

Extreme NetIron Management Configuration Guide, 06.2.00f

Supporting NetIron OS 06.2.00f

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, see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

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

Contents

Preface.....	9
Conventions.....	9
Notes, cautions, and warnings.....	9
Text formatting conventions.....	9
Command syntax conventions.....	10
Documentation and Training.....	10
Getting Help.....	10
Subscribe to Service Notifications.....	11
Providing Feedback.....	11
About This Document.....	13
What's new in this document.....	13
Supported hardware and software.....	13
Supported software.....	14
How command information is presented in this guide.....	14
Configuration Fundamentals.....	15
Interface parameters.....	15
Assigning a port name.....	16
Bulk port naming.....	16
Assigning port name to multiple ports.....	16
Assigning an IP address to a port.....	16
Configuring IP addresses.....	17
Assigning an IP address to an Ethernet port.....	17
Assigning an IP address to a loopback interface.....	18
Assigning an IP address to a virtual interface.....	18
Modifying port speed.....	19
Modifying port mode.....	20
Auto negotiation speed limit.....	20
Disabling or re-enabling a port.....	21
Disabling Source Address Learning on a port.....	21
Changing the default Gigabit negotiation mode.....	21
Changing the negotiation mode.....	22
Designating an interface as the packet source.....	22
Configuring an interface as the source for all Telnet packets.....	22
Cancelling an outbound Telnet session.....	23
Configuring an interface as the source for all SSH packets.....	23
Configuring an interface as the source for all TFTP packets.....	23
Configuring an interface as the source for all TACACS or TACACS+ packets.....	24
Configuring an interface as the source for all RADIUS packets.....	24
Setting IP VPN packets with a TTL value of 1 to be dropped.....	24
Disabling or re-enabling flow control.....	24
Flow control configuration at the global level.....	25
Flow control configuration at the interface level.....	25
Flow control configuration at MIF.....	25
Enabling and disabling interactivity for scripts	25
Entering system administration information.....	27

Setting the system clock.....	27
DST "change" notice for networks using US time zones	28
Creating a command alias.....	28
Removing an alias.....	28
Displaying a list of all configured alias.....	29
Configuring CLI banners.....	29
Setting a message of the day banner.....	29
Setting a privileged EXEC CLI level banner.....	30
Displaying a message on the console when an incoming Telnet session is detected.....	30
Configuring terminal display.....	30
Checking the length of terminal displays.....	31
Enabling or disabling Layer 2 switching	31
Configuring static MAC addresses	32
Disabling the MAC movement console logs.....	33
Changing the MAC age time.....	33
Configuring system max values	34
Maintaining system-max configuration with available system resources.....	38
Management VRF overview.....	38
Source interface and management VRF compatibility.....	39
Supported management applications.....	39
Configuring a global management VRF.....	42
Displaying the management VRF information.....	42
Bootup time.....	45
L2 elements.....	47
L3 elements.....	47
VPLS elements.....	48
Miscellaneous elements.....	48
Bootup time message.....	48
Configuration time.....	49
Monitoring dynamic memory allocation.....	49
Commands that require a reload.....	50
Verifying an image checksum.....	51
Configuring CAM mode globally.....	51
Configuring density mode for the 2x100G and 20x10G CAM.....	52
Configuring IPv6 host CAM mode.....	52
Configuring IPv6 host drop CAM limit.....	52
Configuring -X2 Algorithmic CAM profiles.....	53
CAM partition profiles.....	54
Configuring CAM partition size	84
CAM overflow logging.....	84
Disabling CAM table entry aging.....	85
Data integrity protection.....	85
Configuring detection parameters.....	85
Port transition hold timer.....	89
Port flap dampening.....	90
Configuring port link dampening on an interface.....	90
Configuring port link dampening on a LAG.....	90
Re-enabling a port disabled by port link dampening.....	90
Displaying ports configured with port link dampening.....	90
Port loop detection.....	91

Strict Mode and Loose Mode.....	91
Recovering disabled ports.....	91
Disable duration and loop detection interval.....	92
Enabling loop detection.....	92
Configuring a global loop detection interval.....	93
Configuring the device to automatically re-enable ports.....	93
Clearing loop detection.....	94
Displaying loop detection information.....	94
Discarding loop detection frames in the LACP-blocked port.....	95
Syslog message.....	95
Displaying information for an interface for an Ethernet port.....	95
Displaying the full port name for an Ethernet interface.....	95
Displaying statistics information for an Ethernet port.....	98
Monitoring Ethernet port statistics in real time.....	98
Displaying recent traffic statistics for an Ethernet port.....	102
Displaying and modifying default settings for system parameters.....	103
Network Time Protocol.....	109
Network Time Protocol overview.....	109
Network Time Protocol leap second	111
How Extreme supports leap second handling for NTP.....	111
How NTP works.....	111
NTP server.....	111
NTP client.....	111
NTP peer.....	111
NTP broadcast server.....	112
NTP broadcast client.....	112
Synchronizing time.....	112
Configuring NTP.....	113
Changing to the NTP mode.....	113
Configuring the NTP client.....	114
Configuring the NTP peer.....	115
Configuring NTP on an interface.....	115
Show commands.....	117
Displaying NTP status.....	117
Displaying NTP associations.....	118
Displaying NTP associations details.....	119
Configuration examples.....	120
Packet timestamping.....	122
Supported hardware.....	122
Configuring packet timestamping.....	122
Management VRF for NTP.....	122
Restrictions for NTP support of the management VRF.....	123
Enabling Management VRF for NTP.....	123
Displaying NTP support of the management VRF.....	124
Cisco Discovery Protocol.....	125
Cisco Discovery Protocol overview.....	125
Enabling CDP packet interception.....	125
Displaying CDP packet information.....	126
Clearing CDP statistics and neighbor information.....	127

Network Configuration Protocol.....	129
NETCONF protocol introduction.....	129
Platforms.....	130
Related documentation.....	130
NETCONF in client/server architecture.....	130
RPC request	131
RPC reply.....	131
RPC and error handling.....	132
CLI and SSH subsystem.....	132
Recommendations for NETCONF.....	132
Basic NETCONF operations.....	133
Initial connection.....	133
get operation.....	134
get-config operation.....	137
edit-config operation.....	145
Closing sessions.....	147
NETCONF commands and specifications.....	147
Clients establishing a NETCONF session with NetIron devices.....	151
PuTTY Link.....	151
Netopeer.....	152
Linux OpenSSH.....	154
Data models and mapping.....	155
Example in YANG, XML, and CLI.....	156
Foundry Discovery Protocol.....	157
Foundry Discovery Protocol overview.....	157
Enabling FDP.....	157
Advertising IPv4 or IPv6 management addresses to FDP neighbors.....	158
Verifying FDP.....	158
Clearing FDP statistics and neighbor information.....	160
High Availability.....	161
High availability overview.....	161
Management module redundancy configuration.....	161
Changing the default active chassis slot.....	162
Managing management module redundancy.....	162
File synchronization between active and standby management modules.....	162
Manually switching over to the standby management module.....	164
Rebooting the active and standby management modules.....	164
Monitoring management module redundancy.....	165
Determining management module status.....	165
Monitoring the status change of a module.....	166
Displaying temperature information.....	166
Displaying switchover information.....	166
Flash memory and auxiliary flash card file management commands.....	167
Verifying available flash space on the management module before an image is copied.....	168
Management focus.....	169
Flash memory file system.....	170
Auxiliary flash card file system.....	170
Wildcards.....	171
Formatting a flash card.....	172

Determining the current management focus.....	172
Switching the management focus.....	173
Displaying a directory of the files.....	173
Displaying the contents of a file.....	175
Displaying the hexadecimal output of a file.....	175
Creating a subdirectory.....	176
Removing a subdirectory.....	177
Renaming a file.....	178
Changing the read-write attribute of a file.....	178
Deleting a file.....	179
Recovering ("undeleting") a file.....	180
Appending a file to another file.....	181
Copying files using the copy command.....	181
Copying files using the cp command.....	185
Loading the software.....	186
Saving configuration changes.....	187
File management messages.....	188
LLDP.....	189
Link layer discovery protocol overview.....	189
General operating principles.....	190
Operating modes.....	191
LLDP packets.....	191
TLV support.....	191
Configuration considerations.....	194
Using LLDP.....	195
Enabling LLDP.....	195
Changing the operating mode of a port.....	195
Specifying the maximum number of LLDP neighbors	196
Enabling bridging of LLDP BPDUs when LLDP not enabled.....	197
Enabling LLDP SNMP notifications and Syslog messages.....	197
Specifying the minimum time between SNMP traps and Syslog messages.....	197
Changing the minimum time between LLDP transmissions.....	198
Changing the interval between regular LLDP transmissions.....	198
Changing the holdtime multiplier for transmit TTL.....	198
Changing the minimum time between port reinitializations.....	199
LLDP TLVs advertised by the Extreme device.....	199
Displaying LLDP statistics and configuration settings.....	205
Resetting LLDP statistics.....	209
SNMP.....	211
SNMP overview.....	211
Encryption of SNMP community strings.....	211
Adding an SNMP community string.....	211
Displaying the SNMP community strings.....	212
Using the User-Based Security model.....	213
Configuring your NMS.....	213
Configuring SNMP version 3 on the device.....	213
Defining the engine ID.....	214
Defining an SNMP group.....	214
Defining an SNMP user account.....	215

Displaying the engine ID.....	216
Displaying SNMP groups.....	216
Displaying user information.....	217
Interpreting varbinds in report packets.....	217
Defining SNMP views.....	218
SNMP v3 configuration examples.....	218
Simple SNMP v3 configuration.....	219
More detailed SNMP v3 configuration.....	219
Configuring SNMP traps.....	219
Specifying an SNMP trap receiver.....	219
Specifying a single trap source.....	220
Setting the SNMP trap holddown time.....	221
Disabling SNMP traps.....	221
Configuring SNMP management of VRFs	222
SNMPv3 polling	222
SNMPv1/v2c polling	223
Configuring SNMP ifIndex	224
On Extreme NetIron CES and CER only.....	224
On Extreme NetIron XMR and MLX Series only.....	224
SNMP scalability optimization.....	225
Configuring SNMP throughput optimization.....	225
Configuring SNMP load throttling.....	225
Configuring SNMP to revert ifType to legacy values	226
Configuring snAgentConfigModuleType to return original values.....	226
Preserving interface statistics in SNMP.....	227
NIAP-CCEVS.....	229
NIAP-CCEVS-certified Extreme equipment and IronWare releases.....	229
Web management access to NIAP-CCEVS-certified Extreme equipment.....	230
Warning: local user password changes.....	230

Preface

- Conventions..... 9
- Documentation and Training.....10
- Getting Help.....10
- Providing Feedback.....11

This section describes the text conventions used in this document, where you can find additional information, and how you can provide feedback to us.

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

Documentation and Training

Find Extreme Networks product information at the following locations:

- [Current Product Documentation](#)
- [Release Notes](#)
- [Hardware and software compatibility](#) for Extreme Networks products
- [Extreme Optics Compatibility](#)
- [Other resources](#) such as white papers, data sheets, and case studies

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

Getting Help

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

- Extreme Portal** Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.
- The Hub** A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- Call GTAC** For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

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

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

Subscribe to Service Notifications

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

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

NOTE

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

4. Select **Submit**.

Providing Feedback

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

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

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

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.com.

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

About This Document

- [What's new in this document](#)..... 13
- [Supported hardware and software](#)..... 13
- [How command information is presented in this guide](#)..... 14

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.

There has been no enhancement to this guide for the NI 6.3.00a software release.

On October 30, 2017, Extreme Networks, Inc. acquired the SLX-OS product line from Brocade Communications Systems, Inc. This transitional release includes references to both companies.

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 1 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
		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*.

How command information is presented in this guide

Starting with Extreme NetIron 5.6.00, command syntax and parameter descriptions are removed from commands that are referenced in configuration tasks. To find the full description of a specific command, including all required and optional keywords and variables, refer to the *Extreme NetIron Command Reference* for your software release.

Configuration Fundamentals

• Interface parameters.....	15
• Assigning a port name.....	16
• Bulk port naming.....	16
• Assigning port name to multiple ports.....	16
• Assigning an IP address to a port.....	16
• Configuring IP addresses.....	17
• Modifying port speed.....	19
• Modifying port mode.....	20
• Disabling or re-enabling a port.....	21
• Disabling Source Address Learning on a port.....	21
• Changing the default Gigabit negotiation mode.....	21
• Designating an interface as the packet source.....	22
• Setting IP VPN packets with a TTL value of 1 to be dropped.....	24
• Disabling or re-enabling flow control.....	24
• Enabling and disabling interactivity for scripts	25
• Entering system administration information.....	27
• Setting the system clock.....	27
• Creating a command alias.....	28
• Configuring CLI banners.....	29
• Configuring terminal display.....	30
• Enabling or disabling Layer 2 switching	31
• Configuring static MAC addresses	32
• Disabling the MAC movement console logs.....	33
• Changing the MAC age time.....	33
• Configuring system max values	34
• Maintaining system-max configuration with available system resources.....	38
• Management VRF overview.....	38
• Bootup time.....	45
• Configuration time.....	49
• Monitoring dynamic memory allocation.....	49
• Commands that require a reload.....	50
• Verifying an image checksum.....	51
• Configuring CAM mode globally.....	51
• Data integrity protection.....	85
• Port transition hold timer.....	89
• Port flap dampening.....	90
• Port loop detection.....	91
• Displaying information for an interface for an Ethernet port.....	95
• Displaying statistics information for an Ethernet port.....	98
• Displaying and modifying default settings for system parameters.....	103

Interface parameters

All Extreme device ports are pre-configured with default values that allow the device to be fully operational at initial startup without any additional configuration. In some configuration scenarios, changes to the port parameters may be necessary to adjust to attached devices or other network requirements.

Assigning a port name

A port name can be assigned to help identify interfaces on the network. You can assign a port name to physical ports, virtual routing interfaces, and loopback interfaces.

To assign a name to a port, enter the following command.

```
device(config)# interface e 2/8
device(config-if-e10000-2/8)# port-name Marsha Markey
```

Syntax: **[no]** **port-name** *text*

The *text* parameter is an alphanumeric string. The name can have up to 255 characters on a Extreme device and can include blanks. You do not need to use quotation marks around the string, even when it contains blanks.

Bulk port naming

Beginning with the Extreme NetIron release 6.2.0, you can assign alphanumeric name to multiple ports in one go.

With the bulk port naming feature, you can select a range of ports and assign alphanumeric name to the selected ports. When you assign a name to the selected port range by using the **port-name** command, the name reflects in all ports. The **no port-name** command removes any names assigned to the individual ports.

Assigning port name to multiple ports

Use the **port-name** command to assign a name to a range of ports.

Use the following steps to assign an alphanumeric name to multiple ports in an interface.

1. Enter the global configuration mode.

```
switch# configure terminal
```

2. Enter the interface configuration mode.

```
switch(config)# interface ethernet 2/1 to 2/10
```

3. Assign a port name.

```
switch(config)# port-name 10GPORT
```

Assigning an IP address to a port

To assign an IP address to an interface, enter the following commands.

```
device(config)# interface e 1/8
device(config)# ip address 10.45.6.110 255.255.255.0
```

Syntax: **[no]** **ip address** *ip-addr ip-mask*

or

Syntax: **[no]** **ip address** *ip-addr/mask-bits*

NOTE

You also can enter the IP address and mask in CIDR format, as follows.

```
device(config)# ip address 10.45.6.1/24
```

Configuring IP addresses

You can configure an IP address on the following types of interfaces:

- Ethernet port
- Virtual routing interface (also called a Virtual Ethernet or "VE")
- Loopback interface

By default, you can configure up to 24 IP addresses on each interface.

NOTE

After you configure a virtual routing interface on a VLAN, you cannot configure Layer 3 interface parameters on individual ports in the VLAN. Instead, you must configure the parameters on the virtual routing interface itself. Also, after an IP address is configured on an interface, the hardware is programmed to route all IP packets that are received on the interface. Consequently, all IP packets not destined for this device's MAC address are not bridged and are dropped.

The device supports both classical IP network masks (Class A, B, and C subnet masks, and so on) and Classless Interdomain Routing (CIDR) network prefix masks.

- To enter a classical network mask, enter the mask in IP address format. For example, enter "10.157.22.99 255.255.255.0" for an IP address with a Class-C subnet mask.
- To enter a prefix network mask, enter a forward slash (/) and the number of bits in the mask immediately after the IP address. For example, enter "10.157.22.99/24" for an IP address that has a network mask with 24 significant bits (ones).

By default, the CLI displays network masks in classical IP address format (example: 255.255.255.0). You can change the display to prefix format.

Assigning an IP address to an Ethernet port

To assign an IP address to port 1/1, enter the following commands.

```
device(config)# interface ethernet 1/1
device(config-if-e1000-1/1)# ip address 10.45.6.1 255.255.255.0
```

NOTE

You also can enter the IP address and mask in CIDR format, as follows.

```
device(config-if-e1000-1/1)# ip address 10.45.6.1/24
```

Syntax: [no] interface ethernet slot/port

Syntax: [no] ip address ip-addr ip-mask | ip-addr/mask-bits [ospf-ignore | ospf-passive | secondary]

The **ospf-ignore** and **ospf-passive** parameters modify the device defaults for adjacency formation and interface advertisement. Use one of these parameters if you are configuring multiple IP subnet addresses on the interface but you want to prevent OSPF from running on some of the subnets:

- **ospf-passive** - disables adjacency formation with OSPF neighbors (but does not disable advertisement of the interface into OSPF). By default, when OSPF is enabled on an interface, the software forms OSPF router adjacencies between each primary IP address on the interface and the OSPF neighbor attached to the interface.

- **ospf-ignore** - disables OSPF adjacency formation and advertisement of the interface into OSPF. The subnet is completely ignored by OSPF.

Use the **secondary** parameter if you have already configured an IP address within the same subnet on the interface.

NOTE

When you configure more than one address in the same subnet, all but the first address are secondary addresses and do not form OSPF adjacencies.

Assigning an IP address to a loopback interface

Loopback interfaces are always up, regardless of the states of physical interfaces. They can add stability to the network because they are not subject to route flap problems that can occur due to unstable links between this device and other devices.

You can configure up to 64 loopback interfaces on this device.

NOTE

This number is static. Through **system-max** command, number of supported loopback interfaces can be configured. The range of valid values is between 64 to 1024. You can add up to 24 IP addresses to each loopback interface.

NOTE

If you configure the device to use a loopback interface to communicate with a BGP4 neighbor, you also must configure a loopback interface on the neighbor and configure the neighbor to use that loopback interface to communicate with the device.

To add a loopback interface, enter command as shown in the following example:

```
device(config-bgp-router)# exit
device(config)# int loopback 1
device(config-lbif-1)# ip address 10.0.0.1/24
```

Syntax: [no] interface loopback num

To configure range from 64 to 1024, enter command as shown in the following example:

```
device(config)# system-max loopback-interface ?
device(config)# DECIMAL Valid range 64 to 1024 (default: 64)
```

Assigning an IP address to a virtual interface

A virtual interface is a logical port associated with a Layer 3 Virtual LAN (VLAN) configured on this device.

NOTE

The device uses the lowest MAC address on the device (the MAC address of port 1 or 1/1) as the MAC address for all ports within all virtual interfaces you configure on the device.

To add a virtual interface to a VLAN and configure an IP address on the interface, enter commands such as the following.

```
device(config)# vlan 2 name IP-Subnet_10.1.2.0/24
device(config-vlan-2)# untag e1/1 to 1/4
device(config-vlan-2)# router-interface ve1
device(config-vlan-2)# interface ve1
device(config-vif-1)# ip address 10.1.2.1/24
```

The first two commands create a Layer 3 protocol-based VLAN named "IP-Subnet_10.1.2.0/24" and add a range of untagged ports to the VLAN. The **router-interface** command creates virtual interface 1 as the routing interface for the VLAN. The last two commands change to the interface configuration level for the virtual interface and assign an IP address to the interface.

Syntax: [no] router-interface ve num

Syntax: `[no] interface ve num`

The *num* parameter specifies the virtual interface number. You can specify from 1 to the maximum number of virtual interfaces supported on the device. To display the maximum number of virtual interfaces supported on the device, enter the **show default values** command. The maximum is listed in the System Parameters section, in the Current column of the virtual-interface row.

NOTE

This number is static. Through **system-max** command, number of supported loopback interfaces can be configured. The range of valid values is between 64 to 1024. You can add up to 24 IP addresses to each loopback interface.

Modifying port speed

Each of the 10/100/1000BaseTX ports is designed to auto-sense and auto-negotiate the speed and mode of the connected device. If the attached device does not support this operation, you can manually enter the port speed to operate at either 10 Mbps or 100 Mbps. The default value is 10 or 100 half- or full-duplex.

NOTE

Modifying the port speed of a port that has a pre-configured rate limit policy may result in the inability to remove the port's rate limit policy.

To change the port speed of interface 1/8 from the default of 10/100 auto-sense to 10 Mbps operating at full-duplex, first disable the port. Then, enter the following.

```
device(config)# interface e 1/8
device(config-if-e10000-1/8)# speed-duplex 10-full
```

Syntax: `[no] speed-duplex value`**NOTE**

The speed-duplex configuration is applicable to the first four combination ports of the Extreme NetIron CES 2024F-4X module and not applicable to the remaining fiber ports. This is specific to combination ports when the fiber link is connected.

The *value* can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- 1000-full
- 1000-half
- auto

The default is auto.

NOTE

An auto negotiation port must be connected to another auto negotiation port. If you connect an auto negotiation port to a fixed speed or duplex port, the behavior is undefined. Also, ports must be disabled before changing speed.

Modifying port mode

You can configure a port to accept either full-duplex (bi-directional) or half-duplex (uni-directional) traffic. Port duplex mode and port speed are modified by the same command.

To change the port speed of interface 1/8 from the default of 10/100 auto-sense to 10 Mbps operating at full-duplex, enter the following command.

```
device(config)# interface e 1/8
device(config-if-e10000-1/8)# speed-duplex 10-full
```

Syntax: `speed-duplex value`

The *value* can be one of the following:

- 10-full
- 10-half
- 100-full
- 100-half
- 1000-full
- 1000-half
- auto

The default is auto.

Auto negotiation speed limit

Auto-negotiation is an active method of determining the link mode. Each interface is expected to transmit specific information in a specific format. If an interface that is expecting to use auto-negotiation does not receive this information from the other side, it assumes the other side cannot detect or change its mode.

One of the most common causes of performance issues on 10/100/1000 Mb Ethernet links occurs when one port on the link operates at half-duplex while the other port operates at full-duplex. This occurs when one or both ports on a link are reset and the auto-negotiation process does not result in both link partners having the same configuration. It also can occur when users reconfigure one side of a link and forgets to reconfigure the other side. Both sides of a link should have auto-negotiation on, or both sides should have it off.

The auto negotiation speed limit feature allows the user to reduce or limit the port speed when auto- negotiation is configured. You can set the port to automatically reduce the speed from 1000 Mb to 100 Mb or 10 Mb. The **down-shift** option will reduce the port speed to 100 Mb from 1000 Mb automatically once a 2-wire cable is detected.

The auto negotiation speed limit feature is supported only on FIXED (non SFP) copper ports when auto-neg is ON.

```
device(config)# interface e 1/8
device(config-if-e10000-1/8)# auto
device(config-if-e10000-1/8)# link-config gig copper autoneg-control down-shift
```

Syntax: `[no] link-config gig copper autoneg-control [down-shift | 100m | 10m]`

The **10m** option will limit the port to negotiation to speeds and duplex of 10 Mb.

The **100m** option will limit the port to negotiation of speeds and duplex below 100 Mb.

Disabling or re-enabling a port

The port can be made inactive (disable) or active (enable) by selecting the appropriate status option. The default value for a port is disabled.

To disable port 8 on module 1 of a Extreme device, enter the following command.

```
device(config)# interface e 1/8
device(config-if-e10000-1/8)# disable
```

Syntax: [no] disable

You also can disable or re-enable a virtual routing interface. To do so, enter commands such as the following.

```
device(config)# interface ve 1
device(config-vif-1)# disable
```

To re-enable a virtual routing interface, enter the **enable** command at the Interface configuration level. For example, to re-enable virtual routing interface v1, enter the following command.

```
device(config-vif-1)# enable
```

Syntax: [no] enable

Disabling Source Address Learning on a port

The default operation is for Source Address (SA) Learning to be enabled on all ports. It can be useful to disable SA Learning on a port in situations where high CPU usage is occurring because a large number of packets are being sent to the CPU for SA Learning. For example, it can be useful to disable SA Learning on physical ports that are part of a Virtual Ethernet (VE) interface that has no need to switch packets.

SA Learning can be disabled on a port using the **sa-learning-disable** command as shown in the following.

```
device(config)# interface e 1/8
device(config-if-e10000-1/8)# sa-learning-disable
```

Syntax: [no] sa-learning-disable

Changing the default Gigabit negotiation mode

You can configure the default Gigabit negotiation mode to be one of the following:

- **neg-full-auto** - The port is only for copper-SFP and to support 10/100/1000M tri-speed auto negotiation.
- **auto-full** - The port tries to perform a negotiation with its peer port to exchange capability information. If it is unable to reach an agreed upon speed, the port goes into a fixed speed and keeps the link up.
- **auto-gig** - The port tries to perform a negotiation with its peer port to exchange capability information. This is the default state.
- **neg-off** - The port does not try to perform a negotiation with its peer port.

Unless the ports at both ends of a Gigabit Ethernet link use the same mode (either **auto-gig** or **neg-off**), the ports cannot establish a link. An administrator must intervene to manually configure one or both sides of the link to enable the ports to establish the link.

NOTE

XMR Series, MLX Series, and Extreme MLXe Series support **auto-gig** and **neg-off** options. The **neg-full-auto** and **auto-full** options are not supported on the chassis platforms. Extreme CES and CER-RT series support all four options.

NOTE

Support is provided for the following modules:

- 20x10GE
- 4x10GE-IPSEC

NOTE

Double link flap is observed on the 20x10GE and 4x10GE-IPSEC port modules once the remote peer CER comes back up after reload.

Changing the negotiation mode

You can change the negotiation mode for individual ports as shown in the following.

```
device(config)# interface ethernet 4/1 to 4/4
device(config-mif-4/1-4/4)# gig-default neg-off
```

This command changes the default **auto-gig** setting and sets the negotiation mode to neg-off for ports 4/1 through 4/4.

