-image-file> <user>@<device-IpAddress>:flash:boot
```

This command transfers `boot-image-file` to the device and replaces MP boot image in flash.

Running configuration command options

To download a configuration file and append to running configuration, enter the following command.

```
C:> scp <config-file> <user>@<device-IpAddress>:config:run
```

This command transfers `config-file` to the device and appends to the running configuration.

When a configuration file is loaded using the Secure Copy feature, the following commands within the configuration file are supported.

- **isis metric command**
- **set-overload-bit command**
- **admin-group**
- **cspf-group**
- **bypass-lsp**

For backward compatibility, the following syntax is also supported for this command.

```
C:> scp <config-file> <user>@<device-IpAddress>:runConfig
```

Startup configuration command options

To download a configuration file and replace startup configuration, enter the following command.

```
C:> scp <config-file> <user>@<device-IpAddress>:config:start
```

This command transfers `config-file` to the device and replaces the startup configuration in flash.

For backward compatibility, the following syntax is also supported for this command.

```
C:> scp <config-file> <user>@<device-IpAddress>:startConfig
```

Combined image command options

> NOTE
> SCP combined image command options are not applicable to the CES 2000 Series or CER 2000 Series devices.

To download a combined image file and replace LP and MP primary.

```
C:> scp <combined-image-file>   <user>@<device-IpAddress>:image:primary
```

This command transfers `combined-image-file` to the device and replaces MP and LP Primary image in flash.

To download a combined image file and replace LP primary and MP secondary.

```
C:> scp <combined-image-file>   <user>@<device-IpAddress>:image:mp-sec
```

This command transfers `combined-image-file` to the device and replaces MP Secondary and LP Primary image in flash.

To download a combined image file and replace LP secondary and MP primary, enter the following command.

```
C:> scp <combined-image-file> <user>@<device-IpAddress>:image:lp-sec
```

This command transfers `combined-image-file` to the device and replaces MP Primary and LP Secondary image in flash.

To download a combined image file and replace LP and MP secondary, enter the following command.

```
C:> scp <combined-image-file> <user>@<device-IpAddress>:image:secondary
```

This command transfers `combined-image-file` to the device and replaces MP and LP Secondary image in flash.

MBRIDGE command options

> **NOTE**
> SCP MBRIDGE command options are not applicable to the CES 2000 Series or CER 2000 Series devices.

To download and replace FPGA (mbridge) file in MP, enter the following command.

```
C:> scp <image-file>      <user>@<device-IpAddress>:mbridge
```

This command downloads `image-file` and replaces the mbridge image on the flash.

Switch fabric options

> **NOTE**
> SCP switch fabric options are not applicable to the CES 2000 Series or CER 2000 Series devices.

To download and replace switch fabric file to a single SNM or all in MP, enter the following command.

```
C:> scp <image-file>      <user>@<device-IpAddress>:snm:sbridge:
<snm-index>
```

This command downloads `image-file` and replaces sbridge image on the specified SNM.

To download and replace sbridge image on all SNMs, enter the following command.

```
C:> scp <image-file>      <user>@<device-IpAddress>:snm:sbridge:all
```

This command downloads `image-file` and replaces the sbridge image on all the SNMs.

LP command options

> **NOTE**
> SCP LP command options are not applicable to the CES 2000 Series or CER 2000 Series devices.

To download and over-write the LP boot image on one LP or all LPs, enter the following command.

```
C:> scp <lp-boot-image-file> <user>@<device-IpAddress>:lp:boot:<lp-slot-num>
```

This command transfers `lp-boot-image-file` to the device and replaces the LP boot image in the specified LP slot.

```
C:> scp <lp-boot-image-file> <user>@<device-IpAddress>:lp:boot:all
```

This command transfers `lp-boot-image-file` to the device and replaces LP boot image in all the LP slots

To download and over-write the LP monitor image on one LP or all LPs, enter the following command.

```
C:> scp <lp-monitor-image-file> <user>@<device-IpAddress>:lp:monitor:<lp-slot-num>
```

This command transfers `lp-monitor-image-file` to the device and replaces the LP monitor image in the specified LP slot.