Use the **auto-gig** option to activate auto-negotiation.

```
device(config)# interface ethernet 4/1 to 4/4
device(config-mif-4/1-4/4)# gig-default auto-gig
```

Syntax: [no] gig-default neg-full-auto | auto-gig | neg-off | auto-full

Designating an interface as the packet source

The software uses the lowest-numbered IP address configured on an interface as the source IP address for all Telnet, SSH, NTP, TFTP, TACACS or TACACS+, or RADIUS packets originated from the Extreme device.

You can specify the source interface for one or more of these types of packets.

Configuring an interface as the source for all Telnet packets

Identifying a single source IP address for Telnet packets provides the following benefits:

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

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

```
device(config)# interface loopback 2
device(config-lbif-2)# ip address 10.0.0.2/24
device(config-lbif-2)# exit
device(config)# ip telnet source-interface loopback 2
```

The commands configure loopback interface 2, assign IP address 10.0.0.2/24 to it, then designate it as the source for all Telnet packets from the Extreme device.

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

The following commands configure an IP interface on an Ethernet port and designate the address port as the source for all Telnet packets from the Extreme device.

```
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 telnet source-interface ethernet 1/4
```

Cancelling an outbound Telnet session

If you want to cancel a Telnet session from the console to a remote Telnet server (for example, if the connection is frozen), you can terminate the Telnet session by performing the following steps.

1. At the console, press Ctrl-^ (Ctrl-Shift-6).
2. Press the X key to terminate the Telnet session.

Pressing Ctrl-^ twice in a row causes a single Ctrl-^ character to be sent to the Telnet server. After you press Ctrl-^, pressing any key other than X or Ctrl-^ returns you to the Telnet session.

Configuring an interface as the source for all SSH packets

You can configure the Extreme device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, or Ethernet port as the source for the SSH packets it sends.

For example, to specify an Ethernet port as the interface whose lowest-numbered IP address will be the source address for the SSH packets originated from the Extreme device, enter commands such as the following.

```
device(config)# interface ethernet 1/5
device(config-if-e10000-1/5)# ip address 10.157.22.111/24
device(config-if-e10000-1/5)# exit
device(config)# ip ssh source-interface ethernet 1/5
```

The commands configure Ethernet port 1/5, assign IP address 10.157.22.111/24 to it, then designate it as the source interface.

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

Configuring an interface as the source for all TFTP packets

You can configure the Extreme device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, or Ethernet port as the source for all TFTP packets it sends.

For example, to specify the lowest-numbered IP address configured on a virtual routing interface as the Extreme device's source for all TFTP packets, enter commands such as the following.

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

The commands configure virtual routing interface 1, assign IP address 10.0.0.3/24 to it, then designate the address as the source address for all TFTP packets.

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

The default is the lowest-numbered IP address configured on the port through which the packet is sent. The address therefore changes, by default, depending on the port.

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

You can configure the Extreme device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, or Ethernet port as the source for the TACACS or TACACS+ packets it sends.

For example, to specify a virtual routing interface as the interface whose lowest-numbered IP address will be the source address for the TACACS or TACACS+ packets originated from the Extreme device, enter commands such as the following.

```
device(config)# interface 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 configure virtual routing interface 1, assign IP address 10.0.0.3/24 to it, then designate it as the source interface.

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

Configuring an interface as the source for all RADIUS packets

You can configure the Extreme device to use the lowest-numbered IP address configured on a loopback interface, virtual routing interface, or Ethernet port as the source for the RADIUS packets it sends.

For example, to specify an Ethernet port as the interface whose lowest-numbered IP address will be the source address for the RADIUS packets originated from the Extreme device, enter the following commands.

```
device(config)# interface ethernet 1/5
device(config-if-e10000-1/5)# ip address 10.157.22.111/24
device(config-if-e10000-1/5)# exit
device(config)# ip radius source-interface ethernet 1/5
```

The commands configure Ethernet port 1/5, assign IP address 10.157.22.111/24 to it, then designate it as the source interface.

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

Setting IP VPN packets with a TTL value of 1 to be dropped

This command is for IP VPN packets only. Under normal conditions IP VPN packets with a TTL value equal to 0 are always dropped in hardware regardless of the setting of this command. With this command set, IP VPN packets with TTL value equal to one will also be dropped in hardware.

To enable this command use the following command.

```
device(config)# interface ethernet 4/1 to 4/4
device(config-if-4/1)# hw-drop-bad-ttl-pkt
```

Syntax: `[no] hw-drop-bad-ttl-pkt`

The default value is off.

Disabling or re-enabling flow control

You can configure full-duplex ports on a system to operate with or without flow control (802.3x).

The command to disable or enable flow control is **flow-control** command. This command is enabled by default.

To disable flow control on full-duplex ports on a system, enter the following command.

```
device(config)# no flow-control rx-pause-ignore
```

To turn the feature back on, enter the following command.

```
device(config)# flow-control rx-pause-ignore
```

Syntax: [no] flow-control rx-pause-ignore

Flow control configuration at the global level

The **flow-control rx-pause-ignore** command at the global level will configure rx-pause-ignore and disable flow-control for all the ports.

The command **no flow-control rx-pause-ignore** at the global level will disable rx-pause-ignore and enable flow-control for all the ports.

Flow control configuration at the interface level

The commands **flow-control** or **no flow-control rx-pause-ignore** at the interface level will enable flow-control and disable rx-pause-ignore for a specific interface.

Similarly, the commands **no flow-control** or **flow-control rx-pause-ignore** at the interface level will will disable flow-control and enable rx-pause-ignore for a particular interface.

Flow control configuration at MIF

The commands **mif-flow-control** or **no mif-flow-control rx-pause-ignore** at the interface level will enable flow-control and disable rx-pause-ignore for a particular interface.

Similarly, the commands **no mif-flow-control** or **mif-flow-control rx-pause-ignore** at the interface level will disable flow-control and enable rx-pause-ignore for a particular interface.

Enabling and disabling interactivity for scripts

[Table 2](#) lists certain configuration and action commands that are interactive by default.

Because these commands require a user response, confirmation, or result in multiple changes across the system before the device can complete the configuration changes, they cannot be used in scripts as they are. You can, however, disable the interactive behavior by entering the **prompt** command.

Syntax: [no] prompt

The **no prompt** command will only disable the confirmation prompt for commands in configuration mode. Commands executed in the EXEC mode will continue to prompt for confirmation.

Entering the **no prompt** command allows you to use the commands and actions that are listed in [Table 2](#) within scripts without difficulty. After running a script, you can re-enable the default interactive behavior by entering the **prompt** command.

TABLE 2 Interactive commands

Command type	Command
Configuration	cluster-l2protocol-forward

TABLE 2 Interactive commands (continued)

Command type	Command
	route-only rate-limit policy-map spanning-tree pms disable enable violation deny violation restrict violation shutdown
Action	reboot-standby reset reload switchover power-off lp all slot power-off power-supply index forced hitless-reload mp primary secondary lp primary secondary power-supply monitoring clear all index boot system flash primary secondary

Default behavior for certain configuration commands:

```
device(config)# route-only
This will change the route-only behavior at the global level.
Are you sure? (enter 'y' or 'n'): y
Global 'route-only' committed.
```

Disabling default behavior to allow for script use:

```
device(config)# no prompt
device(config)# no route-only
Global 'no route-only' committed.
```

Re-enabling default behavior:

```
device(config)# prompt

device(config)# route-only
This will change the route-only behavior at the global level.
Are you sure? (enter 'y' or 'n'): y
Global 'route-only' committed.
```

Entering system administration information

You can configure a system name, contact, and location for the Extreme device and save the information locally in the configuration file for future reference. The information is not required for system operation but recommended. When you configure a system name, it replaces the default system name in the CLI command prompt.

To configure a system name, contact, and location, enter commands such as the following.

```
device(config)# hostname home
home(config)# snmp-server contact Suzy Sanchez
home(config)# snmp-server location Centerville
home(config)# end
home# write memory
```

The system name you configure **home** replaces the system name **device**.

Syntax: [no] **hostname** *string*

Syntax: [no] **snmp-server contact** *string*

Syntax: [no] **snmp-server location** *string*

The *string* for the hostname, contact, and location each can be up to 255 alphanumeric characters. The text strings can contain blanks. The SNMP text strings do not require quotation marks when they contain blanks but the host name does.

Setting the system clock

The Extreme device allows you to manually set the system clock. Using the **clock set** command starts the system clock with the time and date you specify. The time counter setting is retained across power cycles.

To set the system time and date to 10:15:05 on October 15, 2005, enter the following command.

```
device# clock set 10:15:05 10-15-05
```

Syntax: [no] **clock set** *hh:mm:ss mm-dd-yy | mm-dd-yyyy*

By default, the Extreme device does not change the system time for daylight savings time. To enable daylight savings time, enter the following command.

```
device# clock summer-time
```

Syntax: [no] **clock summer-time**

You can configure the Extreme device to adjust the time for any one-hour offset from GMT or for one of the following U.S. time zones:

- US Pacific (default)
- Alaska
- Aleutian
- Arizona
- Central
- East-Indiana
- Eastern
- Hawaii
- Michigan
- Mountain

- Pacific
- Samoa

The default is US Pacific.

To change the time zone to Australian East Coast time (which is normally 10 hours ahead of GMT), enter the following command.

```
device(config)# clock timezone gmt gmt+10
```

Syntax: **[no] clock timezone gmt** *time-zone* | **us** *time-zone*

You can enter one of the following values for *time-zone* :

- GMT time zones (**gmt**): gmt+12, gmt+11, gmt+10...gmt+01, gmt+00, gmt-01...gmt-10, gmt-11, gmt-12
- US time zones (**us**): alaska, aleutian, arizona, central, east-indiana, eastern, hawaii, michigan, mountain, pacific, samoa

DST "change" notice for networks using US time zones

The new Daylight Saving Time (DST) change that went into effect on March 11, 2007 affects networks in the US time zones. Because of this change, your network clock might not be correct. If your network uses US time zones, and it needs to maintain the correct time, you must enable the following command.

```
device(config)# clock timezone us pacific
```

Syntax: **[no] clock timezone us { pacific | eastern | central | mountain }**

NOTE

This command must be configured on every device that uses the US DST.

To verify the change, use the following command.

```
device(config)# show clock
```

Creating a command alias

Use the **alias** command to create an alias for a command and to save that alias within the device's configuration.

To create the alias "shro" for the **show ip routes** command, use the following command.

```
device(config)# alias shro = show ip routes
device(config)# write memory
```

Syntax: **[no] alias [name = command]**

The *name* variable is the name that you want to assign to the alias.

The *command* variable is the syntax for the command you want to create an alias for.

The **write memory** command is used to save the alias within the configuration.

Removing an alias

You can remove an alias using the **no** version of the alias command as shown in the following.

```
device(config)# no alias shro
```

Alternately, you can use the **unalias** command as shown in the following.

```
device(config)# unalias shro
```

Syntax: [no] unalias

If the alias you try to remove does not exist, the following error will be displayed.

```
device(config)# unalias wrs
Error: Alias wrs does not exist, unalias failed
```

Displaying a list of all configured alias

The following command allows you to display a list of all configured alias.

```
device# alias
#alias
savemem      write memory
shro         show ip routes
```

Syntax: [no] alias

Configuring CLI banners

The Extreme device can be configured to display a greeting message on users' terminals when they enter the Privileged EXEC CLI level or access the device through Telnet. In addition, a Extreme device can display a message on the Console when an incoming Telnet CLI session is detected.

Setting a message of the day banner

You can configure the Extreme device to display a message on a user's terminal when he or she establishes a Telnet CLI session. For example, to display the message "Welcome to Extreme!" when a Telnet CLI session is established, enter the following.

```
device(config)# banner motd $(Press Return)
Enter TEXT message, End with the character '$'.
Welcome to Extreme! $
```

A delimiting character is established on the first line of the **banner motd** command. You begin and end the message with this delimiting character. The delimiting character can be any character except "(double-quotation mark) and cannot appear in the banner text. In this example, the delimiting character is \$(dollar sign). The text in between the dollar signs is the contents of the banner. The banner text can be up to 2047 characters long and can consist of multiple lines. To remove the banner, enter the **no banner motd** command.

Syntax: [no] banner *delimiting-character* | [motd *delimiting-character*]

NOTE

The **banner *delimiting-character*** command is equivalent to the **banner motd *delimiting-character*** command.

NOTE

The size of the MOTD banner will be restricted (truncated) to 1850 characters when using an SSH client.

Setting a privileged EXEC CLI level banner

You can configure the Extreme device to display a message when a user enters the Privileged EXEC CLI level.

```
device(config)# banner exec_mode # (Press Return)
Enter TEXT message, End with the character '#'.
You are entering Privileged EXEC level
Don't foul anything up! #
```

As with the **banner motd** command, you begin and end the message with a delimiting character; in this example, the delimiting character is # (pound sign). To remove the banner, enter the **no banner exec_mode** command.

Syntax: **[no] banner exec_mode** *delimiting-character*

Displaying a message on the console when an incoming Telnet session is detected

You can configure the Extreme device to display a message on the Console when a user establishes a Telnet session. This message indicates where the user is connecting from and displays a configurable text message.

```
device(config)# banner incoming $(Press Return)
Enter TEXT message, End with the character '$'.
Incoming Telnet Session! $
```

When a user connects to the CLI using Telnet, the following message appears on the Console.

```
Telnet from 10.157.22.63
Incoming Telnet Session!
```

Syntax: **[no] banner incoming** *delimiting-character*

To remove the banner, enter the **no banner incoming** command.

Configuring terminal display

You can configure and display the number of lines displayed on a terminal screen during the current CLI session.

The **terminal length** command allows you to determine how many lines will be displayed on the screen during the current CLI session. This command is useful when reading multiple lines of displayed information, especially those that do not fit on one screen.

To specify the maximum number of lines displayed on one page, enter a command such as the following.

```
device# terminal length 15
```

Syntax: **[no] terminal length** *number-of-lines*

The *number-of-lines* parameter indicates the maximum number of lines that will be displayed on a full screen of text during the current session. If the displayed information requires more than one page, the terminal pauses. Pressing the space bar displays the next page.

The default for *number-of-lines* is 24. Entering a value of 0 prevents the terminal from pausing between multiple output pages:

Checking the length of terminal displays

The **show terminal** command specifies the number of lines that will be displayed on the screen as specified by the **terminal length** , **page display** , and **skip-page-display** commands. It also shows if the **enable skip-page-display** command has been configured. The **enable skip-page-display** command allows you to use the skip-page-display to disable the configured page-display settings.

```
device(config)# show terminal
Length: 24 lines
Page display mode (session): enabled
Page display mode (global): enabled
```

Syntax: show terminal

Enabling or disabling Layer 2 switching

By default, Extreme devices supports routing over Layer 2 switching. You can enable Layer 2 switching globally or on individual port using the **no route-only** command.

NOTE

On the CES 2000 Series and CER 2000 Series, the **route-only** command should not be configured on untagged MPLS uplinks when using it for VPLS or VLL. Otherwise, incoming VPLS or VLL traffic is dropped.

The **no route-only** and **route-only** commands prompts you for whether or not you want to change the "route-only" behavior. You must enter *y* if you want to proceed or *n* if you do not. The prompt is displayed as shown in the following examples of the **no route-only** and **route-only** commands.

NOTE

Always perform a reload after removing a route-only config or enabling route-only. Removing or enabling the route-only option without a reload will cause multicast issues.

To enable Layer 2 switching globally, enter the following.

```
device(config)# no route-only
This will change the route-only behavior at the global level.
Are you sure? (enter 'y' or 'n'): y
Global 'route-only' committed.
```

To globally disable Layer 2 switching on a Extreme device and return to the default (route-only) condition, enter commands such as the following:

```
device(config)# route-only
This will change the route-only behavior at the global level.
Are you sure? (enter 'y' or 'n'): y
Global 'no route-only' committed.
```

Syntax: [no] route-only

NOTE

On the XMR Series and MLX Series devices, **route-only** is the default condition. Because **route-only** is the default condition, it will not be displayed in the configuration. If you use **no route-only** to enable switching, the **no route-only** command will be displayed in the configuration.

NOTE

On the CES 2000 Series device, **route-only** is disabled by default. Therefore, if **route-only** is enabled on a CES 2000 Series device, it will be displayed in the configuration.

To enable Layer 2 switching only on a specific interface, go to the Interface configuration level for that interface, and add the **no route-only** command. The following commands show how to enable Layer 2 switching on port 3/2.

```
device(config)# interface ethernet 3/2
device(config-if-e10000-3/2)# no route-only
```

To re-enable the default **route-only** condition on port 3/2, enter the **route-only** command as shown.

```
device(config-if-e10000-3/2)# route-only
```

When **route-only** is enabled on a physical interface, incoming unknown unicast packets are not sent to the CPU and are dropped locally by the hardware.

NOTE

Configuring **route-only** on a physical interface affects incoming frames only. In other words, interface **route-only** disables Layer 2 switching for incoming frames but does not disable Layer 2 switching for outgoing frames. If the **route-only** interface is a member of a VLAN, the interface will still transmit frames received on other interfaces of that VLAN if those other interfaces still have Layer 2 switching enabled. To prevent this from happening, make sure that any interface you have configured for **route-only** are not also members of VLANs where you are intentionally performing Layer 2 switching.

Configuring static MAC addresses

You can assign static MAC addresses to ports of a Extreme device.

You can manually input the MAC address of a device to prevent it from being aged out of the system address table, to prevent traffic for a specific device, such as a server, from flooding the network with traffic when it is down, and to assign higher priorities to specific MAC addresses.

Static MAC addresses are configured within a specified VLAN including the default VLAN 1. Optionally you can specify a port priority (QoS).

The default and maximum configurable MAC table sizes can differ depending on the device. To determine the default and maximum MAC table sizes for your device, display the system parameter values. Refer to [Displaying and modifying default settings for system parameters](#) on page 103.

NOTE

The absolute maximum number of static MAC addresses is 400.

The ability of the CAM to store depends on the following:

- The number of source MAC address being learned by the CAM.
- The number of destination MAC addresses being forwarded by the CAM
- The distribution of the MAC address entries across ports. For example, if one port is learning all the source MAC addresses, the available of the CAM for that port will be depleted.

In the following example, a static MAC address of 0000.0063.67FF with a priority of 7 is assigned to port 2 of module 1 in VLAN 200.

```
device(config)# vlan 200
device(config)# static-mac-address 0000.0063.67FF e 1/2 priority 7
```

Syntax: [no] static-mac-address *mac-addr* ethernet *portnum* [*priority number*]

The *mac-addr* variable specifies the MAC address that you assigning.

The *portnum* variable specifies the Ethernet port that the MAC address is being assigned to.

Using the **priority** option, you can assign a value to the *number* variable of 0 - 7.

Disabling the MAC movement console logs

The ability to disable MAC movement syslog messages is useful to prevent logging messages appearing on the console display.

Syslog messages are generated when MAC addresses are changed and these messages display on the console port. In situations when there is a lot of MAC movement activity, you can disable these messages from the display using the following steps.

NOTE

This task is applicable to only to MLX Series devices.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Disable console logs when MAC movement syslog messages are generated.

```
device(config)# no mac-move-det-syslog
```

The ability to display MAC movement messages is enabled by default. Disabling the messages removes the **mac-move-det-syslog** command from the running configuration.

The following example shows the MAC movement syslog message output when the **mac-move-det-syslog** command is enabled. Note the syslog message about the MAC address movement.

```
device# configure terminal
device(config)# mac-move-det-syslog
device(config)# show arp

Total number of ARP entries: 2
(In all VRFs)

Entries in default routing instance:
IP   Address      MAC Address      Type      Age Port (Vpls-Id, Vlan)/ Vpls-Id:Peer
1    10.19.19.1     0010.9400.0606   Dynamic   1   1/24
2    172.26.67.1    0024.381c.b900   Dynamic   1   mgmt1
device(config)# exit
device#
SYSLOG: <12>Sep 25 02:43:07 IP/ARP: IP address 19.19.19.1 MAC movement detected,
        changed from MAC 0010.9400.0606 / port 1/24 to MAC 0010.9400.0001 / port 1/24
```

The following example disables the MAC movement syslog message output. No syslog messages are displayed.

```
device# configure terminal
device(config)# no mac-move-det-syslog
device(config)# exit
device# show arp
Total number of ARP entries: 2
(In all VRFs)

Entries in default routing instance:
IP   Address      MAC Address      Type      Age Port (Vpls-Id, Vlan)/ Vpls-Id:Peer
1    10.19.19.1     0010.9400.0001   Dynamic   1   1/24
2    172.26.67.1    0024.381c.b900   Dynamic   2   mgmt1
device#
device#
```

Changing the MAC age time

The MAC age time sets the aging period for ports on the device, defining how long (how many seconds) a port address remains active in the address table.

To change the aging period for MAC addresses from the default of 300 seconds to 600 seconds.