```
C:> scp <lp-monitor-image-file> <user>@<device-IpAddress>:lp:monitor:all
```

This command transfers `lp-monitor-image-file` to the device and replaces LP monitor image in all LP slots.

To download and over-write LP primary image on one LP or all LPs, enter the following commands.

```
C:> scp <lp-primary-file> <user>@<device-IpAddress>:lp:primary:<lp-slot-num>
```

This command transfers `lp-primary-file` to the device and replaces the LP Primary image in the specified LP slot.

```
C:> scp <lp-primary-file> <user>@<device-IpAddress>:lp:primary:all
```

This command transfers `lp-primary-file` to the device and replaces the LP Primary image in all the LP slots.

To download and over-write the LP secondary image on one LP or all LPs, enter the following commands.

```
C:> scp <lp-secondary-file> <user>@<device-IpAddress>:lp:secondary:<lp-slot-num>
```

This command transfers `lp-secondary-file` to the device and replaces LP Secondary image in the specified LP slot.

```
C:> scp <lp-secondary-file> <user>@<device-IpAddress>:lp:secondary:all
```

This command transfers `lp-secondary-file` to the device and replaces the LP Secondary image in all the LP slots.

Bundled FPGA command options

> **NOTE**
> SCP bundled FPGA command options are not applicable to the CES 2000 Series or CER 2000 Series devices.

> **NOTE**
> If force-overwrite is present in the command, the command skips compatibility checks and forcibly replaces the FPGA image, otherwise the command checks for compatibility of the FPGA image and if the check fails the FPGA image is not replaced and error message is returned to the SCP client.

To download and over-write Bundled FPGA image, enter the following commands.

```
C:> scp <fpga-bundle-file> <user>@<device-IpAddress>:lp:fpga-all:<lp-slot-num>
```

This command downloads `fpga-bundle-file` and replaces all the FPGA images (PBIF, STATS, XGMAC and XPP) on the specified LP.

```
C:> scp <fpga-bundle-file> <user>@<device-IpAddress>:lp:fpga-all:all
```

This command downloads `fpga-bundle-file` and replaces all the FPGA images (PBIF, STATS, XGMAC and XPP) on all the LPs.

To download and force overwrite Bundled FPGA image, enter the following.

```
C:> scp <fpga-bundle-file> <user>@<device-IpAddress>:lp:fpga-all:<lp-slot-num>:force-overwrite
```

This command downloads `fpga-bundle-file` and replaces all the FPGA images (PBIF, STATS, XGMAC and XPP) on the specified LP.

```
C:> scp <fpga-bundle-file> <user>@<device-IpAddress>:lp:fpga-all:all:force-overwrite
```

This command downloads `fpga-bundle-file` and replaces all the FPGA images (PBIF, STATS, XGMAC and XPP) on all the LPs.

PBIF FPGA command options

> **NOTE**
> When downloading a PBIF FPGA image onto a CES 2000 Series or CER 2000 Series device, either use the keyword **all** or enter "1" for the LP slot number.

> **NOTE**
> If force-overwrite is present in the command, the command skips compatibility checks and forcibly replaces the FPGA image, otherwise the command checks for compatibility of the FPGA image and if the check fails, the FPGA image is not replaced and error message is returned to the SCP client.

To download and over-write PBIF FPGA image, enter the following command.

```
C:> scp <fpga-pbif-file> <user>@<device-IpAddress>:lp:fpga-pbif:<lp-slot-num>
```

This command downloads `fpga-pbif-file` and replaces the FPGA PBIF image on the specified LP.

```
C:> scp <fpga-pbif-file> <user>@<device-IpAddress>:lp:fpga-pbif:all
```

This command downloads `fpga-pbif-file` and replaces FPGA PBIF image on all the LPs.

To download and force over-write PBIF FPGA image, enter the following command.

```
C:> scp <fpga-pbif-file> <user>@<device-IpAddress>:lp:fpga-pbif:<lp-slot-num>:force-overwrite
```

This command downloads `fpga-pbif-file` and replaces FPGA PBIF image on the specified LP.

```
C:> scp <fpga-pbif-file> <user>@<device-IpAddress>:lp:fpga-pbif:all:force-overwrite
```

This command downloads `fpga-pbif-file` and replaces the FPGA PBIF image on all the LPs.

STATS FPGA command options

> **NOTE**
> STATS FPGA command options are not applicable to the CES 2000 Series or CER 2000 Series devices.

> **NOTE**
> If force-overwrite is present in the command, the command skips compatibility checks and forcibly replaces the FPGA image, otherwise the command checks for compatibility of the FPGA image and if the check fails, the FPGA image is not replaced and error message is returned to the SCP client.

To download and over-write STATS FPGA image, enter the following.

```
C:> scp <fpga-stats-file> <user>@<device-IpAddress>:lp:fpga-stats:<lp-slot-num>
```

This command downloads `fpga-stats-file` and replaces FPGA STATS image on the specified LP.

```
C:> scp <fpga-stats-file> <user>@<device-IpAddress>:lp:fpga-stats:all
```

This command downloads fpga-stats-file and replaces FPGA STATS image on all the LPs.

To download and force over-write STATS FPGA image, enter the following command.

```
C:> scp <fpga-stats-file> <user>@<device-IpAddress>:lp:fpga-stats:<lp-slot-num>:force-overwrite
```

This command downloads fpga-stats-file and replaces FPGA STATS image on the specified LP.

```
C:> scp <fpga-stats-file> <user>@<device-IpAddress>:lp:fpga-stats:all:force-overwrite
```

This command downloads fpga-stats-file and replaces FPGA STATS image on all the LPs.

XGMAC FPGA command options

> **NOTE**
> XGMAC FPGA command options are not applicable to the CES 2000 Series or CER 2000 Series devices.

> **NOTE**
> If force-overwrite is present in the command, the command skips compatibility checks and forcibly replaces the FPGA image, otherwise the command checks for compatibility of the FPGA image and if the check fails, the FPGA image is not replaced and error message is returned to the SCP client.

To download and over-write XGMAC FPGA image, enter the following commands.

```
C:> scp <fpga-xgmac-file> <user>@<device-IpAddress>:lp:fpga-xgmac:<lp-slot-num>
```

This command downloads `fpga-xgmac-file` and replaces FPGA XGMAC image on the specified LP.

```
C:> scp <fpga-xgmac-file> <user>@<device-IpAddress>:lp:fpga-xgmac:all
```

This command downloads `fpga-xgmac-file` and replaces FPGA XGMAC image on all the LPs.

To download and force over-write XGMAC FPGA image, enter the following commands.

```
C:> scp <fpga-xgmac-file> <user>@<device-IpAddress>:lp:
fpga-xgmac:<lp-slot-num>:force-overwrite
```

This command downloads `fpga-xgmac-file` and replaces FPGA XGMAC image on the specified LP.

```
C:> scp <fpga-xgmac-file>    <user>@<device-IpAddress>:lp:fpga-xgmac:all:force-overwrite
```

This command downloads `fpga-xgmac-file` and replaces FPGA XGMAC image on all the LPs.

XPP FPGA command options

> **NOTE**
> XPP FPGA command options are not applicable to the CES 2000 Series or CER 2000 Series devices.

> **NOTE**
> If force-overwrite is present in the command, the command skips compatibility checks and forcibly replaces the FPGA image, otherwise the command checks for compatibility of the FPGA image and if the check fails, the FPGA image is not replaced and error message is returned to the SCP client.

To download and over-write an XPP FPGA image, enter the following commands.

```
C:> scp <fpga-xpp-file> <user>@<device-IpAddress>:lp:fpga-xpp:<lp-slot-num>
```

This command downloads `fpga-xpp-file` and replaces FPGA XPP image on the specified LP.