```
device(config)# mac-age-time 600
```

Syntax: [no] **mac-age-time** *age-time*

The *age-time* can be 0 or a number from 60 - 65535. The zero results in no address aging. The default is 300 (seconds).

Configuring system max values

Table 3 lists the system max values for the several system parameters of the Extreme devices.

TABLE 3 System max values for XMR Series and Extreme MLX Series devices

Parameter	Minimum value for MLX Series	Maximum value for MLX Series	Default value for MLX Series	Minimum value for XMR Series	Maximum value for XMR Series	Default value for XMR Series
config-file-size	2097152	16777216	8388608	2097152	16777216	8388608
gre-tunnels	1	8192	256	1	8192	256
hw-flooding	0	4095	0	0	4095	8
ifl-cam	0	81920	0	0	81920	0
ip-arp	2048	131072	8192	2048	131072	8192
ip-cache	8192	2621440	655360	8192	1048576	204800
ip-filter-system	1024	102400	4096	1024	102400	4096
ip-route	4096	2506752	655360	4096	1048576	204800
ipv4-mcast-cam	0	32768	4096	0	65536	8192
ip-subnet-port	24	128	24	24	128	24
ip-vrf-route	128	655360	5120	128	450560	5120
ipv6-cache	8192	1884160	131072	8192	245760	65536
ipv6-mcast-cam	0	8192	1024	0	16384	2048
ipv6-route	4096	1884160	131072	4096	245760	65536
ipv6-vrf-route	64	98304	1024	64	16384	128
l2-acl-table-entries	64	256	64	64	256	64
mac	4000	1048576	32768	4000	2097152	131072
mgmt-port-acl-size	1	100	20	1	100	20
subnet-broadcast-acl-cam	0	4096	0	0	4096	0
receive-cam	512	8192	1024	512	8192	1024
rstp	1	128	32	1	128	32
session-limit	1024	40960	8192	1024	163840	32768
spanning-tree	1	128	32	1	128	32
virtual-interface	40	4095	255	40	4095	255
vlan	2	4095	512	2	4095	512
vpls-mac	32	262144	2048	32	1000000	8192
vpls-num	512	4096	512	1024	16384	2048
ecmp-pram-block-size	8	32	32	8	32	32

NOTE

If Algorithmic mode is enabled, the system maximum values for ip-cache, ip-route, ipv6-cache, and ipv6-route is limited to the maximum value supported by the specific CAM profile.

TABLE 4 System max values for CES 2000 Series, CER 2000 Series, and CER 2000 Series-RT devices

Parameter	Minimum value for CES 2000 Series	Maximum value for CES 2000 Series	Default value for CES 2000 Series	Minimum value for CER 2000 Series	Maximum value for CER 2000 Series	Default value for CER 2000 Series	Minimum value for CER 2000 Series-RT	Maximum value for CER 2000 Series-RT	Default value for CER 2000 Series-RT
config-file-size	2097152	16777216	8388608	2097152	16777216	8388608	2097152	16777216	8388608
ip-arp	2048	16384	4096	2048	16384	4096	2048	16384	4096
ip-cache	4096	32768	16384	4096	524288	290816	4096	1572864	290816
ip-filter-sys	1024	32768	4096	1024	32768	4096	1024	32768	4096
ip-route	4096	32768	16384	4096	524288	290816	4096	1572864	290816
ip-subnet-port	24	128	24	24	128	24	24	128	24
l2-acl-table-entries	64	256	64	64	256	64	64	256	64
mac	4000	131072	56320	4000	131072	56320	4000	131072	65536
mgmt-port-acl-size	1	100	20	1	100	20	1	100	20
rstp	1	128	32	1	128	32	1	128	32
session-limit	1024	32768	32768	1024	32768	32768	1024	32768	32768
spanning-tree	1	128	32	1	128	32	1	128	32
virtual-interface	40	1024	255	40	4095	255	40	4095	255
vlan	512	4095	512	2	8192	512	2	8192	512
vrf	1	16	1	1	128	16	1	128	16
ipv6-cache	1024	131072	1024	1024	131072	1024	1024	262141	8192
ipv6-route	1024	131072	1024	1024	131072	8192	1024	262141	8192
vrf-route	1024	32768	1024	1024	32768	1024	1024	1572864	1024
ip-tunnels	32	128	32	32	128	32	32	256	32

NOTE

Default values are the same irrespective of the software package on the CES 2000 Series and CER 2000 Series devices.

NOTE

The maximum FIB scalability for CER 2000 Series and CES 2000 Series has been tested using an internet route mix. When using route prefixes concentrated in a narrow prefix length range, the scalability numbers will be lower. It is important to design your network keeping this in mind.

To configure system-max values, use the following command.

Syntax: [no] system-max config-file-size | gre-tunnels | ip-arp | ip-cache | ip-filter-sys | ip-route | ip-static-arp | ipv4-mcast-cam | ip-subnet-port | ip-tunnels | vrf | vrf-route | ipv6-cache | ipv6-mcast-cam | ipv6-route | l2-acl-table-entries | ifl-cam | mac | mgmt-port-acl-size | subnet-broadcast-acl-cam | receive-cam | rstp | session-limit | spanning-tree | trunk-num | virtual-interface | vlan | vpls-mac | vpls-num | ecmp-pram-block-size

The *gre-tunnels* parameter sets the maximum number of GRE tunnels. For minimum, maximum and default values for this parameter refer to [Table 3](#).

The *config-file-size* parameter sets the allowed running and startup-config file sizes. Refer to the appropriate table for your platform. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *ifl-cam* parameter sets the maximum number of Internal Forwarding Lookup Identifiers. These are used when configuring a Local VLL for Dual Tagging. The default value for the *ifl-cam* parameter is 8K. The maximum values for this parameter are different depending on which CAM partition you have configured on your system. For minimum, maximum and default values by CAM partition for this parameter, refer to [Table 5](#).

The *ip-arp* parameter sets the maximum number of ARP entries. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *ip-cache* parameter sets the maximum size of the IP cache. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *ip-filter-sys* parameter sets the maximum number of IP ACL entries. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *ip-route* parameter sets the maximum number of IP Route entries. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

NOTE

There is no need to configure a system-max value for static ARP entries.

The *ip-static-arp* parameter sets the maximum number of static ARP entries. For minimum, maximum and default values for this parameter refer to [Table 3](#).

The *ipv4-mcast-cam* parameter allows you to configure the maximum CAM size for an IPv4 multicast group. For minimum, maximum and default values for this parameter refer to [Table 4](#).

The *ip-subnet-port* parameter sets the maximum number of IP subnets per port. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *vrf-route* parameter sets the maximum number of VRF routes per VRF instance. For minimum, maximum and default values for this parameter refer to [Table 3](#).

The *ipv6-cache* parameter sets the maximum size of the IPv6 cache. For minimum, maximum and default values for this parameter refer to [Table 3](#).

The *ipv6-mcast-cam* parameter allows you to configure the maximum CAM size for an IPv6 multicast group. For minimum, maximum and default values for this parameter refer to [Displaying and modifying default settings for system parameters](#) on page 103.

The *ipv6-route* parameter sets the maximum number of IPv6 routes. For minimum, maximum and default values for this parameter refer to [Table 3](#).

NOTE

The **system-max ipv6-route** command can be configured with a maximum value of 114688 on the MLX Series, but the Extreme device system will only support a maximum value of 114687 for IPv6 routes.

The *l2-acl-table-entries* parameter sets the maximum number of layer-2 ACL entries per ACL table. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *mac* parameter sets the maximum number of MAC entries. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *mgmt-port-acl-size* parameter sets the maximum size for a management port ACL. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *subnet-broadcast-acl-cam* parameter sets the maximum number of IP broadcast ACL CAM entries. For minimum, maximum, and default values for this parameter, refer to [Table 3](#).

The *receive-cam* parameter sets the maximum number of IP Receive ACL software CAM entries. For minimum, maximum and default values for this parameter refer to [Table 3](#).

The *rstp* parameter sets the maximum number of RSTP instances. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *session-limit* parameter sets the maximum number of sessions. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *spanning-tree* parameter sets the maximum number of spanning-tree instances. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *virtual-interface* parameter sets the maximum number of virtual interfaces. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *vlan* parameter sets the maximum number of VLANs. For minimum, maximum and default values for this parameter refer to [Table 3](#) and [Table 4](#).

The *vpls-mac* parameter sets the maximum number of VPLS MAC Entries. For minimum, maximum and default values for this parameter refer to [Table 3](#).

The *vpls-num* parameter sets the maximum number of Virtual Private LAN Services (VPLS). For minimum, maximum and default values for this parameter refer to [Table 3](#).

The *ecmp-pram-block-size* parameter is used as a limiting factor when programming ECMP nexthops (v4/v6/vpnv4/vpnv6). Even though the control plane supports up to 32 nexthops per a route, the actual number of nexthops which are programmed in HW is controlled by this command. If *system-max ecmp-pram-block-size* is configured to a value lesser than the value configured for **ip load-sharing** or **ipv6 load-sharing**, or if *ip load-sharing* or *ipv6 load-sharing* is configured to a value greater than that configured for **system-max ecmp-pram-block-size**, a warning message will be displayed. For minimum, maximum and default values for this parameter, refer to [Table 3](#).

TABLE 5 System maxifl-cam values available by CAM profile on XMR Series and Extreme MLX Series

CAM profile	Minimum value	Maximum value	Default value
Default	0	57344	8192
ipv4	0	114688	8192
ipv6	0	131072	8192
l2-metro	0	114688	8192
mpls-l3vpn	0	114688	8192
mpls-vpls	0	114688	8192
multi-service	0	49152	8192
multi-service-2	0	81920	8192
vpn-vpls	0	114688	8192
ipv4-vpn	0	114688	8192
l2-metro-2	0	114688	8192
mpls-l3vpn-2	0	114688	8192
mpls-vpls-2	0	114688	8192
ipv4-ipv6	0	114688	8192
ipv4-vpls	0	114688	8192

TABLE 5 System maxifl-cam values available by CAM profile on XMR Series and Extreme MLX Series (continued)

CAM profile	Minimum value	Maximum value	Default value
ipv4-ipv6-2	0	81920	8192

Maintaining system-max configuration with available system resources

When system-max values are configured, the Extreme system checks for available system resources. The system resources are required in order to maintain dynamic memory allocation. System-max values are checked at the configuration time, and at the bootup time. If there are insufficient system resources available on the Management Module, this will cause the configuration to be rejected during card bootup. On the Interface Module, insufficient system resources will lead to failure in booting up the card.

Management VRF overview

The management VRF is used to provide secure management access to the device by sending inbound and outbound management traffic through the VRF specified as a global management VRF and through the out-of-band management port, thereby isolating management traffic from the network data traffic.

By default, the inbound traffic is unaware of VRF and allows incoming packets from any VRF, including the default VRF. The outbound traffic is only through the default VRF. The default VRF consists of out-of-band management port and all the LP ports that do not belong to any other VRFs.

Any VRF, except the default VRF, can be configured as a management VRF. When a management VRF is configured, the management traffic is allowed through the ports belonging to the specified VRF and the out-of-band management port. The management traffic through the ports belonging to the other VRFs and the default VRF are dropped and the rejection statistics are incremented.

If the management VRF is not configured, the management applications will follow the default behavior. The management VRF configuration is applicable for both IPv4 and IPv6 management traffic.

NOTE

The IPv6 management VRF is not supported on CES 2000 Series and CER 2000 Series devices.

The management VRF is supported by the following management applications:

- SNMP server
- SNMP trap generator
- Telnet server
- SSH server
- Telnet client
- RADIUS client
- TACACS+ client
- TFTP
- SCP
- Syslog

NOTE

The management VRF is not applicable to inbound and outbound traffic of the **ping** and **traceroute** commands. These commands use the VRF specified in the command or the default VRF, if no VRF is specified.

Source interface and management VRF compatibility

There is a source interface configuration associated with the management applications. When a source interface is configured, the management applications use the lowest configured IP address of the specified interface as source IP address in all the outgoing packets. If the configured interface is not part of the management VRF, the response packet will not reach the destination. If the compatibility check fails while configuring either the management VRF or the source interface, the following warning message will be displayed. However, the configuration command will be accepted.

```
The source-interface for Telnet, TFTP is not part of the management-vrf
```

Supported management applications

This section explains the management VRF support provided by the management applications.

NTP

NTP supports the management-VRF to isolate management traffic from network data traffic.

Enabling management-VRF support for NTP causes the incoming and outgoing traffic to travel through the management VRF or an out-of-bound (OOB) port. To support management-VRF support for NTP, all interfaces must be configured to be part of the global management VRF.

For more details, see the Management VRF for NTP section under the *NTP* chapter.

RADIUS client

When the management VRF is configured, the RADIUS client will send RADIUS requests or receive responses only through the ports belonging to the management VRF and through the out-of-band management port.

Any change in the management VRF configuration will be immediately effective for the RADIUS client.

NOTE

The RADIUS source interface configuration **ip radius source-interface** command must be compatible with the management VRF configuration. Refer to [Source interface and management VRF compatibility](#) on page 39.

SCP

SCP uses SSH as underlying transport. The behavior of SCP is similar to the SSH server. For more information, refer to [SSH server](#) on page 40.

SNMP server

When the management VRF is configured, the SNMP server receives SNMP requests and sends SNMP responses only through the ports belonging to the management VRF and through the out-of-band management port.

Any change in the management VRF configuration becomes immediately effective for the SNMP server.

SNMP trap generator

When the management VRF is configured, the SNMP trap generator sends traps to trap hosts through the ports belonging to the management VRF and through the out-of-band management port.

Any change in the management VRF configuration becomes immediately effective for the SNMP trap generator.

NOTE

The SNMP source interface configuration command **snmp-server trap-source** must be compatible with the management VRF configuration. Refer to [Source interface and management VRF compatibility](#) on page 39.

SSH server

When the management VRF is configured, the incoming SSH connection requests are allowed only from the ports belonging to the management VRF and from the out-of-band management port. Management VRF enforcement is only done during the establishment of a connection. Once the connection is established, no further management VRF enforcement is done.

To allow the incoming SSH connection requests only from the management VRF and not from the out-of-band management port, enter the following command.

```
device# configure terminal
device(config)# ip ssh strict-management-vrf
```

The previous command is applicable only when the management VRF is configured. If not, the command issues the following warning message.

```
Warning - Management-vrf is not configured.
```

For the SSH server, changes in the management VRF configuration or configuring the **ip ssh strict-management-vrf** command will not affect the existing SSH connections and the changes will be applied only to the new incoming connection requests.

By default, when the management VRF is configured, incoming SSH connection requests are only allowed from ports that belong to the management VRF and from the out-of-band management port (of the management VRF, default VRF or user-defined VRF); that is, incoming SSH connection requests from ports that belong to the default VRF or user-defined VRFs are rejected. When this is not acceptable; for example, in the situation where you want to configure the management VRF for RADIUS or SNMP and also want to access the device from a public network using SSH, use the **ip ssh include-all-vrf** command to allow incoming SSH connection requests from ports that belong to any VRF and from the out-of-band management port when the management VRF is configured.

```
device# configure terminal
device(config)# ip ssh include-all-vrf
```

Syslog

When the management VRF is configured, the Syslog module sends log messages only through the ports belonging to the management VRF and the out-of-band management port.

Any change in the management VRF configuration will be immediately effective for Syslog.

NOTE

The Syslog source interface configuration **ip syslog source-interface** command must be compatible with the management VRF configuration. Refer to [Source interface and management VRF compatibility](#) on page 39.

TACACS+ client

When the management VRF is configured, the TACACS+ client establishes connections with TACACS+ servers only through the ports belonging to the management VRF and the out-of-band management port.

For the TACACS+ client, any change in the management VRF configuration will not affect the existing TACACS+ connections and the changes will be applied only to the new TACACS+ connections.

NOTE

The TACACS+ source interface configuration **ip tacacs source-interface** command must be compatible with the management VRF configuration. Refer to [Source interface and management VRF compatibility](#) on page 39.

Telnet sessions

The existing number of inbound and outbound telnet session connections is limited to 5 per device.

With this enhancement, both inbound and outbound telnet sessions is increased to maximum 10 connections. An user can connect 10 telnet session connections on a single device.

Telnet client

When the VRF name is specified in the **telnet vrf** command, the Telnet client initiates Telnet requests only from the ports belonging to the specified VRF.

To configure the VRF name in outbound Telnet sessions, enter the following command at the privileged EXEC level:

```
device(config)# telnet vrf red 10.157.22.39
```

Syntax: **telnet vrf** *vrf-name* *IPv4address* | **ipv6** *IPv6address*

The *vrf-name* variable specifies the name of the pre-configured VRF.

NOTE

The IPv6 management VRF is not supported on CES 2000 Series and CER 2000 Series devices.

Telnet server

When the management VRF is configured, the incoming Telnet connection requests are allowed only from the ports belonging to the management VRF and from the out-of-band management port. Management VRF enforcement is only done during the establishment of a connection. Once the connection is established, no further management VRF enforcement is done.

To allow the incoming Telnet connection requests only from the management VRF and not from the out-of-band management port, enter the following command.

```
device(config)# telnet strict-management-vrf
```

The previous command is applicable only when the management VRF is configured. If not, the command issues the following warning message.

```
Warning - Management-vrf is not configured.
```

For the Telnet server, changes in the management VRF configuration or configuring the **telnet strict-management-vrf** command will not affect the existing Telnet connections and the changes will be applied only to the new incoming connection requests.

TFTP

When the management VRF is configured, TFTP will send or receive the data and acknowledgements only through the ports belonging to the management VRF and through the out-of-band management port.

Any change in the management VRF configuration will be immediately effective for TFTP. You cannot change in the management VRF configuration while TFTP is in progress.

NOTE

The TFTP source interface configuration **ip tftp source-interface** command must be compatible with the management VRF configuration. Refer to [Source interface and management VRF compatibility](#) on page 39.

Configuring a global management VRF

To configure a VRF as a global management VRF, enter the following command.

```
device(config)# management-vrf mvrf
```

Syntax: [no] **management-vrf** *vrf-name*

The *vrf-name* parameter specifies the name of the pre-configured VRF. If the VRF is not pre-configured, the command execution fails and displays the following error message.

```
Error - VRF <vrf-name>
doesn't exist
```

When the management VRF is configured, the software generates the following Syslog message.

```
SYSLOG: VRF <vrf-name>
has been configured as management-vrf
```

Enter the **no** form of the command to remove the management VRF. When the management VRF is deleted, the software generates the following Syslog message.

```
SYSLOG: VRF <vrf-name>
has been un-configured as management-vrf
```

Configuration notes

Consider the following configuration notes:

- If there is a management VRF already configured, you must remove the existing management VRF configuration before configuring a new one. If not, the system displays the following error message.

```
device(config)# management-vrf red
Error - VRF mvrf already configured as management-vrf
```

- If you try to delete a management VRF that was not configured, the system displays the following error message.

```
device(config)# no management-vrf red
Error - VRF red is not the current management-vrf
```

- The deletion or modification of the VRF will fail if the specified VRF is currently configured as the management VRF. Attempting to do so causes the system to return the following error message.

```
device(config)# no vrf mvrf
Error - Cannot modify/delete a VRF which is configured as management-vrf
```

Displaying the management VRF information

To display IP Information for a specified VRF, enter the following command at any level of the CLI.

```
device(config)# show vrf mvrf
Total number of VRFs configured: 1
Status Codes - A:active, D:pending deletion, I:inactive
Name           Default RD      IFL ID  vrf|v4|v6      Routes Interfaces
a              1:1            131071  A | A | A      14
```

```

Total number of IPv4 unicast route for all non-default VRF is 12
Total number of IPv6 unicast route for all non-default VRF is 2
device#show vrf a
VRF a, default RD 1:1, Table ID 1 IFL ID 131071
Label: (Not Allocated), Label-Switched Mode: OFF
Configured as management-vrf
IP Router-Id: 10.2.2.2
No interfaces
  No Export VPN route-target communities
  No Import VPN route-target communities
  No import route-map
  No export route-map
Address Family IPv4
  Max Routes: 5120
  Number of Unicast Routes: 12
  No Export VPN route-target communities
  No Import VPN route-target communities
Address Family IPv6
  Max Routes: 128
  Number of Unicast Routes: 2
  No Export VPN route-target communities
  No Import VPN route-target communities

```

Syntax: `show vrf vrf-name`

The *vrf-name* parameter specifies the VRF for which you want to display IP information.

Table 6 displays a description of the output from the **show vrf** command.

TABLE 6 Output from the **show vrf** command

Field	Description
VRF <i>vrf-name</i>	The name of the VRF.
default RD	The default route distinguisher for the VRF.
Table ID	The table ID for the VRF.
Routes	The total number of IPv4 and IPv6 Unicast routes configured on this VRF.
IFL ID	The Internal Forwarding Lookup Identifier (IFL-ID) for ports in the VRF instance.
Label	The unique VRF label that has been assigned to the specified VRF.
Label-Switched Mode	Indicates whether Label-Switched Mode is ON or OFF.
Configured as management-vrf	Indicates that the specified VRF is configured as a management VRF.
IP Router-Id	The 32-bit number that uniquely identifies the router.
Number of Unicast Routes	The number of Unicast routes configured on this VRF.
import route-map	The name of the import route-map, if any, that is configured for this management VRF.
export route-map	The name of the export route-map if a route-map has been configured for this management VRF.

The **show who** command displays information about the management VRF from which the Telnet and SSH connection has been established.

```

device(config)#show who
Console connections:
  established, monitor enabled, privilege super-user
  10 days 22 hours 46 minutes 30 seconds in idle
Telnet server status: Enabled
Telnet copy-received-cos status: Disabled
Telnet connections (inbound):
  1    established, client ip address 134.141.186.58, privilege super-user
      using vrf default-vrf.
      you are connecting to this session
      2 minutes 54 seconds in idle
  2    closed
  3    closed

```

```

4      closed
5      closed
6      closed
7      closed
8      closed
9      closed
10     closed
Telnet connections (outbound):
11     closed
12     closed
13     closed
14     closed
15     closed
16     closed
17     closed
18     closed
19     closed
20     closed
SSH server status: Enabled
SSH copy-received-cos status: Disabled
SSH connections (inbound):
1      closed
2      closed
3      closed
4      closed
5      closed
6      closed
7      closed
8      closed
9      closed
10     closed
11     closed
12     closed
13     closed
14     closed
15     closed
16     closed
SSH connections (outbound):
17     closed

```

Syntax: show who

To display the packets and sessions rejection statistics due to failure in management VRF validation, enter the following command.

```

device(config)# show management-vrf

Management VRF name : ay
Management Application      Rx Drop Pkts      Tx Drop Pkts
SNMP Engine                 0                  0
RADIUS Client                0                  0
TFTP Client                  0                  0
Traps                        -                  0
SysLogs                     -                  0
NTP Server                   0                  0
NTP Client                   0                  0

TCP Connection rejects:
Telnet      :          0
SSH         :          0
TACACS+ Client :        0

```

Syntax: show management-vrf

Table 7 displays a description of the output from the **show management-vrf** command.

TABLE 7 Output from the **show management-vrf** command

Field	Description
Management VRF name	Displays the configured management VRF name.

TABLE 7 Output from the **show management-vrf** command (continued)

Field	Description
Management Application	Displays the management application names.
Rx Drop Pkts	Displays the number of packets dropped in the inbound traffic.
Tx Drop Pkts	Displays the number of packets dropped in the outbound traffic.
TCP Connection rejects	Displays the number of TCP connections per application rejected due to management VRF validation.

Make sure that the management VRF is configured before executing the **show management-vrf** command. If not, the system will display the following error message.

```
Error - Management VRF is not configured.
```

To clear the management VRF rejection statistics, enter the following command.

```
device(config)# clear management-vrf-stats
```

Syntax: clear management-vrf-stats

Bootup time

At bootup time, the Management Module will repeat the same process as done in the Configuration time. The Management Module calculates the memory required to accept the system-max configuration. The resulting value is checked against the Known-Available-Memory value for both the Management Module and the Interface Module.

After the new system-max value is configured, there are three possible configuration outcomes. The three possible configuration outcomes are described below.

1. The configuration can be accommodated, but leaves only 10% of Available Memory

In this configuration, a check is made against 90% of Available Memory. If the difference between the Required Available Memory and the Available Memory is less than 10% of Available Memory, then the configuration is accepted. The following warning message is displayed on the console if it affects the Management Module or Interface Module.

The following warning message is displayed on the Management Module.

```
WARNING: Configured System-max Leaves less than 10 percent Free Memory Available on MP.
```

The following warning message is displayed on the Interface Module.

```
WARNING: Configured System-max Leaves less than 10 percent Free Memory Available on LP.
```

A syslog message showing the required memory versus the available memory is generated, and a similar warning message is displayed on the Management Module and Interface Module as shown in the following example.

```
device# show log
...
Jan 17 22:55:55:N: WARN: Configured System-max Leaves less than 10 percent Free Memory Available on
MP (162529285 req vs 1625292800 available)
Jan 17 22:55:55:N: WARN: Configured System-max Leaves less than 10 percent Free Memory Available on
LP (162529285 req vs 1625292800 available)
```

NOTE

When the system is booted up again, the percent of free memory is discretionary and is only an estimate.

NOTE

Even if all elements are configured with the maximum allowed value, you may not see the reversion of system-max values that occur on any given Interface Module.

NOTE

Notifications and traps are sent with the same message.

2. The configuration can be easily accommodated.

In this configuration, the Management Module continues to use the configured system-max value, and send the same value to the installed Interface Modules.

3. The configuration cannot be accommodated.

If the configured system-max value cannot be used, the Management Module will locate the elements that can be reverted to a default value. These system-max elements will revert to a default value, and the following message will display on the console.

```
WARN: Configured System-max cannot be accommodated. Resetting revertible elements to default values.
```

A syslog message is generated, and a similar warning message is displayed on the Management Module and Interface Module as shown in the following example.

```
device# show log
...
Jan 17 22:55:55:N: WARN: Configured System-max cannot be accommodated on MP
(1625292801 req vs 1625292800 available). Resetting revertible elements to default
values.
Jan 17 22:55:55:N: WARN: Configured System-max cannot be accommodated on LP
(1625292801 req vs 1625292800 available). Resetting revertible elements to default
values.
```

NOTE

Once the system-max have been reverted, a user might not be able to configure any system-max until configuration for some or all of the revertible system-max elements is removed using "no system-max" CLI.

NOTE

Notifications and traps are sent with the same message.

The following tables show which elements are revertible (Yes or No) in each element category.

L2 elements

TABLE 8 L2 elements

L2 elements	Revertible: yes or no
Mac	yes
Vlan	no
Spanning-tree	no
Rstp	no

L3 elements

TABLE 9 L3 elements

L3 elements	Revertible: yes or no
Arp	no
multicast-route (for v6 only)	yes
pim-mcache	yes
ip-cache	yes
ip-route	yes
ip-subnet-port	no
virtual-interface	no

VPLS elements

TABLE 10 VPLS elements

VPLS elements	Revertible: yes or no
vpls-mac (MAX_VPLS_MAC_INDEX)	yes
vpls-num (MAX_VPLS_NUM_INDEX)	no

Miscellaneous elements

TABLE 11 Miscellaneous elements

Miscellaneous elements	Revertible: yes or no
session-limit	yes
ip-filter-sys	no
mgmt-port-acl-size	no
l2-acl-table-entries	no
ipv6-cache	yes
ipv6-route	yes
IPVRF MAX ROUTES	yes
mgmt-port-acl-size	no
receive-cam	no
IPGRE	no
LSP_ACL	no
SERVICE_LOOKUP	no
IP_SRC_GUARD_CAM	no
IPv4 MCAST CAM	no
IPv6 MCAST CAM	no
SERVER_TRUNKS	no
CONFIG_FILE_SIZE	no

Bootup time message

At bootup time, the following warning and error message is displayed in bold. The warning message and the error message are intermittent. The warning message indicates that when the standby management module (MP) comes up, the active MP syncs the FID entries to the standby MP. If the FID sync fails, the standby MP reboots. The FID sync can fail with a timeout error message if the standby MP is busy processing, and the MP does not respond within an agreeable timeout period and retries. The FIDs are synced successfully on the subsequent reboot. The second error message indicates that when the system is rebooting the standby MP, the baseline sync cannot be completed for multicast. The baseline sync for multicast is aborted. The required state information is not replicated completely from the active MP to the standby MP when the system is rebooting.

```
device>All tasks have completed their initializations
```

```
Start code flash synchronization to standby MP.
```

```
Code flash synchronization to standby MP is done.
```

```
Standby is syncing to Active. Please do not enter anything until Sync complete message is received.
```

```
Start running config synchronization to standby MP.
```

```
Running config synchronization to standby MP is done.
```

```
Warn:alloc_and_distribute_base_fid: Sync to standby MP failed for FID 120 (0078) (err = Timeout),reboot  
it(g_mp_red_wait_done 0) <<<<<<<
```



```
Reset Standby MP
Module is up in slot 3
Module is up in slot 6
Error:process_baseline_sync_status: component id 8 is not in baseline sync <<<<<<<<<<<<<<<<
INFO: Back fan A-1 status is OK now.
INFO: Back fan A-2 status is OK now.
INFO: Back fan B-1 status is OK now.
INFO: Back fan B-2 status is OK now.
Module is up in slot 1
Module is up in slot 5
Module is up in slot 4
Module is up in slot 2
Start code flash synchronization to standby MP.
Code flash synchronization to standby MP is
done.
```

Configuration time

When system-max values are configured, the Management Module calculates the memory required to accept the value. The resulting value is checked against the Known-Available-Memory value, and calculated against the Highest Required Memory value for both the Management Module and the Interface Module.

The Known-Available-Memory is a value with the Lowest Supported Available Memory on a node. For example, if a node can accept a 1 Gigabyte LP, and a 512 MB LP, then the 512 MB LP will be used. The Highest Required Memory is a value with most amount of memory available on a node. For example, if a node has both 2 PPCR LP, and 1 PPCR LP, then the 2 PPCR LP will be used.

If the new system-max value is accepted, then the configuration will also be accepted. The following information will display.

```
device(config)#system-max mac 4000
Reload required. Please write memory and then reload or power cycle.
Failure to reload could cause system instability on failover.
Newly configured system-max will not take effect during hitless-reload.
```

NOTE

You must write this command to memory and perform a system reload for this command to take effect.

If the new system-max value is not accepted, then the configuration is rejected. The following error message is printed on the console.

```
device(config)# system-max ipv4 10000
ERROR: Configured System-max value cannot be accommodated.
```

Monitoring dynamic memory allocation

After a configured system-max value is accepted, it is possible that the dynamic memory allocation may fail in a running system. To monitor the amount of available memory on the Management Module and the Interface Module, a timer will check the memory every 10 seconds. If the available memory falls below 5 percent of the total installed memory, the timer will log the following warning message.

```
device# show log
...
Jan 17 22:55:55:N: WARN: Current Total Free Memory on MP is below 5 percent of Installed Memory.
...
Jan 17 23:53:55:N: WARN: Current Total Free Memory on LP 8 is below 5 percent of Installed Memory.
```

The warning message is displayed at a frequency of 1 log per 5 minutes.

NOTE

Notifications and traps are sent.

When the memory allocation fails, an alert message is logged immediately. The alert message is displayed at a frequency of 1 log per 5 minutes. The following example displays an alert message on the Management Module and the Interface Module.

```
device# show log
...
Jan 17 22:55:55:A: ALERT: Failed to allocate memory on MP
...
Jan 17 23:52:55:A: ALERT: Failed to allocate memory on LP 8
...
```

The NULL value is returned to the calling routine. The calling routine will decide how to proceed after the memory allocation fails.

NOTE

Notifications and traps are sent.

At any time, you can display the status of all recorded memory that is available on the Management Module by entering the **show memory** command. The amount of available memory is displayed in percentage values. The following example displays a show memory output on a Management Module.

```
device#show memory
=====
NetIron XMR active MP slot 33:
Total SDRAM      : 2147483648 bytes
Available Memory  : 1774059520 bytes
Available Memory (%): 82 percent
Free Physical Pages : 428503 pages
<...>
=====
NetIron XMR LP SL 2:
Total SDRAM      : 536870912 bytes
Available Memory  : 45821952 bytes
Available Memory (%): 8 percent
```

Commands that require a reload

Most CLI commands take effect as soon as you enter them. However, a small number of commands require a software reload to take effect. The following table lists some of these commands and more details about these commands is available in the NetIron Command Reference.

To place a configuration change made by one of these commands into effect, you must save the change to the startup-config file, then reload the software. If you reload the software without saving the change to the startup-config file, the device does not make the change.

To reload the software, you must perform a warm start. To perform a warm start, do one of the following:

- Enter the **reload** command at the Privileged EXEC level of the CLI.
- Enter the **boot system** command at the Privileged EXEC level of the CLI.

The following commands require a software reload.

- **cam-mode ip**
- **cam-mode ipvpn**
- **default-max-frame-size**
- **multicast-flooding**

- **system-max**
- **virtual-interface-mac**
- **vll-mtu-enforcement**

Verifying an image checksum

Use the **image-checksum** command to verify the checksum of the application, boot, or monitor images that are saved in code flash and Auxiliary Flash cards.

NOTE

The **image-checksum** command on is not applicable to a combined application image.

To check a monitor image, use the following command.

```
device# image-checksum monitor
OK
```

Syntax: **[no] image-checksum** *file-name*

The *file-name* variable specifies the image file that you want to verify the checksum for.

The following output can be generated by this command

TABLE 12 Output from the image-checksum command

Output	Description
File not found	The device failed to locate the specified file.
Failed to read file	The device failed to obtain the file length from the file system.
Not an image file	The specified file is not an image file.
File read failed	The specified file's actual length is different from the file length stored in the file system.
Checksum failed	The image has a checksum error.
OK	The checksum has been verified for the specified image file.

Configuring CAM mode globally

NOTE

There is no configuration support for **cam-mode** in NetIron CES and NetIron CER. Foundry Direct Routing (FDR) is enabled by default.

The default CAM mode currently supported in static CAM mode, also known as FDR. You can set the CAM mode to dynamic IP CAM using the following command:

```
device(config)# cam-mode ip dynamic
```

You must reload the device for this command to take effect.

Syntax: **[no] cam-mode ip [dynamic | static]**

The **dynamic** parameter sets the IP CAM mode to dynamic.

The **static** parameter sets the IP CAM mode to static. This is the default state.

Configuring density mode for the 2x100G and 20x10G CAM

Setting the CAM to double density mode will automatically disable uRPF. The uRPF is allowed under single density mode. The default setting for the XMR is double density. The default setting for the MLX is single density. You can set the density mode using the following command:

```
device(config)# cam-mode ip urpf-100g
```

You must reload the device for this command to take effect.

NOTE

There is no configuration support for **cam-mode** in NetIron CES and NetIron CER.

From 05.8.00a release onwards, you can achieve -X2 CAM profile numbers by enabling Algorithmic mode using the **cam-mode amod** command (available only on BR-MLX-100Gx2-CFP2-X2, BR-MLX-10Gx20-X2, or BR-MLX-1GX20-U10G-X2 cards). In Algorithmic mode, the line card runs in single density mode and also supports uRPF mode to work without reducing the route scale.

NOTE

By default, BR-MLX-100Gx2-CFP2-X2, BR-MLX-10Gx20-X2, or BR-MLX-1GX20-U10G-X2 cards boot up with -M CAM profile numbers and if uRPF is enabled, the number of routes are reduced by half.

For more information about CAM profile support and uRPF impact with and without Algorithmic mode, refer to [Configuring -X2 Algorithmic CAM profiles](#) on page 53.

Syntax: [no] **cam-mode ip** [**urpf** | **urpf-100g**]

The **urpf** parameter sets the IP CAM partition to single density mode.

The **urpf-100g** parameter sets the IP CAM partition to double density mode.

Configuring IPv6 host CAM mode

NOTE

There is no configuration support for **cam-mode** in NetIron CES and NetIron CER.

The CAM mode for IPv6 routes can be configured to host. You can set the CAM mode to **host** by using the following command.

```
device(config)# cam-mode ipv6 host
```

You must reload the device for this command to take effect.

Syntax: [no] **cam-mode ipv6 host**

The **host** parameter programs the complete 128 bit IPv6 address into the CAM.

Configuring IPv6 host drop CAM limit

To limit the usage of CAM by IPV6 hosts with unresolved ND, enter the **ipv6 max-host-drop-cam** command.

```
device(config)# ipv6 max-host-drop-cam 5
```

Syntax: [no] **ipv6 max-host-drop-cam** [*limit*]

The optional *limit* variable is the IPv6 drop CAM limit for a port per packet processor (PPCR). The limit value can be from 0 through 65535.

Configuring -X2 Algorithmic CAM profiles

The CAM profile numbers for BR-MLX-100Gx2-CFP2-X2, BR-MLX-10Gx20-X2, and BR-MLX-1GX20-U10G-X2 cards are much higher than the actual physical CAM capacity. Algorithmic mode, which employs a prefix-based software algorithm to optimize the CAM space, can be used to accommodate large numbers of longest prefix match (LPM) entries. Algorithmic mode also ensures reduced power consumption.

Depending on the -X2 CAM partition profile configuration, Algorithmic mode supports up to a maximum of 2448K IPv4, 1840K IPv6, 2048K IPv4 VPN, and 400K IPv6 VPN CAM entries.

From 05.8.00a release onwards, Algorithmic mode is available (disabled by default) on BR-MLX-100Gx2-CFP2-X2, BR-MLX-10Gx20-X2, and BR-MLX-1GX20-U10G-X2 cards.

NOTE

Extreme recommends enabling Algorithmic mode only on BR-MLX-100Gx2-CFP2-X2, BR-MLX-10Gx20-X2, or BR-MLX-1GX20-U10G-X2 cards as the cards come with the required factory-installed license. There is no license enforcement to enable Algorithmic mode on BR-MLX-100Gx2-CFP2-M, BR-MLX-10Gx20-M, or BR-MLX-1GX20-U10G-M cards.

By default, BR-MLX-100Gx2-CFP2-X2, BR-MLX-10Gx20-X2, and BR-MLX-1GX20-U10G-X2 cards boot up with -M CAM profile numbers and if uRPF is enabled, the number of routes are reduced by half. You must enable Algorithmic mode using the **cam-mode amod** command to achieve -X2 CAM profile numbers. Algorithmic mode also supports uRPF mode to work without reducing the route scale.

The line card must be reloaded for Algorithmic mode to take effect.

TABLE 13 CAM profile support and uRPF impact with and without Algorithmic mode

Modules	CAM profile support	uRPF impact
BR-MLX-100Gx2-CFP2-M, BR-MLX-10Gx20-M, and BR-MLX-1GX20-U10G-M	<ul style="list-style-type: none"> Non-Algorithmic mode (default): -M CAM profile type Algorithmic mode: Not supported 	The number of routes are reduced by half if uRPF mode is enabled.
BR-MLX-100Gx2-CFP2-X2, BR-MLX-10Gx20-X2, and BR-MLX-1GX20-U10G-X2	<ul style="list-style-type: none"> Non-Algorithmic mode (default): -M CAM profile type Algorithmic mode: -X2 CAM profile type (supported from 05.8.00a release onwards) 	<ul style="list-style-type: none"> Non-Algorithmic mode: The number of routes are reduced by half if uRPF mode is enabled. Algorithmic mode: No change in route scale even if uRPF is enabled.

The system-max values for ip-cache, ip-route, ipv6-cache, and ipv6-route is limited to the maximum value supported by the specific CAM profile. For example, for the multi-service-3 CAM profile, ip-cache system-max is limited to IPv4 + IPv4 VPN size of the profile which is 768k + 608k.

Refer to [Table 19](#) on page 62 for information about -X2 CAM profile partitions.

If the system has a mix of BR-MLX-100Gx2-CFP2-M or BR-MLX-10Gx20-M and BR-MLX-100Gx2-CFP2-X2 or BR-MLX-10Gx20-X2 cards, Algorithmic mode can be enabled on the BR-MLX-100Gx2-CFP2-X2 or BR-MLX-10Gx20-X2 cards by specifying the slot.

To enable Algorithmic mode, enter the following command:

```
device(config)# cam-mode amod slot 2
```

CAM partition profiles

CAM is partitioned on the device by a variety of profiles that you can select depending on your application. The available profiles are described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module).

To implement a CAM partition profile, enter the following command.