```
C:> scp <fpga-xpp-file> <user>@<device-IpAddress>:lp:fpga-xpp:all
```

This command downloads `fpga-xpp-file` and replaces FPGA XPP image on all the LPs.

To download and force over-write XPP FPGA image, enter the following commands.

```
C:> scp <fpga-xpp-file> <user>@<device-IpAddress>:lp:fpga-xpp:< lp-slot-num>:force-overwrite
```

This command downloads fpga-xpp-file and replaces FPGA XPP image on the specified LP.

```
C:> scp <fpga-xpp-file> <user>@<XMR-IpAddress>:lp:fpga-xpp:all:force-overwrite
```

This command downloads fpga-xpp-file and replaces FPGA XPP image on all the LPs.

Delete old file first option

> **NOTE**
> The delete file first option only applies to inbound SCP commands; its purpose is make room in the MP flash by deleting old image files prior to an image download.

An option "delete-first" is provided in the third or fourth token position in the following commands.

```
C:> scp <image-file> <user>@<device-IpAddress>:image:<primary|secondary|mp-sec|lp-sec>:delete-first
C:> scp <image-file> <user>@<device-IpAddress>:flash:<primary|secondary|monitor>:delete-first
C:> scp <image-file> <user>@<device-IpAddress>:lp:<primary|secondary|monitor>:all:delete-first
```

Without **delete-first** option, if the flash is full these commands should display the following message.

"There is not enough space on MP flash. Please clean up MP flash and retry, or use \"delete-first\" option."

When the **delete-first** option is specified, these commands clear space in the MP flash by removing the following files first.

```
image:primary, "primary", "lp-primary-0"
image:secondary, "secondary", "lp-secondary-0"
image:lp-sec, "primary", "lp-secondary-0"
image:mp-sec, "secondary", "lp-primary-0"
flash:primary, "primary"
flash: secondary, "secondary"
flash: monitor, "monitor"
lp:primary:all, "lp-primary-0"
lp:secondary:all, "lp-secondary-0"
lp:monitor:all, "lp-monitor-0
"
```

Before deleting the file the system will check to see if deleting the file or files will create enough space in the flash. If it can create enough space to accommodate the download, the files will be deleted. Otherwise, the command will fail with the following error message.