```
device(config)# cam-partition profile ipv4
```

NOTE

You must reload your device for this command to take effect.

Syntax: `cam-partition profile [ipv4 | ipv4-extended | ipv4-ipv6 | ipv4-ipv6-2 | ipv4-vpls | ipv4-vpn | ipv6 | l2-metro | l2-metro-2 | mpls-l3vpn | mpls-l3vpn-2 | mpls-vpls | mpls-vpls-2 | mpls-vpn-vpls | multi-service | multi-service-2 | multi-service-3 | multi-service-4 | multi-service-7]`

The **ipv4** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module), to optimize the device for IPv4 applications.

The **ipv4-extended** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module), to optimize the device for IPv4 applications with increased IPv4 Routes.

The **ipv4-ipv6** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module), to optimize the device for IPv4 and IPv6 dual stack applications.

The **ipv4-ipv6-2** option that was introduced in NetIron 03.7.00 adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module), to optimize the device for increased IPv4 routes with room for IPv6.

The **ipv4-vpls** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module), to optimize the device for IPv4 and MPLS VPLS applications.

The **ipv4-vpn** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module), to optimize the device for IPv4 and MPLS Layer 3 VPN applications.

The **ipv6** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module), to optimize the device for IPv6 applications.

The **l2-metro** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers, to optimize the device for Layer 2 Metro applications.

The **l2-metro-2** option provides another alternative to **l2-metro** to optimize the device for Layer 2 Metro applications. It adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers.

The **mpls-l3vpn** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers, to optimize the device for Layer 3, BGP, or MPLS VPN applications.

The **mpls-l3vpn-2** option provides another alternative to **mpls-l3vpn** to optimize the device for Layer 3, BGP, or MPLS VPN applications. It adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers.

The **mpls-vpls** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers, to optimize the device for MPLS VPLS applications.

The **mpls-vpls-2** option provides another alternative to **mpls-vpls** to optimize the device for MPLS VPLS applications. It adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers.

The **mpls-vpn-vpls** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers, to optimize the device for MPLS Layer 3 and Layer 2 VPN applications.

The **multi-service** option adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers, to optimize the device for Multi-Service applications.

The **multi-service-2** option provides another alternative to **multi-service** to optimize the device for Multi-Service applications. It adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers.

The **multi-service-3** option provides another alternative to **multi-service** to optimize the device for Multi-Service applications to support IPv6 VRF. It adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers.

The **multi-service-4** option provides another alternative to **multi-service** to optimize the device for Multi-Service applications to support IPv6 VRF. It adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers.

The **multi-service-7** option provides another alternative to **multi-service** to optimize the device for Multi-Service applications to support IPv4 VRF. It adjusts the CAM partitions, as described in [Table 14](#) for the MLX Series (MR2-X management module) and [Table 15](#) for the MLX Series (MR2-M management module) routers.

There are twenty CAM partitioning profiles for the XMR Series and for the MLX Series routers. The profiles for routers are described in MLX Series (MR2-X management module) [Table 14](#) and the profiles for MLX Series (MR2-M management module) routers are described in [Table 15](#).

TABLE 14 CAM partitioning profiles available for MLX Series (MR2-X management module) routers

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	Logical size: 512K	Logical size: 64K	Logical size: 128K	Logical size: 128K	0	Logical size: 48K	Logical size: 4K	Logical size: 48K	Logical size: 4K
ipv4 Profile	Logical size: 1M	0	Logical size: 32K	0	0	Logical size: 112K	0	Logical size: 64K	0
ipv4-extended Profile	512K	64K	128K	128K	0	48K	4K	48K	4K
ipv6 Profile	Logical size: 64K	Logical size: 240K	Logical size: 32K	0	0	Logical size: 16K	Logical size: 24K	Logical size: 16K	Logical size: 12K
I2-metro Profile	Logical size: 256K	0	Logical size: 512K	0	0	Logical size: 64K	0	Logical size: 64K	0
mpls-l3vpn Profile	Logical size: 256K	0	Logical size: 32K	Logical size: 480K	0	Logical size: 64K	0	Logical size: 64K	0
mpls-vpls Profile	Logical size: 256K	0	Logical size: 512K	0	0	Logical size: 64K	0	Logical size: 64K	0
multi-service Profile	Logical size: 256K	Logical size: 32K	Logical size: 192K	Logical size: 256K	0	Logical size: 32K	Logical size: 8K	Logical size: 32K	Logical size: 8K

TABLE 14 CAM partitioning profiles available for MLX Series (MR2-X management module) routers (continued)

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
multi-service-2 Profile	512K	64K	128K	128K	0	48K	4K	48K	4K
multi-service-3 Profile	256K	32K	128K	192K	32K	32K	8K	32K	8K
multi-service-4 Profile	768K	32K	64K	64K	8K	32K	8K	48K	4K
multi-service-7 Profile	768K	32K	64K	64K	8K	32K	8K	48K	4K
mpls-vpn-vpls Profile	Logical size: 128K	0	Logical size: 224K	Logical size: 384K	0	Logical size: 48K	0	Logical size: 64K	0
ipv4-vpn Profile	Logical size: 320K	0	Logical size: 32K	Logical size: 448K	0	Logical size: 64K	0	Logical size: 64K	0
l2-metro-2 Profile	Logical size: 64K	0	Logical size: 608K	0	0	Logical size: 64K	0	Logical size: 64K	0
mpls-l3vpn-2 Profile	Logical size: 128K	0	Logical size: 32K	Logical size: 544K	0	Logical size: 64K	0	Logical size: 64K	0
mpls-vpls-2 Profile	Logical size: 128K	0	Logical size: 576K	0	0	Logical size: 64K	0	Logical size: 64K	0
ipv4-ipv6 Profile	Logical size: 320K	Logical size: 160K	Logical size: 32K	0	0	Logical size: 48K	Logical size: 20K	Logical size: 32K	Logical size: 8K
ipv4-vpls Profile	Logical size: 320K	0	Logical size: 480K	0	0	Logical size: 64K	0	Logical size: 64K	0
ipv4-ipv6-2 Profile	Logical size: 768K	Logical size: 64K	Logical size: 64K	0	0	Logical size: 64K	Logical size: 8K	Logical size: 48K	Logical size: 4K

TABLE 15 CAM partitioning profiles available for MLX Series (MR2-M management module) routers

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	Logical size: 256K	Logical size: 32K	Logical size: 64K	Logical size: 64K	0	Logical size: 24K	Logical size: 2K	Logical size: 48K	Logical size: 4K
ipv4 Profile	Logical size: 512K	0	Logical size: 16K	0	0	Logical size: 56K	0	Logical size: 64K	0
ipv4-extended Profile	256K	32K	64K	64K	0	24K	2K	48K	4K
ipv6 Profile	Logical size: 64K	Logical size: 112K	Logical size: 16K	0	0	Logical size: 8K	Logical size: 12K	Logical size: 16K	Logical size: 12K
l2-metro Profile	Logical size: 128K	0	Logical size: 256K	0	0	Logical size: 32K	0	Logical size: 64K	0
mpls-l3vpn Profile	Logical size: 128K	0	Logical size: 16K	Logical size: 240K	0	Logical size: 32K	0	Logical size: 64K	0

TABLE 15 CAM partitioning profiles available for MLX Series (MR2-M management module) routers (continued)

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
mpls-vpls Profile	Logical size: 128K	0	Logical size: 256K	0	0	Logical size: 32K	0	Logical size: 64K	0
multi-service Profile	Logical size: 128K	Logical size: 16K	Logical size: 96K	Logical size: 128K	0	Logical size: 16K	Logical size: 4K	Logical size: 32K	Logical size: 8K
multi-service-2 Profile	448K	16K	32K	32K	0	24K	2K	48K	4K
multi-service-3 Profile	128K	16K	64K	96K	32K	16K	4K	32K	8K
multi-service-4 Profile	448K	8K	0	32K	8K	16K	4K	48K	4K
multi-service-7 Profile	448K	8K	0	32K	8K	16K	4K	48K	4K
mpls-vpn-vpls Profile	Logical size: 64K	0	Logical size: 112K	Logical size: 192K	0	Logical size: 24K	0	Logical size: 64K	0
ipv4-vpn Profile	Logical size: 160K	0	16K	Logical size: 224K	0	Logical size: 32K	0	Logical size: 64K	0
l2-metro-2 Profile	Logical size: 64K	0	Logical size: 288K	0	0	Logical size: 32K	0	Logical size: 64K	0
mpls-l3vpn-2 Profile	Logical size: 64K	0	Logical size: 16K	Logical size: 272K	0	Logical size: 32K	0	Logical size: 64K	0
mpls-vpls-2 Profile	Logical size: 64K	0	Logical size: 288K	0	0	Logical size: 32K	0	Logical size: 64K	0
ipv4-ipv6 Profile	Logical size: 160K	Logical size: 80K	Logical size: 16K	0	0	Logical size: 24K	Logical size: 10K	Logical size: 32K	Logical size: 8K
ipv4-vpls Profile	Logical size: 160K	0	Logical size: 240K	0	0	Logical size: 32K	0	Logical size: 64K	0
ipv4-ipv6-2 Profile	480K	8K	32K	0	0	32K	4K	48K	4K

TABLE 16 CAM partitioning profiles available for the BR-MLX-100Gx2-X modules disabled/enabled

Profile	IPv4 Note: All IPv4 values listed are when operating in double density mode.	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	L2 Inbound ACL	IPv4 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	262,144	131,072	163,840	131,072	0	16,384	98,304	16,384	49,152	8,192
ipv4 Profile	524,288	0	163,840	0	0	16,384	131,072	0	65,536	0

TABLE 16 CAM partitioning profiles available for the BR-MLX-100Gx2-X modules disabled/enabled (continued)

Profile	IPv4 Note: All IPv4 values listed are when operating in double density mode.	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	L2 Inbound ACL	IPv4 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
ipv4-extended Profile	512K	64K	160K	256K	0	16K	96K	16K	48K	8K
ipv6 Profile	32,768	491,520	163,840	0	0	16,384	32,768	49,152	16,384	24,576
l2-metro Profile	524,288	0	163,840	0	0	16,384	131,072	0	65,536	0
l2-metro-2 Profile	524,288	0	163,840	0	0	16,384	131,072	0	65,536	0
multi-service Profile	147,456	81,920	163,840	294,912	0	16,384	65,536	32,768	32,768	16,384
multi-service-2 Profile	262,144	131,072	163,840	131,072	0	16,384	98,304	16,384	49,152	8,192
multi-service-3 Profile	131,072	65,536	163,840	196,608	131,072	16,384	81,920	24,576	32,768	16,384
multi-service-4 Profile	393,216	32,768	163,840	32,768	65,536	16,384	65,536	32,768	49,152	8,192
multi-service-7 Profile	768K	16K	160K	32K	32K	16K	64K	32K	48K	8K
MPLS-VPN-VPLS Profile	65,536	0	163,840	458,752	0	16,384	131,072	0	65,536	0
MPLS-L3 VPN Profile	131,072	0	163,840	393,216	0	16,384	131,072	0	65,536	0
MPLS-L3 VPN-2 Profile	65,536	0	163,840	458,752	0	16,384	131,072	0	65,536	0
MPLS-VPLS Profile	524,288	0	163,840	0	0	16,384	131,072	0	65,536	0
MPLS-VPLS-2 Profile	524,288	0	163,840	0	0	16,384	131,072	0	65,536	0
IPv4-VPN Profile	163,840	0	163,840	360,448	0	16,384	131,072	0	65,536	0
IPv4-IPv6 Profile	163,840	360,448	163,840	0	0	16,384	49,152	40,960	32,768	16,384

TABLE 16 CAM partitioning profiles available for the BR-MLX-100Gx2-X modules disabled/enabled (continued)

Profile	IPv4 Note: All IPv4 values listed are when operating in double density mode.	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	L2 Inbound ACL	IPv4 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
IPv4-VPLS Profile	524,288	0	163,840	0	0	16,384	32,768	0	0	0
IPv4-IPv6-2 Profile	393,216	65K	163,840	0	0	16,384	98,304	16,384	49,152	8,192
Multi-service-5 Profile	393,216	131,072	65,536	0	0	65,536	98,304	16,384	49,152	8,192
Multi-service-6 Profile	327,680	131,072	65,536	32,768	32768	65,536	98,304	16,384	49,152	8,192
Telemetry-1 Profile	81,920	98,304	131,072	0	0	32,768	81,920	24,576	49,152	8,192

TABLE 17 CAM partitioning profiles available for the BR-MLX-10Gx24-DM modules

Profile	IPv4	IPv6	MAC	IPv4 VPN	IPv4 ACL/MCAST VPLS	IPv6 VPN	IPv6 DAVC	IPv6 ACL	OUT ACL	OUT_ IPv6 ACL	Src_In grs Chk	MCAST VPLS	OUT_ LBL ACL	SRVC LKUP	L2 ACL
Default Profile	128K	16K	128K	32K	48K	0	0	16K	48K	8K	0	NA	NA	64K	16K
ipv4 Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
ipv4-extended Profile	128K	16K	128K	32K	48K	0	0	16K	48K	8K	0	NA	NA	64K	16K
ipv6 Profile	32K	56K	128K	0	16K	0	0	32K	16K	24K	0	NA	NA	64K	16K
mpls-vpn Profile	64K	0	128K	96K	80K	0	0	0	64K	0	0	NA	NA	64K	16K
mpls-vpls Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
l2-metro Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
l2-metro-2 Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
mpls-vpn-2 Profile	32K	0	128K	112K	80K	0	0	0	64K	0	0	NA	NA	64K	16K

TABLE 17 CAM partitioning profiles available for the BR-MLX-10Gx24-DM modules (continued)

Profile	IPv4	IPv6	MAC	IPv4 VPN	IPv4 ACL/ MCAST VPLS	IPv6 VPN	IPv6 DAVC	IPv6 ACL	OUT ACL	OUT_ IPv6 ACL	Src_In grs Chk	MCAST VPLS	OUT_ LBL ACL	SRVC LKUP	L2 ACL
mpls-vpls-2 Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
multi-service Profile	96K	8	128K	64K	48K	0	0	16K	32K	16K	0	NA	NA	64K	16K
mpls-vpn-vpls Profile	64K	0	128K	96K	80K	0	0	0	64K	0	0	NA	NA	64K	16K
ipv4-vpn Profile	64K	0	128K	96K	80K	0	0	0	64K	0	0	NA	NA	64K	16K
ipv4-ipv6 Profile	96K	40	128K	0	64K	0	0	8K	32K	16K	0	NA	NA	64K	16K
ipv4-vpls Profile	256K	0	128K	0	80K	0	0	0	64K	0	0	NA	NA	64K	16K
ipv4-ipv6-2 Profile	224K	8K	128K	0	48K	0	0	16K	48K	8K	0	NA	NA	64K	16K
multi-service-2 Profile	6K	8K	128K	16K	64K	0	0	8K	48K	8K	0	NA	NA	64K	16K
multi-service-3 Profile	192K	8K	128K	48K	48K	16K	0	16K	32K	16K	0	NA	NA	64K	16K
multi-service-4 Profile	128K	8K	128K	32K	48K	8K	0	16K	48K	8K	0	NA	NA	64K	16K
multi-service-7 Profile	128K	8K	128K	32K	48K	8K	0	16K	48K	8K	0	NA	NA	64K	16K

TABLE 18 CAM partitioning profiles available for the BR-MLX-10Gx20-M (1G/10G combo), or BR-MLX-1Gx20-U10G-M, and BR-MLX-100Gx2-CFP2-M

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	512K	64K	128K	128K	0	48K	4K	48K	4K
ipv4 Profile	1024K	0	32K	0	0	112K	0	64K	0
ipv4-extended Profile	512K	64K	128K	128K	0	48K	4K	48K	4K
ipv6 Profile	64K	240K	32K	0	0	16K	24K	16K	12K
l2-metro Profile	384K	0	512K	0	0	32K	0	64K	0

TABLE 18 CAM partitioning profiles available for the BR-MLX-10Gx20-M (1G/10G combo), or BR-MLX-1Gx20-U10G-M, and BR-MLX-100Gx2-CFP2-M (continued)

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
l2-metro-2 Profile	128K	0	640K	0	0	32K	0	64K	0
multi-service Profile	256K	32K	192K	256K	0	32K	8K	32K	8K
multi-service-2 Profile	512K	64K	128K	128K	0	48K	4K	40K	6K
multi-service-3 Profile	256K	32K	128K	192K	64K	32K	8K	32K	8K
multi-service-4 Profile	768K	32K	64K	64K	32K	32K	8K	48K	4K
multi-service-5 Profile	768K	32K	64K	0	0	96K	8K	48K	4K
multi-service-6 Profile	640K	96K	32K	32K	16K	48K	8K	48K	4K
multi-service-7 Profile	768K	32K	64K	64K	32K	32K	8K	48K	4K
telemetry-1 Profile	512K	64K	32K	0	0	96K	20K	32K	8K
mpls-vpn-vpls Profile	128K	0	224K	384K	0	48K	0	64K	0
mpls-l3vpn Profile	256K	0	32K	480K	0	64K	0	64K	0
mpls-l3vpn-2 Profile	128K	0	32K	544K	0	64K	0	64K	0
mpls-vpls Profile	256K	0	512K	0	0	64K	0	64K	0
mpls-vpls-2 Profile	128K	0	576K	0	0	64K	0	64K	0
ipv4-vpn Profile	320K	0	32K	448K	0	64K	0	64K	0
ipv4-ipv6 Profile	320K	160K	32K	0	0	48K	20K	32K	8K
ipv4-vpls Profile	320K	0	480K	0	0	64K	0	64K	0
ipv4-ipv6-2 Profile	768K	64K	64K	0	0	64K	8K	40K	6K

TABLE 19 CAM partitioning profiles available for the BR-MLX-10Gx20-X2 (1G/10G combo), or BR-MLX-1Gx20-U10G-X2, and BR-MLX-100Gx2-CFP2-X2

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	1,458,176	425,984	229,376	622,592	0	147,456	12,288	49,152	4,096
ipv4 Profile	2,506,752	49,152	65,536	0	0	278,528	0	65,536	0
ipv4-extended Profile	1792K	416K	224K	240K	0	144K	12K	48K	4K
ipv6 Profile	262,144	1,884,160	65,536	0	0	32,768	61,440	16,384	12,288
l2-metro Profile	2,097,152	0	557,056	0	0	32,768	0	65,536	0
l2-metro-2 Profile	2,097,152	0	589,824	0	0	16,384	0	65,536	0
multi-service Profile	1,048,576	720,896	196,608	720,896	0	131,072	20,480	32,768	8,192
multi-service-2 Profile	1,359,872	524,288	262,144	608K	0	128K	12K	40K	6K
multi-service-3 Profile	768K	512K	192K	608K	400K	128K	20K	32K	8K
multi-service-4 Profile	1024K	304K	128K	768K	304K	144K	24K	48K	4K
multi-service-5 Profile	1840K	704K	128K	0	0	128K	28K	48K	4K
multi-service-6 Profile	1120K	768K	128K	368K	224K	144K	24K	48K	4K
multi-service-7 Profile	976K	208K	128K	1024K	192K	144K	24K	48K	4K
telemetry-1 Profile	1936K	608K	64K	0	0	160K	28K	32K	8K
mpls-vpn-vpls Profile	512K	0	416K	1840K	0	96K	0	64K	0
mpls-l3vpn Profile	704K	0	128K	1712K	0	240K	0	64K	0
mpls-l3vpn-2 Profile	512K	0	128K	2048K	0	240K	0	64K	0
mpls-vpls Profile	2048K	0	416K	0	0	96K	0	64K	0
mpls-vpls-2 Profile	2048K	0	544K	0	0	32K	0	64K	0
ipv4-vpn Profile	1024K	0	160K	1424K	0	224K	0	64K	0

TABLE 19 CAM partitioning profiles available for the BR-MLX-10Gx20-X2 (1G/10G combo), or BR-MLX-1Gx20-U10G-X2, and BR-MLX-100Gx2-CFP2-X2 (continued)

Profile	IPv4	IPv6	MAC or VPLS MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
ipv4-ipv6 Profile	1536K	1024K	64K	0	0	128K	36K	32K	8K
ipv4-vpls Profile	2048K	0	416K	0	0	96K	0	64K	0
ipv4-ipv6-2 Profile	2048K	1024K	32K	0	0	16K	4K	40K	6K

TABLE 20 CAM partitioning profiles available for the BR-MLX-40Gx4-M module disable/enable

Profile	IPv4	IPv6	MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Default Profile	256K	65,536	65,536	65,536	0	28,672	6,144	49,152	8,192
IPv4 Profile	524,288	0	16,384	0	0	53,248	0	65,536	0
IPv4-extended Profile	256K	32K	64K	64K	0	36K	6K	48K	8K
IPv6 Profile	65,536	229,376	16,384	0	0	16,384	20,480	16,384	24,576
MPLS L3 VPN Profile	131,076	0	16,384	196,608	0	53,248	0	65,536	0
MPLS VPLS Profile	524,288	0	180,224	0	0	16,384	0	65,536	0
L2 Metro Profile	524,288	0	180,224	0	0	16,384	0	65,536	0
L2 Metro 2 Profile	524,288	0	180,224	0	0	16,384	0	65,536	0
MPLS L3 VPN 2 Profile	65,536	0	16,384	229,376	0	53,248	0	65,536	0
MPLS VPLS 2 Profile	524,288	0	180,224	0	0	16,384	0	65,536	0
Multi-Service	131,072	32,768	98,304	163,840	0	20,480	6,144	32,768	16,384
MPLS VPN VPLS Profile	131,072	0	114,688	196,608	0	28,672	0	65,536	0
IPv4 VPN Profile	163,840	0	16,384	180,224	0	53,248	0	65,536	0
IPv4 IPv6 Profile	163,840	180,224	16,384	0	0	16,384	18,432	32,768	16,384
IPv4 VPLS Profile	524,288	0	180,224	0	0	16,384	0	65,536	0
IPv4 IPv6 2 Profile	491,520	16,384	32,768	0	0	32,768	8,192	40,960	12,288

TABLE 20 CAM partitioning profiles available for the BR-MLX-40Gx4-M module disable/enable (continued)

Profile	IPv4	IPv6	MAC	IPv4 VPN	IPv6 VPN	IPv4 or L2 Inbound ACL	IPv6 Inbound ACL	IPv4 or L2 Outbound ACL	IPv6 Outbound ACL
Multi-service 2 Profile	458,752	16,384	32,768	16,384	0	32,768	8,192	40,960	12,288
Multi-service 3	131,072	32,768	65,536	114,688	49,152	24,576	8,192	32,768	16,384
Multi-service 4	425,984	16,384	32,768	16,384	16,384	32,768	8,192	49,152	8,192
Multi-Service 5	458,752	32,768	16,384	0	0	49,152	4,096	49,152	8,192
Multi-Service 5	327,680	65,536	32,768	16,384	16,384	32,768	8,192	49,152	8,192
Multi-Service 7	416K	8K	32K	16K	8K	40K	8K	48K	8K
Telemetry 1 Profile	327,680	98,304	16,384	0	0	32,768	12,288	32,768	16,384

TABLE 21 CAM partitioning profiles available for the NI-MLX-10Gx8-M modules

Profile	IPv4 (Shared with Src_Ingrs_Chk) Note at end *	IPv6	IPv4 VPN	MAC	IPv4 VPN	MAC	IPv6 VPN	IPv4/L2 ACL (Share with v4 MCAST_V PLS)	IPv6 ACL	Super-v4 ACL (Share with v4/l2 ACL)*	Super-v6 ACL (Share with v6 ACL)*	SRVC LKUP	OUT ACL (Shared with OUT_L BL ACL)	OUT_I Pv6 ACL
Default profile	8 (8x16K*2 = 256K)	4 (4x16K/2 = 32K)	4(4x16K*1 = 64K)	0	0	8 (8x8K*1 = 64K)	0	6 (6x8K/2 = 24K)	2(2x8K/8 = 2K)	0-12K	0-2K	1 (1x8K*1 = 8K)	11 (11x8K/2 = 44K)	4 (4x8K/8 = 4K)
IPv4 Optimized	16 (512K)	0	0	0	0	2(16K)	0	14(56K)	0	0-28K	0	1(8K)	15(60K)	0
IPv4-extended	8 (8x16K*2 = 256K)	4 (4x16K/2 = 32K)	4(4x16K*1 = 64K)	0	0	8 (8x8K*1 = 64K)	0	6 (6x8K/2 = 24K)	2(2x8K/8 = 2K)	0-12K	0-2K	1 (1x8K*1 = 8K)	11 (11x8K/2 = 44K)	4 (4x8K/8 = 4K)
IPv6 Optimized	2 (64K)	14 (112K)	0	0	0	2(16K)	0	2(8K)	12(12K)	0-4K	0-12K	1(8K)	3(12K)	12(12K)
MPLS VPN Optimized	4 (128K)	0	12(192K)	0	6 (6x8K*1 = 48K)	2(16K)	0	8(32K)	0	0-16K	0	1(8K)	15(60K)	0
MPLS VPLS Optimized	4 (128K)	0	0	12(12x16K*1 = 192K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60K)	0
L2 Metro	4 (128K)	0	0	12(192K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60K)	0

TABLE 21 CAM partitioning profiles available for the NI-MLX-10Gx8-M modules (continued)

Profile	IPv4 (Shared with Src_Ingrs _Chk) Note at end *	IPv6	IPv4 VPN	MAC	IPv4 VPN	MAC	IPv6 VPN	IPv4/L2 ACL (Share with v4 MCAST_V PLS)	IPv6 ACL	Super- v4 ACL (Share with v4/I2 ACL)*	Super- v6 ACL (Share with v6 ACL)*	SRVC LKUP	OUT ACL (Shared with OUT_L BL ACL)	OUT_I Pv6 ACL
Optimi zed														
L2 Metro Optimi zed #2	2 (64K)	0	0	0	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
MPLS VPN Optimi zed #2	2 (64K)	0	14(22 4K)	0	6(48K)	2(16K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
MPLS VPLS Optimi zed #2	2 (64K)	0	0	14(22 4K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
Multi- Service	4 (128K)	2 (16K)	8(128 K)	2(32K)	0	8(64K)	0	4(16K)	4(4 K)	0-8K	0-4K	1(8K)	7(28K)	8(8K)
MPLS VPN +VPLS	2 (64K)	0	12(19 2K)	2(32K)	0	10(80 K)	0	6(24K)	0	0-12K	0	1(8K)	15(60 K)	0
IPv4 + VPN	5 (160K)	0	11(17 6K)	0	6(48K)	2(16K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
IPv6 + IPv4	5 (160K)	10 (80K)	0	1(16K)	0	0	0	6(24K)	10(10K)	0-12K	0-10 K	1(8K)	7(28K)	8(8K)
IPv4 + VPLS	5 (160K)	0	0	11(17 6K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
IPv4 + Ipv6 2	15 (480K)	1 (8K)	0	0	0	4(32K)	0	8(32K)	4(4 K)	0-16K	0-4K	1(8K)	11(44 K)	4(4K)
Multi- service 2	14 (448K)	2 (16K)	0	0	4(32K)	4(32K)	0	6(24K)	2(2 K)	0-12K	0-2K	1(8K)	11(44 K)	4(4K)
Multi- service 3	4 (128K)	2 (16K)	6 (96K)	4(64K)	0	0	8(8x8K /2 = 32K)	4(16K)	4(4 K)	0-8K	0-4K	1(8K)	7(28K)	8(8K)
Multi- service 4	14 (448K)	1 (8K)	1 (16K)	0	2(16K)	4(32K)	2(8K)	4(16K)	4(4 K)	0-8K	0-4K	1(8K)	11(44 K)	4(4K)
Multi- Service 5	12 (384K)	2 (16K)	0	2(32K)	0	0	0	12(48K)	4(4 K)	0-24K	0-4K	1(8K)	11(44 K)	4(4K)
Multi- Service -7	14 (448K)	1 (8K)	1 (16K)	0	2 (16K)	4 (32K)	2 (8K)	4 (16K)	4 (4K)	0-8K	0-4K	1 (8K)	11(44 K)	4 (4K)
L3- Optimi zed	12(384K)	4(32 K)	0	0	4(32K)	4(32K)	4(16K)	2(8K)	2(2 K)	0-4K	0-2K	1(8K)	11(44 K)	4(4K)

TABLE 22 CAM partitioning profiles available for the BR-MLX-10Gx8-X modules

Profile	IPv4 (Shared with Src_Ingrs _Chk) Note at end *	IPv6	IPv4 VPN	MAC	IPv4 VPN	MAC	IPv6 VPN	IPv4/L2 ACL (Share with v4 MCAST_V PLS)	IPv6 ACL	Super- v4 ACL (Share with v4/L2 ACL)*	Super- v6 ACL (Share with v6 ACL)*	SRVC LKUP	OUT ACL (Shared with OUT_L BL ACL)	OUT_I Pv6 ACL
Default profile	16 (16X16K *2 = 512K)	8 (8X16 6K/2 = 64K)	8 (8X16 K*1 = 128K)	0	0	8 (8X16 K*1 = 128K)	0	6 (6X16K/2 = 48K)	2 (2X 16K /8 = 4K)	0-24K	0-4K	10 (10x8K *1 = 80K)	16 (16X8 K/2 = 64K)	6 (6X8K/ 8 = 6K)
IPv4 Optimi zed	32(512K)	0	0	0	0	2(32K)	0	14(112K)	0	0-56K	0	10(80 K)	22(88 K)	0
IPv4- extend ed	16 (16X16K *2 = 512K)	8 (8X16 6K/2 = 64K)	8 (8X16 K*1 = 128K)	0	0	8 (8X16 K*1 = 128K)	0	6 (6X16K/2 = 48K)	2 (2X 16K /8 = 4K)	0-24K	0-4K	10 (10x8K *1 = 80K)	16 (16X8 K/2 = 64K)	6 (6X8K/ 8 = 6K)
IPv6 Optimi zed	2(64K)	30(240K)	0	0	0	2(32K)	0	2(16K)	12(24K)	0-8K	0-24 K	10(80 K)	4(16K)	18(18K)
MPLS VPN Optimi zed	8(256K)	0	24(384K)	0	6 (6X16 K*1 = 96K)	2(32K)	0	8(64K)	0	0-32K	0	10(80 K)	22(88 K)	0
MPLS VPLS Optimi zed	8(256K)	0	0	24(24x 16K*1 =384K)	0	8(128 K)	0	8(64K)	0	0-32K	0	10(80 K)	22(88 K)	0
L2 Metro Optimi zed	8(256K)	0	0	24(384K)	0	8(128 K)	0	8(64K)	0	0-32K	0	10(80 K)	22(88 K)	0
L2 Metro Optimi zed #2	2(64K)	0	0	30(480K)	0	8(128 K)	0	8(64K)	0	0-32K	0	10(80 K)	22(88 K)	0
MPLS VPN Optimi zed #2	4(128K)	0	28(448K)	0	6(96K)	2(32K)	0	8(64K)	0	0-32K	0	10(80 K)	22(88 K)	0
MPLS VPLS Optimi zed #2	4(128K)	0	0	28(448K)	0	8(128 K)	0	8(64K)	0	0-32K	0	10(80 K)	22(88 K)	0
Multi- Service	8(256K)	4(32K)	16(256K)	4(64K)	0	8(128 K)	0	4(32K)	4(8K)	0-16K	0-8K	10(80 K)	10(40 K)	12(12K)
MPLS VPN +VPLS	4(128K)	0	24(384K)	4(64K)	0	10(160K)	0	6(48K)	0	0-24K	0	10(80 K)	22(88 K)	0
IPv4 + VPN	10(320K)	0	22(352K)	0	6(96K)	2(32K)	0	8(64K)	0	0-32K	0	10(80 K)	22(88 K)	0

TABLE 22 CAM partitioning profiles available for the BR-MLX-10Gx8-X modules (continued)

Profile	IPv4 (Shared with Src_Ingrs _Chk) Note at end *	IPv6	IPv4 VPN	MAC	IPv4 VPN	MAC	IPv6 VPN	IPv4/L2 ACL (Share with v4 MCAST_V PLS)	IPv6 ACL	Super- v4 ACL (Share with v4/L2 ACL)*	Super- v6 ACL (Share with v6 ACL)*	SRVC LKUP	OUT ACL (Shared with OUT_L BL ACL)	OUT_I Pv6 ACL
IPv6 + IPv4	10(320K)	20(160K)	0	2(32K)	0	0	0	6(48K)	10(20K)	0-24K	0-20K	10(80K)	10(40K)	12(12K)
IPv4 + VPLS	10(320K)	0	0	22(352K)	0	8(128K)	0	8(64K)	0	0-32K	0	10(80K)	22(88K)	0
IPv4 + Ipv6 2	24(768K)	8(64K)	0	0	0	4(64K)	0	8(64K)	4(8K)	0-32K	0-8K	10(80K)	16(64K)	6(6K)
Multi- service 2	16(512K)	8(64K)	8(128K)	0	0	8(128K)	0	6(48K)	2(4K)	0-24K	0-4K	10(80K)	16(64K)	6(6K)
Multi- service 3	8(256K)	4(32K)	12(192K)	8(128K)	0	0	8(64K)	4(32K)	4(8K)	0-16K	0-8K	10(80K)	10(40K)	12(12K)
Multi- service 4	24(768K)	4(32K)	4(64K)	0	0	4(64K)	4(32K)	4(32K)	4(8K)	0-16K	0-8K	10(80K)	16(64K)	6(6K)
Multi- Service 5	24(768K)	4(32K)	0	4(64K)	0	0	0	12(96K)	4(8K)	0-48K	0-8K	10(80K)	16(64K)	6(6K)
Multi- Service 6	20(640K)	12(96K)	0	0	2(32K)	2(32K)	2(16K)	6(48K)	4(8K)	0-24K	0-8K	10(80K)	16(64K)	6(6K)
Multi- Service 7	24(768K)	4(32K)	4(64K)	0	0	4(64K)	4(32K)	4(32K)	4(8K)	0-16K	0-8K	10(80K)	16(64K)	6(6K)
L3- Optimi- zed	24(768K)	8(64K)	0	0	4(64K)	4(64K)	4(32K)	2(16K)	2(4K)	0-8K	0-4K	10(80K)	16(64K)	6(6K)

TABLE 23 CAM partitioning profiles available for the NI-MLX-10Gx8-D modules

Profile	IPv4 (Shared with Src_Ingrs _Chk) Note at end *	IPv6	IPv4 VPN	MAC	IPv4 VPN	MAC	IPv6 VPN	IPv4/L2 ACL (Share with v4 MCAST_V PLS)	IPv6 ACL	Super- v4 ACL (Share with v4/L2 ACL)*	Super- v6 ACL (Share with v6 ACL)*	SRVC LKUP	OUT ACL (Shared with OUT_L BL ACL)	OUT_I Pv6 ACL
Default	4 (4x16K*2 = 128K)	2 (2x16K/ 2 = 16K)	2 (2x16K *1 = 32K)	0	0	8 (8x8K* 1 =64K)	0	6(6x8K/2 = 24K)	2(2x8K /8 = 2K)	0-12K	0-2K	1 (1x8K* 1 = 8K)	11 (11x8K/2 = 44K)	4 (4x8K/ 8 = 4K)
IPv4 Optimi- zed	8(256K)	0	0	0	0	2(16K)	0	14(56K)	0	0-28K	0	1(8K)	15(60K)	0

TABLE 23 CAM partitioning profiles available for the NI-MLX-10Gx8-D modules (continued)

Profile	IPv4 (Shared with Src_Ingrs _Chk) Note at end *	IPv6	IPv4 VPN	MAC	IPv4 VPN	MAC	IPv6 VPN	IPv4/L2 ACL (Share with v4 MCAST_V PLS)	IPv 6 AC L	Super- v4 ACL (Share with v4/L2 ACL)*	Super- v6 ACL (Share with v6 ACL)*	SRVC LKUP	OUT ACL (Shared with OUT_L BL ACL)	OUT_I Pv6 ACL
IPv4- extend ed	4 (4x16K*2 = 128K)	2 (2x16K/ 2 = 16K)	2 (2x16K *1 = 32K)	0	0	8 (8x8K* 1 =64K)	0	6(6X8K/2 = 24K)	2(2 x8K /8 = 2K)	0-12K	0-2K	1 (1x8K* 1 = 8K)	11 (11X8 K/2 = 44K)	4 (4X8K/ 8 = 4K)
IPv6 Optimi zed	2(64K)	6(48 K)	0	0	0	2(16K)	0	2(8K)	12(12 K)	0-4K	0-12K	1(8K)	3(12K)	12(12 K)
MPLS VPN Optimi zed	2(64K)	0	6(96K)	0	6 (6x8K* 1 = 48K)	2(16K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
MPLS VPLS Optimi zed	2(64K)	0	0	6 (6x16K *1 = 96K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
L2 Metro Optimi zed	2(64K)	0	0	6(96K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
L2 Metro Optimi zed #2	2(64K)	0	0	6(96K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
MPLS VPN Optimi zed #2	2(64K)	0	6(96K)	0	6(48K)	2(16K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
MPLS VPLS Optimi zed #2	2(64K)	0	0	6(96K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
Multi- Service	2(64K)	1(8K)	4(64K)	1(16K)	0	8(64K)	0	4(16K)	4 (4K)	0-8K	0-4K	1(8K)	7(28K)	8(8K)
MPLS VPN +VPLS	2(64K)	0	6(96K)	0	0	10(80 K)	0	6(24K)	0	0-12K	0	1(8K)	15(60 K)	0
IPv4 + VPN	3(96K)	0	5(80K)	0	6(48K)	2(16K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
IPv6 + IPv4	3(96K)	4(32 K)	0	1(16K)	0	0	0	6(24K)	10(10 K)	0-12K	0-10K	1(8K)	7(28K)	8(8K)
IPv4 + VPLS	2(64K)	0	0	6(96K)	0	8(64K)	0	8(32K)	0	0-16K	0	1(8K)	15(60 K)	0
IPv4 + Ipv6 2	6(192K)	2(16 K)	0	0	0	4(32K)	0	8(32K)	4 (4K)	0-16K	0-4K	1(8K)	11(44 K)	4(4K)

TABLE 23 CAM partitioning profiles available for the NI-MLX-10Gx8-D modules (continued)

Profile	IPv4 (Shared with Src_Ingrs _Chk) Note at end *	IPv6	IPv4 VPN	MAC	IPv4 VPN	MAC	IPv6 VPN	IPv4/L2 ACL (Share with v4 MCAST_V PLS)	IPv 6 AC L	Super- v4 ACL (Share with v4/L2 ACL)*	Super- v6 ACL (Share with v6 ACL)*	SRVC LKUP	OUT ACL (Shared with OUT_L BL ACL)	OUT_I Pv6 ACL
Multi- service 2	6(192K)	2(16 K)	0	0	4(32K)	4(32K)	0	6(24K)	2(2 K)	0-12K	0-2K	1(8K)	11(44 K)	4(4K)
Multi- service 3	2(64K)	1(8K)	3(48K)	2(32K)	0	0	8 (8x8K/ 2 =32K)	4(16K)	4 (4K)	0-8K	0-4K	1(8K)	7(28K)	8(8K)
Multi- service 4	6(192K)	1(8K)	1(16K)	0	0	4(32K)	4(16K)	4(16K)	4 (4K)	0-8K	0-4K	1(8K)	11(44 K)	4(4K)
Multi- Service 5	6(192K)	1(8K)	0	1(16K)	0	0	0	12(48K)	4(4 K)	0-48K	0-8K	1(8K)	11(44 K)	4(4K)
Multi- Service 7	6(192K)	1(8K)	1(16K)	0	0	4(32K)	4(16K)	4(16K)	4(4 K)	0-8K	0-4K	1(8K)	11(44 K)	4(4K)
L3- Optimi- zed	6(192K)	2(16 K)	0	0	4(32K)	4(32K)	4(16K)	2(8K)	2(2 K)	0-4K	0-2K	10(80 K)	16(64 K)	6(6K)

Supernet CAM partition sharing

TCAM sharing within a particular CAM section is supported.

TCAM allocation is optimized to allow dynamic allocation of resources to each level within a particular resource pool. If one level runs out of TCAM resources, it can use resources that have been allocated to another level and remain unused. This feature is applicable to IPv4, IPv6, and Layer 3 VPN routes.

NOTE

CAM Sharing is not shared across resource pools, such as IPv4, IPv6 or Layer 3 VPN. Only shared between levels of each pool. For example: IPv4 may not use CAM resources from the IPv6 resource pool.

Displaying CAM partition

The **show cam-partition** command provides information about available CAM in three formats: raw size, user size, and reserved size.

```
device# show cam-partition
CAM partitioning profile: default
Slot 1 XPP20SP 0:
# of CAM device          = 4
Total CAM Size           = 917504 entries (63Mbits)
IP: Raw Size 524288, User Size 524288(0 reserved)
  Subpartition 0: Raw Size 12288, User Size 12288, (0 reserved)
  Subpartition 1: Raw Size 468107, User Size 468107, (0 reserved)
  Subpartition 2: Raw Size 37335, User Size 37335, (0 reserved)
  Subpartition 3: Raw Size 5140, User Size 5140, (0 reserved)
  Subpartition 4: Raw Size 778, User Size 778, (0 reserved)
IPv6: Raw Size 131072, User Size 65536(0 reserved)
  Subpartition 0: Raw Size 12288, User Size 6144, (0 reserved)
  Subpartition 1: Raw Size 107496, User Size 53748, (0 reserved)
```

```

Subpartition 2: Raw Size 9332, User Size 4666, (0 reserved)
Subpartition 3: Raw Size 1284, User Size 642, (0 reserved)
Subpartition 4: Raw Size 384, User Size 192, (0 reserved)
IP VPN Raw Size 131072, User Size 131072(0 reserved)
Subpartition 0: Raw Size 2048, User Size 2048, (0 reserved)
Subpartition 1: Raw Size 116886, User Size 116886, (0 reserved)
Subpartition 2: Raw Size 9333, User Size 9333, (0 reserved)
Subpartition 3: Raw Size 1285, User Size 1285, (0 reserved)
Subpartition 4: Raw Size 384, User Size 384, (0 reserved)
MAC: Raw Size 131072, User Size 131072(0 reserved)
Subpartition 0: Raw Size 10, User Size 10, (0 reserved)
Subpartition 1: Raw Size 32, User Size 32, (0 reserved)
Subpartition 2: Raw Size 131030, User Size 131030, (0 reserved)
Session: Raw Size 98304, User Size 49152(0 reserved)
Subpartition 0: Raw Size 79872, User Size 39936, (0 reserved)
Subpartition 1: Raw Size 2048, User Size 1024, (0 reserved)
Subpartition 2: Raw Size 16384, User Size 8192, (0 reserved)
IPv6 Session: Raw Size 32768, User Size 4096(0 reserved)
Subpartition 0: Raw Size 15872, User Size 1984, (0 reserved)
Subpartition 1: Raw Size 512, User Size 64, (0 reserved)
Subpartition 2: Raw Size 16384, User Size 2048, (0 reserved)
Out Session: Raw Size 196608, User Size 98304(49152 reserved)
Out IPv6 Session: Raw Size 65536, User Size 8192(4096 reserved)
Slot 1 XPP20SP 0:
IP Section(Left):      0(000000) - 262143(03ffff)
IP Section(Right):     0 (000000) - 262143 (03ffff)
IP SNet 0:(Left):      0(000000) - 12287(002fff)
IP SNet 1:(Left):     12288(003000) - 262143(03ffff)
IP SNet 1:(Right):     0 (000000) - 218250 (03548a)
IP SNet 2:(Right):    218251 (03548b) - 255585 (03e661)
IP SNet 3:(Right):    255586 (03e662) - 260725 (03fa75)
IP SNet 4:(Right):    260726 (03fa76) - 261503 (03fd7f)
IP SNet 5:(Right):    261504 (03fd80) - 261631 (03fdff)
IP SNet 6:(Right):    261632 (03fe00) - 261695 (03fe3f)
IP SNet 7:(Right):    261696 (03fe40) - 261727 (03fe5f)
IP SNet 8:(Right):    261728 (03fe60) - 261759 (03fe7f)
IP SNet 9:(Right):    261760 (03fe80) - 261791 (03fe9f)
IP SNet 10:(Right):   261792 (03fea0) - 261807 (03feaf)
IP SNet 11:(Right):   261808 (03feb0) - 261823 (03febf)
IP SNet 12:(Right):   261824 (03fec0) - 261839 (03fecf)
IP SNet 13:(Right):   261840 (03fed0) - 261855 (03fedf)
IP SNet 14:(Right):   261856 (03fee0) - 261871 (03feef)
IP SNet 15:(Right):   261872 (03fef0) - 261887 (03fef7)
IP SNet 16:(Right):   261888 (03ff00) - 261903 (03ff0f)
IP SNet 17:(Right):   261904 (03ff10) - 261919 (03ff1f)
IP SNet 18:(Right):   261920 (03ff20) - 261935 (03ff2f)
IP SNet 19:(Right):   261936 (03ff30) - 261951 (03ff3f)
IP SNet 20:(Right):   261952 (03ff40) - 261967 (03ff4f)
IP SNet 21:(Right):   261968 (03ff50) - 261983 (03ff5f)
IP SNet 22:(Right):   261984 (03ff60) - 261999 (03ff6f)
IP SNet 23:(Right):   262000 (03ff70) - 262015 (03ff7f)
IP SNet 24:(Right):   262016 (03ff80) - 262031 (03ff8f)
IP SNet 25:(Right):   262032 (03ff90) - 262047 (03ff9f)
IP SNet 26:(Right):   262048 (03ffa0) - 262063 (03ffaf)
IP SNet 27:(Right):   262064 (03ffb0) - 262079 (03ffbf)
IP SNet 28:(Right):   262080 (03ffc0) - 262095 (03ffcf)
IP SNet 29:(Right):   262096 (03ffd0) - 262111 (03ffdf)
IP SNet 30:(Right):   262112 (03ffe0) - 262127 (03ffef)
IP SNet 31:(Right):   262128 (03fff0) - 262143 (03ffff)
IPv6 Section : 262144 (040000) - 393215 (05ffff)
IPv6 SNet 0: 262144 (040000) - 274431 (042fff)
IPv6 SNet 1: 274432 (043000) - 381927 (05d3e7)
IPv6 SNet 2: 381928 (05d3e8) - 391259 (05f85b)
IPv6 SNet 3: 391260 (05f85c) - 392543 (05fd5f)
IPv6 SNet 4: 392544 (05fd60) - 392927 (05fedf)
IPv6 SNet 5: 392928 (05fee0) - 393055 (05fff5)
IPv6 SNet 6: 393056 (05ff60) - 393119 (05fff9)
IPv6 SNet 7: 393120 (05ffa0) - 393151 (05ffbf)
IPv6 SNet 8: 393152 (05ffc0) - 393183 (05ffdf)
IPv6 SNet 9: 393184 (05ffe0) - 393215 (05ffff)
IP VPN Section: 393216 (060000) - 524287 (07ffff)
IP VPN SNet 0: 393216 (060000) - 395263 (0607ff)

```

```

IP VPN SNet 1: 395264 (060800) - 512149 (07d095)
IP VPN SNet 2: 512150 (07d096) - 521482 (07f50a)
IP VPN SNet 3: 521483 (07f50b) - 522767 (07fa0f)
IP VPN SNet 4: 522768 (07fa10) - 523151 (07fb8f)
IP VPN SNet 5: 523152 (07fb90) - 523279 (07fc0f)
IP VPN SNet 6: 523280 (07fc10) - 523343 (07fc4f)
IP VPN SNet 7: 523344 (07fc50) - 523375 (07fc6f)
IP VPN SNet 8: 523376 (07fc70) - 523407 (07fc8f)
IP VPN SNet 9: 523408 (07fc90) - 523439 (07fc9f)
IP VPN SNet 10: 523440 (07fcb0) - 523455 (07fcbf)
IP VPN SNet 11: 523456 (07fcc0) - 523471 (07fccf)
IP VPN SNet 12: 523472 (07fcd0) - 523487 (07fcd9)
IP VPN SNet 13: 523488 (07fce0) - 523503 (07fce9)
IP VPN SNet 14: 523504 (07fcf0) - 523519 (07fcff)
IP VPN SNet 15: 523520 (07fd00) - 523535 (07fd0f)
IP VPN SNet 16: 523536 (07fd10) - 523551 (07fd1f)
IP VPN SNet 17: 523552 (07fd20) - 523567 (07fd2f)
IP VPN SNet 18: 523568 (07fd30) - 523583 (07fd3f)
IP VPN SNet 19: 523584 (07fd40) - 523599 (07fd4f)
IP VPN SNet 20: 523600 (07fd50) - 523615 (07fd5f)
IP VPN SNet 21: 523616 (07fd60) - 523631 (07fd6f)
IP VPN SNet 22: 523632 (07fd70) - 523647 (07fd7f)
IP VPN SNet 23: 523648 (07fd80) - 523663 (07fd8f)
IP VPN SNet 24: 523664 (07fd90) - 523679 (07fd9f)
IP VPN SNet 25: 523680 (07fda0) - 523695 (07fdaf)
IP VPN SNet 26: 523696 (07fdb0) - 523711 (07fdbf)
IP VPN SNet 27: 523712 (07fdc0) - 523727 (07fdcf)
IP VPN SNet 28: 523728 (07fdd0) - 523743 (07fddf)
IP VPN SNet 29: 523744 (07fde0) - 523759 (07fdef)
IP VPN SNet 30: 523760 (07fdf0) - 523775 (07fdf9)
IP VPN SNet 31: 523776 (07fe00) - 524287 (07ffff)
MAC Section : 524288 (080000) - 655359 (09ffff)
MAC Forwarding: 524288 (080000) - 655317 (09ffd5)
MAC Flooding : 655318 (09ffd6) - 655327 (09ffdf)
Misc Protocol : 655350 (09fff6) - 655381 (0a0015)
Session Section: 655360 (0a0000) - 753663 (0b7fff)
IP Multicast : 655360 (0a0000) - 671743 (0a3fff)
Broadcast ACL : 673792 (0a4800) - 675839 (0a4fff)

Receive ACL : 671744 (0a4000) - 673791 (0a47ff)
Rule-based ACL: 673792 (0a4800) - 753663 (0b7fff)
IPv6 Session Sec: 753664 (0b8000) - 786431 (0bffff)
IP Multicast : 753664 (0b8000) - 770047 (0bbfff)
Receive ACL : 770048 (0bc000) - 770559 (0bc1ff)
Rule-based ACL: 770560 (0bc200) - 786431 (0bffff)
Out Session : 786432 (0c0000) - 983039 (0effff)
Out IPv6 Session: 983040 (0f0000) - 104857 (0fffff)
...
```

Syntax: show cam-partition

The output displays the CAM partitioning profile name, slot number, number of CAM device, and total CAM size. It also displays the raw size, user size, and reserved size for each of the CAM sub-partitions.

In Algorithmic mode, the subpartitions and subnets for IP, IP VPN, IPv6, and IPv6 VPN are not displayed in the output of the **show cam-partition** command.

```

device# show cam-partition
CAM partitioning profile: default

XPP100GEXE 0:
# of CAM device           = 1
Total CAM Size            = 4456448 entries (340Mbits)

MAC: Raw Size 229376, User Size 229376(0 reserved)
  Subpartition 0: Raw Size 4, User Size 4, (0 reserved)
  Subpartition 1: Raw Size 8, User Size 8, (0 reserved)
  Subpartition 2: Raw Size 229354, User Size 229354, (0 reserved)
  Subpartition 3: Raw Size 10, User Size 10, (0 reserved)
```

```

Session: Raw Size 294912, User Size 147456(0 reserved)
  Subpartition 0: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 1: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 2: Raw Size 276480, User Size 138240, (0 reserved)
  Subpartition 3: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 4: Raw Size 2048, User Size 1024, (0 reserved)
  Subpartition 5: Raw Size 16384, User Size 8192, (0 reserved)
  Subpartition 6: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 7: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 8: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 9: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 10: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 11: Raw Size 0, User Size 0, (0 reserved)

IPv6 Session: Raw Size 98304, User Size 12288(0 reserved)
  Subpartition 0: Raw Size 81920, User Size 10240, (0 reserved)
  Subpartition 1: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 2: Raw Size 16384, User Size 2048, (0 reserved)
  Subpartition 3: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 4: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 5: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 6: Raw Size 0, User Size 0, (0 reserved)

Out Session: Raw Size 98304, User Size 49152(0 reserved)

Out IPv6 Session: Raw Size 32768, User Size 4096(0 reserved)

Internal Forwarding Lookup: Raw Size 131072, User Size 131072(0 reserved)
  Subpartition 0: Raw Size 131071, User Size 131071, (0 reserved)
  Subpartition 1: Raw Size 1, User Size 1, (0 reserved)

IP: Raw Size 1458176, User Size 1458176(0 reserved)

IP VPN Raw Size 622592, User Size 622592(0 reserved)

IPv6: Raw Size 425984, User Size 425984(0 reserved)

MAC Section      : 163840 (028000) - 393215 (05ffff)
  Misc Protocol   : 163840 (028000) - 163849 (028009)
  MAC Forwarding: 163840 (028000) - 163849 (028009)
  MAC Flooding   : 163840 (028000) - 163849 (028009)
  PORT BUM RL    : 163840 (028000) - 163849 (028009)
  Misc Protocol   : 163850 (02800a) - 393203 (05fff3)
  MAC Forwarding: 163850 (02800a) - 393203 (05fff3)
  MAC Flooding   : 163850 (02800a) - 393203 (05fff3)
  PORT BUM RL    : 163850 (02800a) - 393203 (05fff3)
  Misc Protocol   : 393204 (05fff4) - 393211 (05fffb)
  MAC Forwarding: 393204 (05fff4) - 393211 (05fffb)
  MAC Flooding   : 393204 (05fff4) - 393211 (05fffb)
  PORT BUM RL    : 393204 (05fff4) - 393211 (05fffb)
  Misc Protocol   : 393212 (05fffc) - 393215 (05ffff)
  MAC Forwarding: 393212 (05fffc) - 393215 (05ffff)
  MAC Flooding   : 393212 (05fffc) - 393215 (05ffff)
  PORT BUM RL    : 393212 (05fffc) - 393215 (05ffff)

Session Section : 393216 (060000) - 688127 (0a7fff)
  Rule-based ACL : 411648 (064800) - 688127 (0a7fff)
  Receive ACL    : 409600 (064000) - 411647 (0647ff)
  IP Multicast   : 393216 (060000) - 409599 (063fff)

IPv6 Session Sec: 688128 (0a8000) - 786431 (0bffff)
  IP Multicast   : 688128 (0a8000) - 704511 (0abfff)
  Receive ACL    : 688128 (0a8000) - 704511 (0abfff)
  Rule-based ACL : 688128 (0a8000) - 704511 (0abfff)
  IP Multicast   : 704512 (0ac000) - 786431 (0bffff)

```



```

Receive ACL   : 704512 (0ac000) - 786431 (0bffff)
Rule-based ACL: 704512 (0ac000) - 786431 (0bffff)

Out Session   : 786432 (0c0000) - 884735 (0d7fff)

Out IPv6 Session: 884736 (0d8000) - 917503 (0dffff)

Internal Forward: 917504 (0e0000) - 104857 (0ffffff)
  IFL Main      : 104857 (0ffffff) - 104857 (0ffffff)
  IFL Openflow C: 104857 (0ffffff) - 104857 (0ffffff)
  IFL Main      : 917504 (0e0000) - 104857 (0ffffffe)
  IFL Openflow C: 917504 (0e0000) - 104857 (0ffffffe)

IP Section(Left): 104857 (100000) - 250675 (263fff)

IP VPN Section: 250675 (264000) - 312934 (2fbfff)

IPv6 Section   : 312934 (2fc000) - 355532 (363fff)

XPP100GEXE 1:
# of CAM device           = 1
Total CAM Size            = 4456448 entries (340Mbits)

MAC: Raw Size 229376, User Size 229376(0 reserved)
  Subpartition 0: Raw Size 4, User Size 4, (0 reserved)
  Subpartition 1: Raw Size 8, User Size 8, (0 reserved)
  Subpartition 2: Raw Size 229354, User Size 229354, (0 reserved)
  Subpartition 3: Raw Size 10, User Size 10, (0 reserved)

Session: Raw Size 294912, User Size 147456(0 reserved)
  Subpartition 0: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 1: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 2: Raw Size 276480, User Size 138240, (0 reserved)
  Subpartition 3: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 4: Raw Size 2048, User Size 1024, (0 reserved)
  Subpartition 5: Raw Size 16384, User Size 8192, (0 reserved)
  Subpartition 6: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 7: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 8: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 9: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 10: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 11: Raw Size 0, User Size 0, (0 reserved)

IPv6 Session: Raw Size 98304, User Size 12288(0 reserved)
  Subpartition 0: Raw Size 81920, User Size 10240, (0 reserved)
  Subpartition 1: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 2: Raw Size 16384, User Size 2048, (0 reserved)
  Subpartition 3: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 4: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 5: Raw Size 0, User Size 0, (0 reserved)
  Subpartition 6: Raw Size 0, User Size 0, (0 reserved)

Out Session: Raw Size 98304, User Size 49152(0 reserved)

Out IPv6 Session: Raw Size 32768, User Size 4096(0 reserved)

Internal Forwarding Lookup: Raw Size 131072, User Size 131072(0 reserved)
  Subpartition 0: Raw Size 131071, User Size 131071, (0 reserved)
  Subpartition 1: Raw Size 1, User Size 1, (0 reserved)

IP: Raw Size 1458176, User Size 1458176(0 reserved)

IP VPN Raw Size 622592, User Size 622592(0 reserved)

IPv6: Raw Size 425984, User Size 425984(0 reserved)

```

```

MAC Section      : 163840 (028000) - 393215 (05ffff)
  Misc Protocol  : 163840 (028000) - 163849 (028009)
  MAC Forwarding: 163840 (028000) - 163849 (028009)
  MAC Flooding   : 163840 (028000) - 163849 (028009)
  PORT BUM RL    : 163840 (028000) - 163849 (028009)
  Misc Protocol  : 163850 (02800a) - 393203 (05ffff3)
  MAC Forwarding: 163850 (02800a) - 393203 (05ffff3)
  MAC Flooding   : 163850 (02800a) - 393203 (05ffff3)
  PORT BUM RL    : 163850 (02800a) - 393203 (05ffff3)
  Misc Protocol  : 393204 (05ffff4) - 393211 (05ffffb)
  MAC Forwarding: 393204 (05ffff4) - 393211 (05ffffb)
  MAC Flooding   : 393204 (05ffff4) - 393211 (05ffffb)
  PORT BUM RL    : 393204 (05ffff4) - 393211 (05ffffb)
  Misc Protocol  : 393212 (05ffffc) - 393215 (05fffff)
  MAC Forwarding: 393212 (05ffffc) - 393215 (05fffff)
  MAC Flooding   : 393212 (05ffffc) - 393215 (05fffff)
  PORT BUM RL    : 393212 (05ffffc) - 393215 (05fffff)

Session Section : 393216 (060000) - 688127 (0a7fff)
  Rule-based ACL : 411648 (064800) - 688127 (0a7fff)
  Receive ACL    : 409600 (064000) - 411647 (0647ff)
  IP Multicast   : 393216 (060000) - 409599 (063fff)

IPv6 Session Sec: 688128 (0a8000) - 786431 (0bffff)
  IP Multicast   : 688128 (0a8000) - 704511 (0abfff)
  Receive ACL    : 688128 (0a8000) - 704511 (0abfff)
  Rule-based ACL: 688128 (0a8000) - 704511 (0abfff)
  IP Multicast   : 704512 (0ac000) - 786431 (0bffff)
  Receive ACL    : 704512 (0ac000) - 786431 (0bffff)
  Rule-based ACL: 704512 (0ac000) - 786431 (0bffff)

Out Session     : 786432 (0c0000) - 884735 (0d7fff)

Out IPv6 Session: 884736 (0d8000) - 917503 (0dffff)

Internal Forward: 917504 (0e0000) - 104857 (0ffffff)
  IFL Main       : 104857 (0ffffff) - 104857 (0ffffff)
  IFL Openflow C: 104857 (0ffffff) - 104857 (0ffffff)
  IFL Main       : 917504 (0e0000) - 104857 (0fffffe)
  IFL Openflow C: 917504 (0e0000) - 104857 (0fffffe)

IP Section(Left): 104857 (100000) - 250675 (263fff)

IP VPN Section: 250675 (264000) - 312934 (2fbfff)

IPv6 Section    : 312934 (2fc000) - 355532 (363fff)

```

Table 24 describes the output parameters of the **show cam-partition** command.

TABLE 24 Output parameters of the show cam-partition command

Field	Description
CAM partitioning profile	Shows the CAM profile name.
Slot	Shows the slot number.
# of CAM device	Shows the number of the CAM device.
Total CAM Size	Shows the total available CAM size.
Raw Size	Shows the value double that of the CAM partition standard entry count. A standard entry contains 64 bits for the data and 64 bits for the mask. The raw size may cover invalid entries.

TABLE 24 Output parameters of the show cam-partition command (continued)

Field	Description
User Size	Shows the actual number of entries that the application can use. For a 128-bit application, such as Layer 4 ACL and IPV6, two standard entries equal one user entry. The user size may also cover invalid entries.
reserved	Shows the number of entries not usable in a specific sub-partition.
IP	Shows the raw size, user size, and reserved size for the IP CAM partition and its subpartitions. In Algorithmic mode, subpartitions are not displayed.
IPv6	Shows the raw size, user size, and reserved size for the IPv6 CAM partition and its subpartitions. In Algorithmic mode, subpartitions are not displayed.
IP VPN	Shows the raw size, user size, and reserved size for the IP VPN CAM partition and its subpartitions. In Algorithmic mode, subpartitions are not displayed.
MAC	Shows the raw size, user size, and reserved size for the MAC CAM partition and its subpartitions.
Session	Shows the raw size, user size, and reserved size for the session CAM partition and its subpartitions.
IPv6 Session	Shows the raw size, user size, and reserved size for the IPv6 session CAM partition and its subpartitions.
Out Session	Shows the raw size, user size, and reserved size for the out session CAM partition and its subpartitions.
Out IPv6 Session	Shows the raw size, user size, and reserved size for the out IPv6 session CAM partition and its subpartitions.
IP Section	Shows the CAM partition size of the IP section and its subnets. In Algorithmic mode, subnets are not displayed.
IPv6 Section	Shows the CAM partition size of the IPv6 section and its subnets. In Algorithmic mode, subnets are not displayed.
IP VPN Section	Shows the CAM partition size of the IP VPN section and its subnets. In Algorithmic mode, subnets are not displayed.
MAC Section	Shows the CAM partition size of the MAC section.
MAC Forwarding	Shows the CAM partition size of the MAC forwarding section.
MAC Flooding	Shows the CAM partition size of the MAC flooding section.
Misc Protocol	Shows the CAM partition size of the miscellaneous protocol section.
Session Section	Shows the CAM partition size of the session section.
IP Multicast	Shows the CAM partition size of the IP multicast ACL.
Broadcast ACL	Shows the CAM partition size of the IP broadcast ACL.
Receive ACL	Shows the CAM partition size of the IP receive ACL.
Rule-based ACL	Shows the CAM partition size of the rule-based ACL.
IPv6 Session Sec	Shows the CAM partition size of the IPv6 session section.

Displaying CAM partition for IPv6 VPN

The IPv6 VPN CAM partition is created when multi-service-3 or multi-service-4 CAM profile is configured. The IPv6 VPN CAM partition contains 10 sub partitions. The sub-partition is allocated with a fixed size, but can be dynamically changed. If the size of sub-partition is dynamically changed, the output from the **show cam-partition** command is affected. The following example displays information about IPv6 VPN CAM partition when the current CAM profile is multi-service-3:

```
device# show cam-partition
CAM partitioning profile: multi-service-3
Slot 1 XPP20SP 0:
# of CAM device           = 4
Total CAM Size            = 917504 entries (63Mbits)
.....
IPv6 VPN: Raw Size 131072, User Size 65536(0 reserved)
```

```

Subpartition 0: Raw Size 2048, User Size 1024, (0 reserved)
Subpartition 1: Raw Size 117734, User Size 58867, (0 reserved)
Subpartition 2: Raw Size 9333, User Size 4666, (0 reserved)
Subpartition 3: Raw Size 1285, User Size 642, (0 reserved)
Subpartition 4: Raw Size 384, User Size 192, (0 reserved)
.....
Slot 1 XPP20SP 0:
.....
IPv6 VPN Section: 524288 (080000) - 655359 (09ffff)
IPv6 VPN SNet 0: 524288 (080000) - 526335 (0807ff)
IPv6 VPN SNet 1: 526336 (080800) - 644069 (09d3e5)
IPv6 VPN SNet 2: 644070 (09d3e6) - 653402 (09f85a)
IPv6 VPN SNet 3: 653403 (09f85b) - 654687 (09fd5f)
IPv6 VPN SNet 4: 654688 (09fd60) - 655071 (09fedf)
IPv6 VPN SNet 5: 655072 (09fee0) - 655199 (09ff5f)
IPv6 VPN SNet 6: 655200 (09ff60) - 655263 (09ff9f)
IPv6 VPN SNet 7: 655264 (09ffa0) - 655295 (09ffbfbf)
IPv6 VPN SNet 8: 655296 (09ffc0) - 655327 (09ffdf)
IPv6 VPN SNet 9: 655328 (09ffe0) - 655359 (09ffff)

```

NOTE

In Algorithmic mode, the subpartitions and subnets are not displayed in the output of the **show cam-partition** command.

Output from show CAM partition usage command

The **show cam-partition usage** command shows the CAM size available per partition, the amount free, and the percent used. This information is shown here for slot 1.

```

device# show cam-partition usage
CAM partitioning profile: multi-service-3
Slot 1 XPP20SP 0:
Slot 1 XPP20SP 0:
    [IP]262144(size), 262129(free), 00.00%(used)
    :SNet 0: 2048(size), 2036(free), 00.58%(used)
    :SNet 1:237830(size), 237828(free), 00.00%(used)
    :SNet 2: 18667(size), 18667(free), 00.00%(used)
    :SNet 3: 2570(size), 2570(free), 00.00%(used)
    :SNet 4: 389(size), 389(free), 00.00%(used)
    :SNet 5: 128(size), 128(free), 00.00%(used)
    :SNet 6: 64(size), 64(free), 00.00%(used)
    :SNet 7: 32(size), 32(free), 00.00%(used)
    :SNet 8: 32(size), 32(free), 00.00%(used)
    :SNet 9: 32(size), 32(free), 00.00%(used)
    :SNet 10: 16(size), 16(free), 00.00%(used)
    :SNet 11: 16(size), 16(free), 00.00%(used)
    :SNet 12: 16(size), 16(free), 00.00%(used)
    :SNet 13: 16(size), 16(free), 00.00%(used)
    :SNet 14: 16(size), 16(free), 00.00%(used)
    :SNet 15: 16(size), 16(free), 00.00%(used)
    :SNet 16: 16(size), 16(free), 00.00%(used)
    :SNet 17: 16(size), 16(free), 00.00%(used)
    :SNet 18: 16(size), 16(free), 00.00%(used)
    :SNet 19: 16(size), 16(free), 00.00%(used)
    :SNet 20: 16(size), 16(free), 00.00%(used)
    :SNet 21: 16(size), 16(free), 00.00%(used)
    :SNet 22: 16(size), 16(free), 00.00%(used)
    :SNet 23: 16(size), 16(free), 00.00%(used)
    :SNet 24: 16(size), 16(free), 00.00%(used)
    :SNet 25: 16(size), 16(free), 00.00%(used)
    :SNet 26: 16(size), 16(free), 00.00%(used)
    :SNet 27: 16(size), 16(free), 00.00%(used)
    :SNet 28: 16(size), 16(free), 00.00%(used)
    :SNet 29: 16(size), 16(free), 00.00%(used)
    :SNet 30: 16(size), 16(free), 00.00%(used)
    :SNet 31: 16(size), 15(free), 06.25%(used)
    [IPV6] 32768(size), 32762(free), 00.01%(used)
    :SNet 0: 1024(size), 1022(free), 00.19%(used)
    :SNet 1: 28756(size), 28754(free), 00.00%(used)

```

```

:SNNet 2: 2332(size), 2332(free), 00.00%(used)
:SNNet 3: 320(size), 320(free), 00.00%(used)
:SNNet 4: 192(size), 192(free), 00.00%(used)
:SNNet 5: 64(size), 64(free), 00.00%(used)
:SNNet 6: 32(size), 32(free), 00.00%(used)
:SNNet 7: 16(size), 16(free), 00.00%(used)
:SNNet 8: 16(size), 15(free), 06.25%(used)
:SNNet 9: 16(size), 15(free), 06.25%(used)
[IP VPN]196608(size), 196532(free), 00.03%(used)
:SNNet 0: 2048(size), 1999(free), 02.39%(used)
:SNNet 1:177113(size), 177086(free), 00.01%(used)
:SNNet 2: 14000(size), 14000(free), 00.00%(used)
:SNNet 3: 1927(size), 1927(free), 00.00%(used)
:SNNet 4: 384(size), 384(free), 00.00%(used)
:SNNet 5: 128(size), 128(free), 00.00%(used)
:SNNet 6: 64(size), 64(free), 00.00%(used)
:SNNet 7: 32(size), 32(free), 00.00%(used)
:SNNet 8: 32(size), 32(free), 00.00%(used)
:SNNet 9: 32(size), 32(free), 00.00%(used)
:SNNet 10: 16(size), 16(free), 00.00%(used)
:SNNet 11: 16(size), 16(free), 00.00%(used)
:SNNet 12: 16(size), 16(free), 00.00%(used)
:SNNet 13: 16(size), 16(free), 00.00%(used)
:SNNet 14: 16(size), 16(free), 00.00%(used)
:SNNet 15: 16(size), 16(free), 00.00%(used)
:SNNet 16: 16(size), 16(free), 00.00%(used)
:SNNet 17: 16(size), 16(free), 00.00%(used)
:SNNet 18: 16(size), 16(free), 00.00%(used)
:SNNet 19: 16(size), 16(free), 00.00%(used)
:SNNet 20: 16(size), 16(free), 00.00%(used)
:SNNet 21: 16(size), 16(free), 00.00%(used)
:SNNet 22: 16(size), 16(free), 00.00%(used)
:SNNet 23: 16(size), 16(free), 00.00%(used)
:SNNet 24: 16(size), 16(free), 00.00%(used)
:SNNet 25: 16(size), 16(free), 00.00%(used)
:SNNet 26: 16(size), 16(free), 00.00%(used)
:SNNet 27: 16(size), 16(free), 00.00%(used)
:SNNet 28: 16(size), 16(free), 00.00%(used)
:SNNet 29: 16(size), 16(free), 00.00%(used)
:SNNet 30: 16(size), 16(free), 00.00%(used)
:SNNet 31: 512(size), 512(free), 00.00%(used)
[MAC]131072(size), 131061(free), 00.00%(used)
:Protocol: 10(size), 6(free), 40.00%(used)
:Forwarding:131054(size), 131047(free), 00.00%(used)
:Flooding: 8(size), 8(free), 00.00%(used)
[IPv6 VPN] 65536(size), 15(free), 99.97%(used)
:SNNet 0: 20(size), 0(free), 100.00%(used)
:SNNet 1: 65500(size), 0(free), 100.00%(used)
:SNNet 2: 2(size), 2(free), 00.00%(used)
:SNNet 3: 2(size), 2(free), 00.00%(used)
:SNNet 4: 2(size), 2(free), 00.00%(used)
:SNNet 5: 2(size), 2(free), 00.00%(used)
:SNNet 6: 2(size), 2(free), 00.00%(used)
:SNNet 7: 2(size), 2(free), 00.00%(used)
:SNNet 8: 2(size), 1(free), 50.00%(used)
:SNNet 9: 2(size), 2(free), 00.00%(used)
[Session] 32768(size), 32767(free), 00.00%(used)
:IP Multicast: 8192(size), 8192(free), 00.00%(used)
:Receive ACL: 1024(size), 1023(free), 00.09%(used)
:Rule ACL: 23552(size), 23552(free), 00.00%(used)
:IP Source Guard Permit: 0(size), 0(free), 00.00%(used)
:IP Source Guard Denial: 0(size), 0(free), 00.00%(used)
[IPv6 Session] 8192(size), 8192(free), 00.00%(used)
:IP Multicast: 2048(size), 2048(free), 00.00%(used)
:Receive ACL: 0(size), 0(free), 00.00%(used)
:Rule ACL: 6144(size), 6144(free), 00.00%(used)
[Internal Forwarding Lookup] 8192(size), 8185(free), 00.08%(used)
[Out Session] 28672(size), 28672(free), 00.00%(used)
[Out V6 Session] 8192(size), 8192(free), 00.00%(used)

```

The type of CAM partitioning profile configured is displayed in the "CAM partitioning profile line. The "multi-service-3" or "multi-service-4" profile indicates that the system will allocate a partition for IPv6 VPN.

The output displays the size of the available CAM, amount of CAM currently free, and what percentage of the available CAM is used currently.

(size): The effective user size obtained by subtracting the reserved size from the user size.

(free): The amount of CAM currently available.

(used): The percentage of CAM currently being used.

In Algorithmic mode, the subnets of IP, IP VPN, IPv5, and IPv6 VPN are not displayed in the output of the **show cam-partition usage** command.