```
"There will not be enough space on MP flash even after deleting the target files. Please clean up MP flash
and retry."
```

> **NOTE**
> Other commands will not check for available space in the flash or delete the file before downloading. In other words, the delete-first option is not supported for commands not mentioned above.

I

SCP client

Secure copy (SCP) supports file transfer between local and a remote hosts. It combines the file-transfer element of BSD remote copy (RCP) with the authentication and encryption provided by the Secure shell (SSH) protocol.

The SCP client feature on NetIron devices helps to transfer files to and from the SCP server and maintains the confidentiality of the data being transferred by blocking packet sniffers from extracting valuable information from the data packets. You can use SCP client to do the following:

- Download a boot file, NetIron application image file, signature file, license file, startup configuration file, or running configuration from an SCP server
- Upload a NetIron application image file, startup configuration file, or running configuration to an SCP server

SCP client uploads the file to the SCP server (that is, the SSH server) by providing files to be uploaded. You can specify file attributes, such as permissions and time-stamps as part of file data when you use SCP client to upload files. SCP client supports the same copy features as the time stamps, TFTP client feature on NetIron devices, but the SSH2 protocol secures data transfer.

SCP client support limitations

SCP client sessions are limited by file size and by whether other SCP client sessions are running and by whether SC server sessions are in progress.

The following limitations apply to SCP client sessions:

- An SCP copy of the running or startup configuration file from a device to Linux may fail if the configuration size is less than 700 bytes.
- Only one SCP client session is supported at a time.
- An SCP client session cannot be initiated if an SCP server session is in progress.
- An SSH client outbound session cannot be initiated if an SCP client session is in progress from the same terminal.
- There is only one outbound SSH session. SCP client and SSH client share this session. Only one application can use it at a time. When an additional attempt to use this session is made while the session is being used, an error message is displayed.
- Only one file transfer operation can be allowed at a given time. When SCP client file transfer is on-going, TFTP transfer will be prohibited, and vice-versa.
- Simplified upgrade (and LP auto-upgrade) through SCP client is not be supported in NetIron release 5.8.00. Image upgrade through SCP client can be achieved by downloading the image and saving it to the CF, and utilizing the existing CLI commands that are used to upgrade images using the locally saved image files. The same process can be used for all the images that are required for full system upgrade.
- Whenever SCP client is used for downloading a file to flash and the destination file name is given as any of the reserved file names on flash, the command line interface (CLI) will display an error. This avoids corrupting of the already available image files in flash when the user attempts to download any non-image file and download it as an image name. However, this restriction is only for downloading to flash and not applicable to slot 1 and 2.

Supported SCP client configurations

SCP client automatically uses the authentication methods, encryption algorithm, and data compression level configured for SSH.

For example, if password authentication is enabled for SSH, you are prompted for a user name and password before SCP allows a file to be transferred.

The following conditions also apply:

- SCP is enabled by default and can be enabled or disabled using the **ip ssh scp disable | enable** command.
- If SSH is disabled, SCP is disabled automatically.
- The SCP client session uses one SSH outbound client session.
- Because the SCP client internally uses the SSH2 client for creating outbound SSH sessions from the device, all configurations related to the SSH2 client are required for SCP client support, as described here:
  - The SSH2 server on the device must be enabled by creating an SSH server DSA or RSA key pair; otherwise, the SSH2 client cannot be used.
  - You can use the **crypto key client { generate | zeroize } dsa** command to generate or delete an SSH-client-DSA key pair. The SSH-client-DSA public key is stored in the file - $$sshdsapub.key.
  - You can use the **crypto key client generate rsa [modulus 1024 | 2048]** command to generate an SSH-client-RSA key pair. The SSH-client-RSA public key is stored in the file $$sshrsapub.key.
  - You can use the **crypto key client zeroize rsa** command to delete an SSH-client-RSA key pair.

Uploading an image to an SCP server

To securely upload image files to a secure copy (SCP) server, copy an image from a device to the SCP server.

```
device# copy flash scp 10.20.99.146 ~/xmr05800.bin primary
```

Uploading configuration files to an SCP server

To securely upload startup and running configuration files to a secure copy (SCP) server.

1. Copy a startup configuration file to the SCP server.

```
Device#copy startup-config  scp 10.20.1.1 icx-74-startup
```

The startup configuration file is uploaded to the SCP server and you are notified when the transfer is complete.

```
device#copy startup-config scp 172.20.133.116 run.txt
User name:tester
Password:
Connecting to remote host......
Connection Established. Uploading .Done.
startup-config upload via SCP complete.

Connection Closed
device#
```

2. Copy a running configuration file to the SCP server.

```
Device#copy running-config  scp 10.20.1.1 icx-74-run
```

Downloading configuration files from an SCP server

To securely download startup and running configuration files from a secure copy (SCP) server to a device.

1. Copy a startup configuration file from the SCP server.

```
device# copy scp startup-config 10.20.1.1 icx-74-startup
```

2. Copy a running configuration file from the SCP server.