```
device# show cam-partition usage
CAM partitioning profile: default

XPP100GEXE 0:

      [MAC] 229376(size), 229371(free), 0. 0%(used)
      :Protocol: 10(size), 5(free), 50. 0%(used)
      :Forwarding: 229354(size), 229354(free), 0. 0%(used)
      :Flooding: 8(size), 8(free), 0. 0%(used)
      :Port BUM RL: 4(size), 4(free), 0. 0%(used)

      [Session] 147456(size), 147455(free), 0. 0%(used)
:IP Source Guard Denial: 0(size), 0(free), 0. 0%(used)
:IP Source Guard Permit: 0(size), 0(free), 0. 0%(used)
      :Rule-based ACL: 138240(size), 138240(free), 0. 0%(used)
      :Broadcast ACL: 0(size), 0(free), 0. 0%(used)
      :Receive ACL: 1024(size), 1023(free), 0. 9%(used)
      :IP Multicast: 8192(size), 8192(free), 0. 0%(used)
      :IP Multicast 1G: 0(size), 0(free), 0. 0%(used)
      :IP Multicast 2GM: 0(size), 0(free), 0. 0%(used)
      :Open Flow CatchAll: 0(size), 0(free), 0. 0%(used)
:Open Flow UnProtected: 0(size), 0(free), 0. 0%(used)
      :Open Flow Normal: 0(size), 0(free), 0. 0%(used)
      :Open Flow Protected: 0(size), 0(free), 0. 0%(used)

[IPv6 Session] 12288(size), 12288(free), 0. 0%(used)
      :Rule ACL: 10240(size), 10240(free), 0. 0%(used)
      :Receive ACL: 0(size), 0(free), 0. 0%(used)
      :IPv6 Multicast: 2048(size), 2048(free), 0. 0%(used)
      :IPv6 Open Flow CatchAll: 0(size), 0(free), 0. 0%(used)
:IPv6 Open Flow UnProtected: 0(size), 0(free), 0. 0%(used)
      :IPv6 Open Flow Normal: 0(size), 0(free), 0. 0%(used)
      :IPv6 Open Flow Protected: 0(size), 0(free), 0. 0%(used)

[Out Session] 49152(size), 49152(free), 0. 0%(used)

[Out V6 Session] 4096(size), 4096(free), 0. 0%(used)

[Internal Forwarding Lookup] 131072(size), 131072(free), 0. 0%(used)
      :IFL Main: 131071(size), 131071(free), 0. 0%(used)
      :IFL Openflow CatchAll: 1(size), 1(free), 0. 0%(used)

      [IP]1458176(size), 1458165(free), 0. 0%(used)

      [IP VPN] 622592(size), 622592(free), 0. 0%(used)

      [IPv6] 425984(size), 425982(free), 0. 0%(used)

XPP100GEXE 1:
```

```

[MAC] 229376(size), 229371(free), 0. 0%(used)
:Protocol: 10(size), 5(free), 50. 0%(used)
:Forwarding: 229354(size), 229354(free), 0. 0%(used)
:Flooding: 8(size), 8(free), 0. 0%(used)
:Port BUM RL: 4(size), 4(free), 0. 0%(used)

[Session] 147456(size), 147455(free), 0. 0%(used)
:IP Source Guard Denial: 0(size), 0(free), 0. 0%(used)
:IP Source Guard Permit: 0(size), 0(free), 0. 0%(used)
:Rule-based ACL: 138240(size), 138240(free), 0. 0%(used)
:Broadcast ACL: 0(size), 0(free), 0. 0%(used)
:Receive ACL: 1024(size), 1023(free), 0. 9%(used)
:IP Multicast: 8192(size), 8192(free), 0. 0%(used)
:IP Multicast 1G: 0(size), 0(free), 0. 0%(used)
:IP Multicast 2GM: 0(size), 0(free), 0. 0%(used)
:Open Flow CatchAll: 0(size), 0(free), 0. 0%(used)
:Open Flow UnProtected: 0(size), 0(free), 0. 0%(used)
:Open Flow Normal: 0(size), 0(free), 0. 0%(used)
:Open Flow Protected: 0(size), 0(free), 0. 0%(used)

[IPv6 Session] 12288(size), 12288(free), 0. 0%(used)
:Rule ACL: 10240(size), 10240(free), 0. 0%(used)
:Receive ACL: 0(size), 0(free), 0. 0%(used)
:IPv6 Multicast: 2048(size), 2048(free), 0. 0%(used)
:IPv6 Open Flow CatchAll: 0(size), 0(free), 0. 0%(used)
:IPv6 Open Flow UnProtected: 0(size), 0(free), 0. 0%(used)
:IPv6 Open Flow Normal: 0(size), 0(free), 0. 0%(used)
:IPv6 Open Flow Protected: 0(size), 0(free), 0. 0%(used)

[Out Session] 49152(size), 49152(free), 0. 0%(used)

[Out V6 Session] 4096(size), 4096(free), 0. 0%(used)

[Internal Forwarding Lookup] 131072(size), 131072(free), 0. 0%(used)
:IFL Main: 131071(size), 131071(free), 0. 0%(used)
:IFL Openflow CatchAll: 1(size), 1(free), 0. 0%(used)

[IP]1458176(size), 1458165(free), 0. 0%(used)

[IP VPN] 622592(size), 622592(free), 0. 0%(used)

[IPv6] 425984(size), 425982(free), 0. 0%(used)

```

Syntax: `show cam-partition usage slot slot-number`

Displaying CAM information

The following commands display CAM information.

Show cam l2vpn

To display all VLL or VPLS MAC entries, including local entries (Port or VLAN or MAC from end points) and remote entries (VC or MAC from VLL or VPLS peers) enter the following command.

```

device# show cam l2vpn 2/1
Slot Index   MAC                Age  Port  IFL/  VC Label  Out Port Remote DA/  PRAM
      (Hex)                                     VLAN                                     SA  (Hex)
2    9fff6    0000.0034.5678  Dis  2/4   4096    74565    2/2     0    DA   8f
2    9fff7    0000.0034.5566  Dis  2/2   500     N/A     Filter   0    DA   8e

```

Syntax: `show cam l2vpn slot/port [MAC address]`

Show cam ipvpn

To display IPv4 VPN CAM entries, including local (Port+VLAN+IP) and remote (VC+IP) entries, enter the following command.

```
device# show cam ipvpn 2/1
Slot Index IP_Address Port Vlan VC Lbl MAC Age Out Vlan Out Port
2 0x60000 10.2.3.4/32 2/6 18 N/A N/A Dis 10 3/5
2 0x60001 224.7.8.9/32 N/A N/A 4660 0000.0080.0600 Dis 20 3/5
```

Syntax: `show cam ipvpn slot/port [IP prefix]`

Show cam l4

To display all CAM entries on a Layer 4 interface, enter the following command.

```
device# show cam l4 4/1
LP Index Src IP SPort Pro Age IFL/ Out IF Group PRAM
(Hex) (Dest IP DPort) VLAN Action (Hex)
4 a4000 0.0.0.0 0 17 Dis 0 CPU 31 00084
(10.0.0.0 3784 )
4 a4800 0.0.0.0 0 0 Dis 0 Pass 0 000c1
(10.9.4.255 0 )
4 a4802 0.0.0.0 0 0 Dis 0 Pass 0 000c2
(10.10.4.255 0 )
4 a4804 0.0.0.0 0 0 Dis 0 Pass 0 000c3
(10.33.33.255 0 )
4 a4806 0.0.0.0 0 0 Dis 0 Pass 0 000c4
(10.10.10.255 0 )
4 a4808 0.0.0.0 0 0 Dis 0 Pass 0 000c5
(10.20.20.255 0 )
4 a480a 0.0.0.0 0 0 Dis 0 Pass 0 000c6
(10.13.13.255 0 )
4 a480c 0.0.0.0 0 0 Dis 0 Pass 0 000c7
(10.41.41.255 0 )
4 a480e 0.0.0.0 0 0 Dis 0 Pass 0 000c8
(10.21.21.21 0 )
4 a4810 0.0.0.0 0 0 Dis 0 Pass 0 000c9
(10.55.55.255 0 )
```

Syntax: `show cam l4 slot/port`

The *slot/port* parameter specifies the port for which you want to display the CAM entries.

[Table 25](#) describes the output parameters of the **show cam l4 slot/port** command.

TABLE 25 Output parameters of the show cam l4 command

Field	Description
LP	Shows the number of the interface module.
Index (Hex)	Shows the row number of this entry in the IP route table.
Src IP Dest IP	Shows the source IP address and the destination IP address.
SPort DPort	Shows the source port ID and the destination port ID.
Pro	Shows the type of the protocol (TCP, UDP) used.
Age	Shows whether the age is enabled or disabled.
IFL/ VLAN	Shows the VLAN to which the port belongs.
Out IF Action	Shows the state of outgoing interface action.
Group	Shows the group address.
PRAM (Hex)	Shows the ACL PRAM entries.

Show cam label-out

To display Outbound Label ACL CAMs, enter the following command.

```
device# show cam label-out 2/1
Slot Index Port Outer Lbl Inner Lbl MAC Action
2 0xc0000 2/1 1024 1025 0000.0034.5678 Drop
2 0xc0002 2/1 1027 1028 0000.0034.5577 Drop
```

Syntax: `show cam label-out slot/port`

Show IFL CAM partition

To display information about the IFL CAM partition, enter the following command.

```
device# show cam ifl 2/1
Slot Index Port Outer VLAN Inner VLAN PRAM IFL ID
(Hex) (Hex)
2 00c5fff 2/1 100 200 185fff 4096
```

Syntax: `show cam ifl slot/port`

Show CAM IP

To display IP CAM information, enter the following command.

```
device# show cam ip 3/1
LP Index IP Address MAC Age IFL/ Out IF PRAM
(Hex) VLAN (Hex)
3 02fef(L) 10.33.32.0/32 N/A Dis N/A Drop 00094
3 02ff0(L) 10.33.32.255/32 N/A Dis N/A Mgmt 0009d
3 02ff1(L) 10.33.32.1/32 N/A Dis N/A Mgmt 0009c
3 02ff2(L) 10.11.11.0/32 N/A Dis N/A Drop 00094
3 02ff3(L) 10.11.11.255/32 N/A Dis N/A Mgmt 0009b
3 02ff4(L) 10.11.11.3/32 N/A Dis N/A Mgmt 0009a
3 02ff5(L) 10.5.5.5/32 N/A Dis N/A Mgmt 00096
3 02ff6(L) 224.0.0.22/32 N/A Dis N/A Mgmt 00093
3 02ff7(L) 224.0.0.18/32 N/A Dis N/A Mgmt 00092
3 02ff8(L) 224.0.0.13/32 N/A Dis N/A Mgmt 00091
3 02ff9(L) 224.0.0.9/32 N/A Dis N/A Mgmt 00090
3 02ffa(L) 224.0.0.6/32 N/A Dis N/A Mgmt 0008f
3 02ffb(L) 224.0.0.5/32 N/A Dis N/A Mgmt 0008e
3 02ffc(L) 224.0.0.4/32 N/A Dis N/A Mgmt 0008d
3 02ffd(L) 224.0.0.2/32 N/A Dis N/A Mgmt 0008c
3 02ffe(L) 224.0.0.1/32 N/A Dis N/A Mgmt 0008b
3 02fff(L) 10.255.255.255/32 N/A Dis N/A Mgmt 0008a
3 35488(R) 10.33.32.0/24 N/A Dis N/A CPU 0009f
3 35489(R) 10.11.11.0/24 N/A Dis N/A CPU 0009e
3 3548a(R) 10.5.5.5/32 N/A Dis N/A Drop 00094
3 3ffff(R) 0.0.0.0/0 N/A Dis N/A Drop 00094
```

Syntax: `show cam ip slot/port`

The *slot/port* parameter specifies the port for which you want to display the CAM entries.

[Table 26](#) describes the output parameters of the `show cam ip slot/port` command.

TABLE 26 Output parameters of the show cam ip command

Field	Description
LP	Shows the number of the interface module.
Index (Hex)	Shows the row number of this entry in the IP route table.
IP Address	Shows the IP address of the interface.

TABLE 26 Output parameters of the show cam ip command (continued)

Field	Description
MAC	Shows the MAC address of the interface.
Age	Shows whether the age is enabled or disabled.
IFL/ VLAN	Shows the VLAN to which the port belongs.
Out IF	Shows the state of outgoing interface action.
PRAM (Hex)	Shows the ACL PRAM entries.

Show CAM IPv6

To display IPv6 CAM information, enter the following command

```
device# show cam ipv6 3/20
LP Index IPV6 Address   MAC Age IFL/ Out IF PRAM
  (Hex)                                     VLAN      (Hex)
3 22ffc 2001:db8::/128   N/A Dis N/A Mgmt  000dc
3 22ffe 2001:db8::1/128 N/A Dis N/A Mgmt  000db
3 2e8a6 2001:db8::/64    N/A Dis N/A CPU   000dd
3 2ffde fe80::/10       N/A Dis N/A CPU   00086
3 2fffe ::/0            N/A Dis N/A Drop  00085
```

Syntax: `show cam ipv6 slot/port`

The *slot/port* parameter specifies the port for which you want to display the CAM entries.

[Table 27](#) describes the output parameters of the **show cam ipv6 slot/port** command.

TABLE 27 Output parameters of the show cam ipv6 command

Field	Description
LP	Shows the number of the interface module.
Index (Hex)	Shows the row number of this entry in the IP route table.
IPV6 Address	Shows the IPv6 address of the interface.
MAC	Shows the MAC address of the interface.
Age	Shows whether the age is enabled or disabled.
IFL/ VLAN	Shows the VLAN to which the port belongs.
Out IF	Shows the state of outgoing interface action.
PRAM (Hex)	Shows the ACL PRAM entries.

Displaying IPv6 VPN CAM information

The **show cam ipv6-vpn** command displays CAM information for an IPv6 VPN CAM entry on a single port, or for all ports on a device. IPv6 VPN CAM contains the destination IPv6 VPN address and layer 3 VPN ID. To display information for an IPv6 VPN CAM entry, enter the following command:

```
device# show cam ipv6-vpn 1/1
LP Index IPV6 VPN Address   MAC      Out IF   Age
  (Hex)      ID              IFL ID   PRAM
1  407f0 2001:db8:1::/128   N/A      Filter  Dis
                        (21847)  1d615)
1  407f2 2001:db8:2::/128   N/A      Drop    Dis
                        (21846)  5af6d)
```

Syntax: `show cam ipv6-vpn slot/port`

Show cam v6acl

The **show cam v6acl** command displays IPv6 ACL CAM sessions configured on the device. The VLAN column is expanded to display either VLAN or IFL ID as shown in the following example.

```
device# show cam v6acl 1/1
LP Index Src IP Addr          SPort IFL/VLAN ID
      Dst IP Addr          DPort Pro Age Out IF PRAM
1  74000 2001:db8:1::/64          0      536977
      2001:db8:2::/64          0      6 Dis Pass 000a4
1  74008 2001:db8:1::/64          0      536977
      2001:db8:2::/64          0      6 Dis Pass 000a5
1  74010 2001:db8:1::/64          0      536977
      2001:db8:2::/64          0      6 Dis Pass 000a6
1  74018 2001:db8:1::/64          0      536977
      2001:db8:2::/64          0      6 Dis Pass 000a7
1  74020 2001:db8:1::/64          0      536977
      2001:db8:2::/64          0      6 Dis Pass 000a8
1  74028 2001:db8:1::/64          0      536977
```

Syntax: **show cam ipv6-vpn slot/port**

Displaying IPv6 host drop CAM limit

Run the **show ipv6** command to display information about the IPv6 host drop CAM limit.

```
device# show ipv6
Global Settings
IPv6 Router-Id: 10.23.23.1 load-sharing path: 4
unicast-routing enabled, ipv6 allowed to run, hop-limit 64
reverse-path-check disabled
host drop cam limit 5
urpf-exclude-default disabled
session-logging-age 5
No Inbound Access List Set
No Outbound Access List Set
source-route disabled, forward-source-route disabled, icmp-redirect disabled
Configured Static Routes: 2
```

Syntax: **show ipv6**

Show IFL CAM ISID partition

To display information about 802.1AH for ISID, enter the following command:

```
device#show cam ifl-isid 1/1
Slot Index Port Outer VLAN Itag ISID PRAM IFL ID IPV4/V6
      (Hex)                               (Hex)      Routing
1  0085fe8 1/14 27 37 185fe8 1 0/0
1  0085fe9 1/13 26 36 185fe9 1 0/1
1  0085fea 1/16 25 35 185fea 1 1/0
1  0085feb 1/15 24 34 185feb 1 1/1
```

Syntax: **show cam ifl-isid slot/port**

This output includes an IPv4/ IPv6 Routing column. The IPv4/IPv6 Routing column indicates whether IPv4 or IPv6 is enabled or disabled on the interface. The number 1 represents enabled, and the number 0 represents disabled. For example, if 0/0 is displayed, then IPv4/IPv6 is disabled. If 0/1 is displayed, then IPv4 is disabled/ IPv6 is enabled. The IPv4/IPv6 Routing column is also displayed in the output of the **show cam ifl** command and **show cam ifl-mps** command.

Configuring CAM partition size

When you configure a tftp file size into the device, the device can only perform a parameter check based on the default CAM profile configured. In this situation, it is possible that you have configured a CAM partition size that conflicts with the physical CAM size. The following **system-max** commands may cause a conflict with the physical CAM size:

- system-max
 - ifl-cam
 - ip-source-guard-cam
 - ipv4-mcast-cam
 - ipv6-mcast-cam
 - lsp-out-acl-cam
 - subnet-broadcast-acl-cam
 - receive cam

When you have configured a CAM size that conflicts with the physical CAM size, a partition is created with the maximum possible CAM indices assigned to it. The following Syslog message is generated:

```
device# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 27 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
               I=informational N=notification W=warning
Dynamic Log Buffer (50 lines):
Sep  9 18:48:23:A:CAM IPv6 VPN SNet 9 partition warning: request 32, actual 0, slot 1, ppcr 0
```

CAM overflow logging

At system initialization, a threshold value is calculated for each sub-partition. If a partition does not have any sub-partitions, the value is based on the entire partition size. If a partition has movable sub-partition boundaries, the threshold value is also based on the entire partition size. By default, the threshold value is 5% of the total entry count. A minimum logging interval (default of 5 minutes) is also set for each partition to check usage. For example, let us say CAM overflow logging duration was set to 5 minutes and the overflow log is generated during a CAM write at 2:00 pm, then any further CAM writes will not cause an overflow log until 5 minutes have elapsed. So the next CAM overflow logging would occur on a CAM write after 2:05 pm. When the interval elapses, if the number of unused CAM entries drops below the threshold percentage value, a log message is generated during a CAM write.

```
CAM partition <partition name including sub-partition ID if applicable> warning: total <total count>
, free <current free count>
, slot <1 based slot number>
, ppcr <0 based ppcr id>
```

After the log message is generated, the sub-partition time stamp is updated to the current time.

Configuring minimum logging interval and threshold value

You can configure a minimum logging interval and threshold value for CAM partition logging using the following command.

```
device(config)# cam-partition logging 10% 5
```

Syntax: [no] cam-partition logging *threshold percentage %* | *interval in minutes*

You can configure the *threshold percentage %* variable to change the threshold value from the default 95%.

The *interval in minutes* variable allows you to set the minimum logging interval. Default 5 minutes.

NOTE

Because IP and IPv6 sub-partitions can dynamically grow and shrink, for these partitions, logging is implemented at the entire partition level. An SNMP trap is generated with the logging message.

Disabling CAM table entry aging

By default if no traffic hits a programmed flow-based content addressable memory (CAM) table entry, the CAM entry is removed from the system's CAM table. Depending on your network needs, however, you might have to disable the default behavior and force the system to retain CAM entries even when no traffic hits them. You can stop and start the CAM aging feature by using the **hw-aging** command in the global configuration mode.

Syntax: **hw-aging** **disable** | **enable**

The **disable** option prevents CAM entries from aging out. Even if no traffic hits a particular CAM entry, the entry remains in the CAM table.

The **enable** option returns the system to the default mode and unused CAM age out of the CAM table.

Data integrity protection

Data integrity protection provides a way to detect and report potential problems with the internal data path of the network processor. It also allows you to tune the detection and reporting of these types of problems. In addition, a show command is provided to display the status of the system.

Configuring detection parameters

Several parameters can be configured to support this data integrity protection: rolling window time frame, event thresholds for ingress and egress buffer events, and event thresholds for Control Static Random Access Memory (CSRAM) and Longest Prefix Match (LPM) memories. The configurations are applied system-wide.

Rolling Window Time Frame

Data integrity protection implements a rolling window to calculate the most recent history of errors. The rolling window time frame is the period of time error events are recorded. Data integrity protection polls for events every 500 milliseconds and updates the current window.

```
device(config)# system np rolling-window 10
```

Syntax: **[no] system np rolling-window** *window size*

The *window size* parameter sets the rolling window time frame. The allowable window time is 10 to 60 seconds. Setting to 0 seconds will disable error monitoring.

The **no** form of the command returns the threshold to the default setting.

Event Threshold Configuration

The data integrity protection implements configurable thresholds for generating a syslog and trap. There is one threshold for ingress buffer events and one threshold for egress buffer events. Once crossed, a syslog and trap will be generated.

To prevent excessive log and traps there is a 10 minute period before another syslog or trap is generated. Setting a threshold to zero disables error detection for the monitor point on all network processors.

The default threshold values are different for ingress and egress. The ingress error count is based on the errors detected on each 32-bit word. The egress error count is based on the number of packets with one or more errors.

The **system np ingress-threshold** command configures the ingress buffer error reporting threshold.

```
device(config)# system np ingress-threshold 20
```

Syntax: **[no] system np ingress-threshold** *threshold*

The *threshold* range is 0 to 120 events. Setting the threshold to 0 disables the monitor point for all network processors. The default setting is 20 events.

The **no** form of the command returns the threshold to default.

The **system np egress-threshold** command configures the egress buffer error reporting threshold.

```
device(config)# system np egress-threshold 20
```

Syntax: **[no] system np egress-threshold** *threshold*

The *threshold* range is 0 to 120 events. Setting the threshold to 0 disables the monitor point for all network processors. The default setting is 3 events.

The **no** form of the command returns the threshold to default.

Configuring the threshold parameters for CSRAM

Use the **system np control-ram-threshold** *threshold* command to configure the CSRAM threshold parameter when monitoring low level memory events occurring within the CSRAM memory module of the network processor.

Configure the Rolling Window Time Frame. Refer to the [Rolling Window Time Frame](#) on page 85.

NOTE

Configuring the threshold parameters for CSRAM is supported only on the CER 2000 Series and the CES 2000 Series platforms.

1. In privileged EXEC mode, issue the **configure terminal** command to enter the global configuration mode.
2. Enter the **system np** command to configure the NP error reporting threshold parameter on the device.

```
device (config)#system np
```

3. Enter the **system np control** command to configure the CSRAM error reporting threshold parameter on the device.

```
device (config)#system np control
control-ram-threshold  Configure the Control SRAM error reporting threshold
```

4. Enter the **system np control-ram-threshold** *threshold* command with the appropriate threshold parameter.

```
device (config)#system np control-ram-threshold 20
```

5. Enter the **system np control-ram-threshold 0** command to disable the monitoring of low level memory events.

```
device (config)#system np control-ram-threshold 0
```

6. Enter the **no system np control-ram-threshold** *threshold* command to reset the threshold value to default.

```
device (config)#no system np control-ram-threshold 20
```

NOTE

By default, the feature is enabled with default configuration values. When the default configuration values are present, the **show run** command does not display CSRAM error configuration information. The **show run** command displays CSRAM error configuration information only when a non-default value is configured or when the feature is disabled.

The following example configures the CSRAM error reporting threshold parameter to 20 events on the device.

```
device(config)#system np
control-ram-threshold  Configure the Control SRAM error reporting threshold
egress-threshold       Configure the egress buffer error reporting threshold
ingress-threshold      Configure the ingress buffer error reporting threshold
lpm-ram-threshold      Configure the LPM memory error reporting threshold
rolling-window         Configure the rolling window time frame
device(config)#system np control
control-ram-threshold  Configure the Control SRAM error reporting threshold
device(config)#system np control-ram-threshold 20
DECIMAL 0-120 events (default:10 disable: 0)
device(config)#system np control-ram-threshold 20
```

Configuring the threshold parameters for LPM memory

Use the **system np lpm-ram-threshold threshold** command to configure the LPM memory threshold parameter when monitoring low level memory events occurring within the LPM memory module of the network processor.

Configure the Rolling Window Time Frame. Refer to the [Rolling Window Time Frame](#) on page 85.

NOTE

Configuring the threshold parameters for LPM memory is supported only on the CER 2000 Series and the CES 2000 Series platforms.

1. In privileged EXEC mode, issue the **configure terminal** command to enter the global configuration mode.
2. Enter the **system np** command to configure the NP error reporting threshold parameter on the device.