```
device# copy scp running-config 10.20.1.1 icx-74-run
```

Joint Interoperability Test Command

JITC overview

The Joint Interoperability Test Command (JITC) mode on a device is compliant with the standards established by JITC, a United States military organization that oversees testing of national security systems and information technology systems for hardware, software and components. Some services include developmental, operational, and validation testing.

The JITC mode implemented on a device enforces default behavior for some features to ensure strict JITC certification compliance.

AES-CTR encryption mode support for SSH

The Advanced Encryption Standard - Cipher Block Chaining (AES-CBC) encryption mode for Secure Shell (SSH) is vulnerable to certain plain-text attacks. JITC mandates that AES-CBC mode be disabled, and only AES-CTR mode be used. The JITC mode of operation implements this requirement in both SSH client and server modes.

In the JITC mode, by default, the AES-CBC encryption mode for SSH is disabled and the AES-CTR (Counter) encryption mode is enabled. To enable the JITC mode, use the **jitc enable** command. The **ip ssh encryption disable-aes-cbc** command that disables the AES-CBC mode can be seen in the running configuration. The encryption algorithms such as aes256-ctr, aes192-ctr, or aes128-ctr are enabled and the CBC mode ciphers are removed. The AES-CTR encryption mode for SSH is available in JITC and non-JITC mode.

The following table lists encryption algorithm keys in different operation modes in the preferred order by the sender.

TABLE 56 SSH ciphers with JITC support

Mode	Ciphers supported (listed in order of precedence)
Standard mode	aes256-ctr, aes192-ctr, aes128-ctr, aes256-cbc, aes192-cbc, aes128-cbc, 3des-cbc
Standard mode with the **ip ssh encryption aes-only** command enabled	aes256-ctr, aes192-ctr, aes128-ctr, aes256-cbc, aes192-cbc, aes128-cbc
JITC mode with the **ip ssh encryption aes-only** command and **ip ssh encryption disable- aes-cbc** command enabled.	aes256-ctr, aes192-ctr, aes128-ctr

The AES-CBC mode can be re-enabled by issuing the **no ip ssh encryption disable-aes-cbc** command, which will bring back the pre-existing CBC ciphers (aes256-cbc, aes192-cbc, aes128-cbc) along with the CTR ciphers.

> NOTE
> The AES-CTR mode must be configured both on the client and server sides to establish an SSH connection.

SHA1 authentication support for NTP

In the JITC mode, the symmetric key scheme supported for cryptographic authentication of messages uses the SHA1 keyed hash algorithm instead of the MD5 authentication scheme. The MD5 authentication for Network Time Protocol (NTP) is disabled by default in the JITC mode, using the **disable authentication md5** command. The **disable authentication md5** command is displayed in the running configuration. The SHA1 authentication scheme is available to define the authentication key for NTP. SHA1 authentication must be

enabled manually using the **authentication-key** command with the *key-id* **sha1** *key-string* options. The SHA1 authentication configuration is available in JITC and non-JITC mode.

> **NOTE**
> Extreme does not recommend re-enabling the MD5 authentication in JITC mode using the **no disable authentication md5** command.

Acknowledgements

This appendix presents the acknowledgements for portions of code from various vendors that are included in the Extreme devices covered in this manual.

Cryptographic software

MPL 1.1

The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at http://www.mozilla.org/MPL/ .

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is the Network Security Services libraries.

The Initial Developer of the Original Code is Red Hat, Inc. Portions created by the Initial Developer are Copyright (C) 2009 the Initial Developer. All Rights Reserved.

Portions created by Netscape Communications Corporation are Copyright (C) 1994-2000 Netscape Communications Corporation. All Rights Reserved.

Contributor(s):

OpenSSL license

Copyright (c) 1998-2001 The OpenSSL Project. All rights reserved.

1. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

2. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

3. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation or other materials provided with the distribution.

4. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (http://www.openssl.org/)"

5. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org .

6. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

7. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/)"

    THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com) All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL. This library is free for commercial and non-commercial use as long as the following conditions are aheared to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com). Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

1. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

2. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.

3. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. All advertising materials mentioning features or use of this software must display the following acknowledgment: "This product includes cryptographic software written by Eric Young(eay@cryptsoft.com)" The word 'cryptographic' can be left out if the rouines from the library being used are not cryptographic related.

4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgment: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

    THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE. The licence and distribution terms for any publically available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution licence.

Cryptographic software

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Original SSLeay License

Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)/. All rights reserved.

This package is an SSL implementation written by Eric Young (eay@cryptsoft.com). The implementation was written so as to conform with Netscape's SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

1. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

2. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.

3. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation or other materials provided with the distribution.

4. All advertising materials mentioning features or use of this software must display the following acknowledgement:

   "This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)" The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related:-).

5. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement:

   "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

   THIS SOFTWARE IS PROVIDED BY ERIC YOUNG "AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

   The license and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution license [including the GNU Public License.]