```
device (config)#system np
```

3. Enter the **system np lpm** command to configure the LPM memory error reporting threshold parameter on the device.

```
device (config)#system np lpm
lpm-ram-threshold  Configure the LPM memory error reporting threshold
```

4. Enter the **system np lpm-ram-threshold threshold** command with the appropriate threshold parameter.

```
device (config)#system np lpm-ram-threshold 30
```

5. Enter the **system np lpm-ram-threshold 0** command to disable the monitoring of low level memory events.

```
device (config)#system np lpm-ram-threshold 0
```

6. Enter the **no system np lpm-ram-threshold threshold** command to reset the threshold value to default.

```
device (config)#no system np lpm-ram-threshold 30
```

NOTE

By default, the feature is enabled with default configuration values. When the default configuration values are present, the **show run** command does not display LPM memory error configuration information. The **show run** command displays LPM memory error configuration information only when a non-default value is configured or when the feature is disabled.

The following example configures the LPM memory error reporting threshold parameter to 20 events on the device.

```
device(config)#system np
control-ram-threshold  Configure the Control SRAM error reporting threshold
egress-threshold        Configure the egress buffer error reporting threshold
ingress-threshold       Configure the ingress buffer error reporting threshold
lpm-ram-threshold       Configure the LPM memory error reporting threshold
rolling-window          Configure the rolling window time frame
device(config)#system np lpm
lpm-ram-threshold       Configure the LPM memory error reporting threshold
device(config)#system np lpm-ram-threshold 20
DECIMAL 0-120 events (default:10 disable: 0)
device(config)#system np lpm-ram-threshold 20
```

Showing Status

The **show np buffer-errors** command displays the count of error events for the rolling window.

```
device# show np buffer-errors
```

Ports	Ingress		Egress	
	Current	Cumulative	Current	Cumulative
1/1- 1/24	15	37	0	0
1/25 - 1/48	0	0	0	0
2/1 - 2/2	0	0	0	0

Syntax: show np buffer-errors

Displaying CSRAM error statistics

You can display the CSRAM error statistics information for the configured ports. Each port range corresponds to one network processor.

Use the **show np control-ram-errors** command to display CSRAM error statistics for the packet processor.

```
device# show np control-ram-errors
CSRAM
Ports      Current Cumulative
1/1 - 1/24      0         0
2/1 - 2/2      0         0
```

The following information is displayed:

- The ports that are configured with CSRAM threshold parameters.
- The current column displays the number of errors recorded in the rolling window time frame. The rolling window time frame records the most recent number of errors. Refer to the [Rolling Window Time Frame](#) on page 85 for more information. The threshold configuration parameter is applied on the current value of the recorded error events.
- The cumulative column displays the number of errors recorded from the time the feature was first enabled.

Displaying LPM memory error statistics

You can display the LPM memory error statistics information for the configured ports. Each port range corresponds to one network processor.

The **show np lpm-ram-errors** command displays LPM memory error statistics for the packet processor.

```
device#show np lpm-ram-errors
```

LPM 0 Ports	LPM 1		LPM 2			
	Current	Cumulative	Current	Cumulative	Current	Cumulative
1/1 - 1/24	0	3	0	3	0	3
2/1 - 2/2	0	3	0	3	0	3

The following information is displayed:

- The ports that are configured with LPM memory threshold parameters.
- The current column displays the number of errors recorded in the rolling window time frame. The rolling window time frame records the most recent number of errors. Refer to the [Rolling Window Time Frame](#) on page 85 for more information. The threshold configuration parameter is applied on the current value of the recorded error events.
- The cumulative column displays the number of errors recorded from the time the feature was first enabled.

Port transition hold timer

Using the **delay-link-event** command will delay the sending of port "up" or "down" events to Layer 2 protocols. While link down events are reported immediately in syslog, their effect on higher level protocols such as OSPF is delayed according to how the delay-link-event is configured. This command affects the physical link events. However, the resulting logical link events are also delayed. This is a per-interface command.

NOTE

When a Layer 2 protocol packet is received before the delay-link-event is expired, NetIron will reply to the received Layer 2 protocol without the delay-link-event. After the delay-link-event is expired, NetIron will retransmit the previous Layer 2 event.

For example, if VSRP is enabled on the port, the ownership will not change until the port status has remained up or down for the configured amount of time to ensure that minor transient states of a link do not unintentionally cause a disruptive topology change in the network.

NOTE

All LAG ports must have the same delayed-link-down-event configuration.

The following command will delay the sending of port "down" event for 100 ms when a port state is detected "down". If the port state is detected "up" afterwards within 100 ms, the delayed "down" event is cancelled; otherwise, the "down" event is sent after 100 ms. This allows the upper layer applications not to be affected by a port state flapping.

```
device (config-if-e1000-1/2)# delay-link-event 2 down
```

Syntax: [no] delay-link-event time up | down

The **time** parameter is the number of 50-ms units. The default is 0. The valid range is from 0 to 200.

The **up** parameter means only "up" events are delayed.

The **down** parameter means that only the down events are delayed.

If neither the **up** or **down** parameter is specified, both up and down events are delayed. This is the default.

Port flap dampening

Port flap dampening allows you to configure a wait period before a port, whose link goes down then up, becomes enabled.

If the port link state toggles, from down to up or from up to down, for a specified number of times within a specified period, the interface is physically disabled for the specified wait period. Once the wait period expires, the port's link state is re-enabled. However, if the wait period is set to zero (0) seconds, or you want to re-enable the port before the wait period expires, the port must be manually re-enabled.

Configuring port link dampening on an interface

This feature is configured at the interface level.

```
device(config)#interface ethernet 2/1
device(config-if-e10000-2/1)#link-error-disable 10 3 10
```

Syntax: [no] link-error-disable toggle-threshold sampling-time-in-sec wait-time-in-sec

The **toggle-threshold** is the number of times a port's link state goes from up to down and down to up before the wait period is activated. The default is 0. Enter a valid value range from 1-50.

The **sampling-time-in-sec** is the amount of time during which the specified toggle threshold can occur before the wait period is activated. The default is 0 seconds. Enter a value between 1 and 65565 seconds.

The **wait-time-in-sec** is the amount of time the port remains disabled (down) before it becomes enabled. Entering 0 indicates that the port will stay down until an administrative override occurs. Enter a value between 0 and 65565 seconds.

Configuring port link dampening on a LAG

You can configure the port link dampening feature on the primary port of a LAG at the interface level using the **link-error-disable** command. Once configured on the primary port of the LAG, the feature is enabled on all port that are members of the LAG. You cannot disable the feature from a member of the LAG.

Enter commands such as the following on the primary port of a LAG.

```
device(config)#interface ethernet 2/1
device(config-if-e10000-2/1)#link-error-disable 10 3 10
```

Re-enabling a port disabled by port link dampening

A port disabled by the port link dampening is automatically re-enabled once the wait period expires; however, if the wait period is set to zero (0) seconds or you want to re-enable the port before the configured wait period expires, you must re-enable the port by entering the **link-error-disable** command on the disabled port as shown in the following.

```
device(config)#interface ethernet 2/1
device(config-if-e10000-2/1)#link-error-disable 10 3 10
```

NOTE

You must enter the **link-error-disable** command with the **toggle-threshold**, **sampling-time-in-sec**, and **wait-time-in-sec** variables defined to re-enable the port. Using the **link-error-disable** command without the variables, will not bring the port back up.

Displaying ports configured with port link dampening

Ports that have been disabled due to the port link dampening feature are not identified in a **show running-config** command.

Use the **show interface link-error-disable** command to display the ports that have the port link dampening feature enabled.

```
device(config-if-e10000-8/1)#show interfaces link-error-disable
Port 8/1: link-error-disabled (Config: 2 toggles per 3 sec, wait time 1 sec)
Port 8/3: not link-error-disabled (Config: 2 toggles per 2 sec, wait time 30 sec)
Port 8/4: not link-error-disabled (Config: 2 toggles per 2 sec, wait time 30 sec)
```

TABLE 28 Output for the show interface link-error-disable command

Field	Description
port	The port that has been configured
link-error-disabled	The port that has been disabled by this feature
not link-error-disabled	The "not" means the port has not been disabled due to this feature
toggle	The number of times a port's link state goes from up to down and down to up before the wait period is activated
wait time	The amount of time the port remains disabled (down) before it becomes enabled

Issuing the **disabled-only** option with the command displays only the ports that have been disabled by the port link dampening feature.

```
device(config-if-e10000-8/1)#show interfaces link-error-disable disabled-only
Port 8/1: link-error-disabled (Config: 2 toggles per 3 sec, wait time 1 sec)
```

Syntax: **show interface link-error-disable [disabled-only]**

Entering the **show interface link-error-disable** command displays all the ports that have the port link dampening feature enabled. Add the **disabled-only** keyword for a list of ports disabled by this feature.

Port loop detection

This feature allows the Extreme device to disable a port that is on the receiving end of a loop by sending test packets. You can configure the time period during which test packets are sent.

Strict Mode and Loose Mode

There are two types of loop detection; Strict Mode and Loose Mode. In Strict Mode, a port is disabled only if a packet is looped back to that same port. Strict Mode overcomes specific hardware issues where packets are echoed back to the input port. In Strict Mode, loop detection must be configured on the physical port.

In Loose Mode, loop detection is configured on the VLAN of the receiving port. Loose Mode disables the receiving port if packets originate from any port or VLAN on the same device. The VLAN of the receiving port must be configured for loop detection in order to disable the port.

Recovering disabled ports

Once a loop is detected on a port, it is placed in a disabled state. The port will remain disabled until one of the following occurs:

- You manually disable and enable the port at the interface level of the CLI.
- You enter the **clear loop-detection** command. The **clear loop-detection** command clears the loop detection statistics and enables all disabled ports.
- The device automatically re-enables the port. To set your device to automatically re-enable disabled ports, refer to [Configuring the device to automatically re-enable ports](#) on page 93.

Disable duration and loop detection interval

By default, the ports are shutdown permanently until user enables it manually. You can configure the disable duration from 1 minute to 1440 minutes (24 hours)

By default, the Loop Detection time Interval between the loop detection BPDU is 1 second. You can configure the loop detection PDU interval from 100 ms to 10 seconds.

Configuration notes

Loopback detection packets are sent and received on both tagged and untagged ports. Therefore, this feature cannot be used to detect a loop across separate devices.

The following information applies to Loose Mode loop detection:

- Loop detection is configured on the VLAN. Different VLANs may disable different ports. A disabled port affects every VLAN using it.
- Loose Mode disables the receiving port if packets originate from any port or member port of a VLAN on the same device
- The VLAN of the receiving port must be configured for loop detection in order to disable the port.
- Loose Mode floods test packets to the entire VLAN. This can impact system performance if too many VLANs are configured for Loose Mode loop detection.

The following information applies to Strict Mode loop detection:

- A port is disabled only if a packet is looped back to that same port.
- Loop detection must be configured on the physical port.
- Strict Mode overcomes specific hardware issues where packets are echoed back to the input port.

NOTE

Extreme recommends that you limit the use of Loose Mode. If you have a large number of VLANs or VLAN groups, configuring loop detection on all of them can significantly affect system performance because of the flooding of test packets to all configured VLANs. An alternative to configuring loop detection in a VLAN-group of many VLANs is to configure a separate VLAN with the same tagged port and configuration, and enable loop detection on this VLAN only.

NOTE

When loop detection is used with Layer 2 loop prevention protocols, such as Spanning Tree Protocol (STP), the Layer 2 protocol takes higher priority. Loop detection cannot send or receive probe packets if ports are blocked by Layer 2 protocols, so it does not detect Layer 2 loops when STP is running because loops within a VLAN have been prevented by STP. Loop detection running in Loose Mode can detect and break Layer 3 loops because STP cannot prevent loops across different VLANs. In these instances, the ports are not blocked and loop detection is able to send out probe packets in one VLAN and receive packets in another VLAN. In this way, loop detection running in Loose Mode disables both ingress and egress ports.

Enabling loop detection

Use the loop-detection command to enable loop detection on a physical port (Strict Mode) or a VLAN (Loose Mode). Loop detection is disabled by default. The following example shows a Strict Mode configuration.

```
device(config)#interface ethernet 1/1
device(config-if-e1000-1/1)#loop-detection
```

The following example shows a Loose Mode configuration.

```
device(config)#vlan 20
device(config-vlan-20)#loop-detection
```

The following example shows a Loose Mode configuration for a VLAN group.

```
device(config)#vlan-group 10
device(config-vlan-group-10)#add-vlan 1 to 100
device(config-vlan-group-10)#loop-detection
```

By default, the port will send test packets every one second, or the number of seconds specified by the **loop-detection-interval** command. Refer to [Configuring a global loop detection interval](#) on page 93.

Syntax: **[no] loop-detection**

Use the **no** form of the command to disable loop detection.

Configuring a global loop detection interval

The loop detection interval specifies how often a test packet is sent on a port. When loop detection is enabled, the loop detection time unit is 0.1 second, with a default of 10 (one second). The range is from 1 (one tenth of a second) to 100 (10 seconds). You can use the **show loop-detection status** command to view the loop detection interval.

To configure the global loop detection interval, enter a command such as the following.

```
device(config)#loop-detection-interval 50
```

This command sets the loop-detection interval to 5 seconds (50 x 100 ms).

To revert to the default global loop detection interval of 10, enter one of the following.

```
device(config)#loop-detection-interval 10
```

or

```
device(config)#no loop-detection-interval 50
```

Syntax: **[no] loop-detection-interval** *number*

Where *number* is a value from 1 to 100. The system multiplies your entry by 0.1 to calculate the interval at which test packets will be sent.

Configuring the device to automatically re-enable ports

To configure the Extreme device to automatically re-enable ports that were disabled because of a loop detection, enter the following command. The default is 0.

```
device(config)#loop-detection disable-duration 1440
```

The above command will cause the Extreme device to automatically re-enable ports that were disabled for a duration of 24 hours because of a loop detection. This configuration applies to all the ports that are configured the loop detection (strict or loose).

Syntax: **[no] loop-detection disable-duration** *num*

Use the **no** form of the command to disable this feature.

Where *num* is the number of minutes from 0 to 1440. When 0 is specified, it is permanently off.

Clearing loop detection

To clear loop detection statistics and re-enable all ports that are in disabled state because of a loop detection, enter the following command.

```
device #clear loop-detection
```

Syntax: `clear loop-detection [vlan | ethernet] vlanid/port-num`

Where *port-num* enables the specified port.

Where *vlan-id* enables all the ports disabled by loop detection for this VLAN

Displaying loop detection information

Use the **show loop-detection** command to display the loop detection status.

```
device(config-vlan-100)#show loop-detection
loop detection packets interval: 10 (unit 100 msec)
loop detection disable duration: 10 (In minutes, 0 means permanently disabled)
Ports mode loop detection
=====
port-num    disable-count
1/12        0
1/11        0
Vlan mode loop detection
=====
vlan-id     disable-count
100         2
10          0
200         0
Ports disabled by loop detection
=====
port        age(minutes)  disable cause
1/11 1          Disabled by VLAN: 100 loopdetect 1/11
1/12 1          Disabled by VLAN: 100 loopdetect 1/12
```

Syntax: `show loop-detection`

TABLE 29 Port loop detection output description

Parameter	Description
loop detection packets interval	Specifies how often a test packet is sent on a port.
loop detection disable duration	Specifies the device to automatically re-enable ports that were disabled for the configured duration because of a loop detection
ports mode	The VLAN or port that port loop detection was configured on.
loop detection disabled ports	<p>The ports that are disabled by port loop detection.</p> <ul style="list-style-type: none"> port - The port number that was disabled by port loop detection. age - The time duration after which port will be automatically re-enabled. If the age is "0", it means port is not configured to be automatically re-enabled. disable cause - Specifies all the ports that were disabled by loop detection (either strict or loose).

Discarding loop detection frames in the LACP-blocked port

When loop detection is enabled and a loop is detected in the network, the looped packet port is disabled.

In a dynamic LAG scenario on a trunk, loop detection frames are sent out on the active primary port of a trunk group. A packet received in the LACP-blocked port of the transmitting port triggers loop detection on the trunk. Loop detection discards the loop detection frames received in the LACP-blocked port, keeps the port in the up state, and prevents the entire LAG from shutting down.

Syslog message

The following message is logged when a port is disabled due to loop detection. This message will also appear on the console.

```
SYSLOG: Jan 27 18:16:42:<14>Jan 27 18:16:42 LOOP_DETECT LOG: Port Down 1/10 - Loop detected on VLAN: 150
```

Displaying information for an interface for an Ethernet port

To display information for a show interface for an ethernet port, enter the following command at any CLI level.

```
device# show interface ethernet 9/1
GigabitEthernet2/3 is up, line protocol is up
  STP Root Guard is disabled, STP BPDU Guard is disabled
  Hardware is GigabitEthernet, address is 0000.0098.4900 (bia 0000.0098.492a)
  Configured speed auto, actual 1Gbit, configured duplex fdx, actual fdx
  Member of VLAN 1 (untagged), 5 L2 VLANs (tagged), port is in dual mode (default      vlan), 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
  arp-inspection-trust configured to OFF
  mirror disabled, monitor disabled
  Not member of any active trunks
  Not member of any configured trunks
  Port name is ->7.blade1.shelf1.access.aprd
    Port state change time: Jan 21 02:40:21, (0 days, 00:07:16 ago)
  MTU 1544 bytes, encapsulation ethernet
  300 second input rate: 1509512 bits/sec, 713 packets/sec, 0.15% utilization
  300 second output rate: 1992071 bits/sec, 751 packets/sec, 0.20% utilization
  712896623 packets input, 204984611768 bytes, 0 no buffer
  Received 1315502 broadcasts, 53313 multicasts, 711527808 unicasts
  0 input errors, 0 CRC, 0 frame, 0 ignored
  0 runts, 29433839 giants
  NP received 712896745 packets, Sent to TM 712839428 packets
  NP Ingress dropped 57317 packets
  796106728 packets output, 366570033985 bytes, 0 underruns
  Transmitted 2045784 broadcasts, 32330616 multicasts, 761730328 unicasts
  0 output errors, 0 collisions
  NP transmitted 796106833 packets, Received from TM 796534170 packets
```

Syntax: `show interface [ethernet slot-port [to slot-port]]`

You can display information for all ports in a device by using the **show interface** command without options, or use the **ethernet slot-port** option to limit the display to a single port, or add the **to slot-port** option for a range of ports.

Displaying the full port name for an Ethernet interface

To display the full port name for an ethernet interface using the CLI, enter the following command.

```
device# show interface brief slot 3
Port  Link Port-State  Dupl Speed Trunk Tag Priori MAC      Name      Type
```

```

3/1   Up    Forward   Full 100G   None No   level0 0000.0002.025c default-port
3/2   Up    Forward   Full 100G   None No   level0 0000.0002.025d default-port

```

Syntax: show interface brief slot/port

If the port is logically UP (meaning not LK-DISABLE or LACP-BLOCKED or OAM-DISABLE or DOT1X-BLOCKED), then:

- If the port is untagged then the L2 Port state field indicates the STP State of Port in the untagged VLAN context.
- If the port is tagged or in dual mode (both tagged and untagged), then it is marked forwarding as a single port state cannot be determined.

In case Port is logically down, L2 Port State indicates the reason for Logical Port down condition (LK-DISABLE or LACP-BLOCKED or OAM-DISABLE or DOT1X-BLOCKED)

Using the **show interface brief wide** command long port names are displayed. If the **show interface brief wide** command is not used only partial names are displayed in cases of long port names.

```

device# show interface brief wide
Port  Link Port-State  Speed Tag MAC          Name
2/1   Up    Forward   10G   No  0000.00f7.0230  port-connected-to-chicago
2/2   DisabNone        None  No  0000.00f7.0231
2/3   DisabNone        None  No  0000.00f7.0232
2/4   DisabNone        None  No  0000.00f7.0233
Port  Link Port-State  Speed Tag MAC          Name
mgmt1 Up    Forward   100M  Yes 0000.00f7.0200
Port  Link Port-State  Speed Tag MAC          Name
lb1   Up    N/A       N/A   N/A N/A
device#

```

Syntax: show interface brief wide slot/port**TABLE 30** Display of show interface ethernet port

Field	Description
<i>Module type</i> port# is <i>state</i>	The <i>module type</i> variable specifies a type of interface module, such as 10GigabitEthernet. The <i>port#</i> variable specifies the port number for the interface module. The <i>state</i> variable if the interface module is up or down.
Line protocol is <i>status</i>	The <i>status</i> variable specifies if the line protocol is up or down. If the interface is down due to Remote Fault, the reason is indicated as: "(remote fault)". If a port is down because of a Local Fault, the reason is indicated as: "(local fault)".
STP Root Guard is <i>status</i>	The <i>status</i> variable specifies if the STP Root Guard is enabled or disabled.
STP BPDU Guard is <i>status</i>	The <i>status</i> variable specifies if the STP BPDU Guard is enabled or disabled.
Hardware is <i>module type</i>	The <i>module type</i> variable specifies a type of interface module, such as # GigabitEthernet.
Address is <i>MAC- address</i>	The <i>MAC- address</i> variable specifies the MAC address of the port.
Configured speed and actual speed	The speed that the module has been configured to operate at, and the actual speed it is currently operating at.
Configured port speed and actual duplex	The port capacity that the module has been configured to operate at, and the actual speed it is currently operating at.
Member of VLAN # (untagged) <i>port#</i> L2 VLANS (tagged) Port is in <i>dual mode/untagged/tagged</i> mode Port state is <i>status</i>	The <i>VLAN#</i> (untagged) variable specifies a port that is a member of only 1 VLAN. The <i>port#</i> L2 VLANS (tagged) variable specifies a port that is a member of multiple ports and untagged. A port is in <i>dual- mode</i> specifies member VLAN ports as untagged and tagged. The default mode is dual-mode. The <i>status</i> variable identifies the flow of traffic as forwarding or disabled.
STP configured to <i>status</i>	The <i>status</i> variable specifies if the STP is ON or OFF.

TABLE 30 Display of show interface ethernet port (continued)

Field	Description
Priority level Flow control <i>status</i>	The priority level assigned to the port-based VLAN. The priority level is on scale from 0-7. The default is 0. The <i>status</i> variable is enabled or disabled.
Priority force <i>status</i>	The <i>status</i> variable specifies if the priority force on a port is disabled or enabled.
Drop precedence level <i>value</i>	Identifies the TOS or DSCP value in the IPv4 or IPv6 packet header. The <i>value</i> variable specifies the drop precedence on a scale from 0-3. Packets that contain a DSCP value of 0 are least likely to be dropped and packets with a value of 3 are most likely to be dropped. The default value is 0.
Drop precedence force <i>status</i>	The <i>status</i> variable specifies the drop precedence force as enabled or disabled. Identifies the drop precedence if the force command is configured for a specific ingress port.
arp-inspection-trust configured to <i>status</i>	The <i>status</i> variable specifies if arp-inspection-trust feature is configured ON or OFF. The default trust setting for a port is untrusted.
Mirror <i>status</i>	The <i>status</i> variable specifies if the port mirror command is configured as enabled or disabled.
Monitor <i>status</i>	The <i>status</i> variable specifies if the port monitor command is configured as enabled or disabled.
Trunk membership	The <i>Trunk membership</i> variable identifies the interface module as a member of a primary or secondary port. This specifies members of an active port or not a member of an active port.
Configured trunk membership	The <i>Configured trunk membership</i> variable identifies the interface module as a member of any configured trunk or not a member of a configured trunk.
Port name	The <i>port name</i> variable identifies the name assigned to the port.
MTU # bytes , encapsulation ethernet	Maximum Transmission Unit (MTU) refers to the size of the largest packet or frame that a known layer can pass forward. The # bytes variable refers to size of the packet or frame.
# seconds input rate: <i>value</i> bits/sec, <i>value</i> packets/sec, % utilization	The # seconds input rate refers to: <ul style="list-style-type: none"> The <i>value</i> of bits received per second. The <i>value</i> of packets received per second. The % utilization specifies the port's bandwidth used by received traffic.
# seconds output rate: <i>value</i> bits/sec, <i>value</i> packets/sec, % utilization	The # seconds output rate refers to: <ul style="list-style-type: none"> The <i>value</i> of bits transmitted per second. The <i>value</i> of packets transmitted per second. The % utilization specifies the port's bandwidth used by transmitted traffic.
<i>value</i> packets input, <i>value</i> bytes, <i>value</i> no buffer	<ul style="list-style-type: none"> The <i>value</i> variable specifies the number of packets received. The <i>value</i> variable specifies the number of bytes received. The <i>value</i> no buffer variable specifies the total number of packets that have been discarded by the MAC device, due to temporary inability to store the packets before forwarding to the Network Processor (NP).
Received <i>value</i> broadcasts, <i>value</i> multicasts, <i>value</i> unicasts	The <i>value</i> variable specifies the amount of traffic the interface module is receiving on broadcasts, multicasts, and unicast traffic.
<i>value</i> input errors, <i>value</i> CRC, <i>value</i> frame, <i>value</i> ignored	<ul style="list-style-type: none"> The <i>value</i> input errors variable specifies the number of received packets with errors. The <i>value</i> CRC variable specifies the number of packets discarded by the MAC device due to detected CRC error. The <i>value</i> variable specifies the number of received packets with alignment errors. The <i>value</i> variable specifies the number of received packets that are discarded. These parameters are not currently supported and will always display 0.
<i>value</i> runts, <i>value</i> giants	The <i>value</i> runts variable specifies the number of small packets that are less than 64 bytes. The <i>value</i> giants variable specifies the number of large packets greater than 1518 bytes. These parameters are not currently supported and will always display 0.

TABLE 30 Display of show interface ethernet port (continued)

Field	Description
NP received	The number of packets received on the Network Processor (NP).
NP transmitted	The number of packets sent from the Network Processor to the Traffic Manager (TM).
NP Ingress dropped	The number of ingress packets dropped on the Network Processor.
<i>value</i> packets output <i>value</i> bytes	<ul style="list-style-type: none"> The <i>value</i> variable specifies the number of transmitted packets. The <i>value</i> variable specifies the number of transmitted bytes.
Transmitted <i>value</i> broadcasts, <i>value</i> multicasts, <i>value</i> unicasts	The <i>value</i> variable specifies the amount of traffic the interface module transmitted on broadcasts, multicasts, and unicast traffic.
<i>value</i> output errors, <i>value</i> collisions	<ul style="list-style-type: none"> The <i>value</i> variable specifies the number of transmitted packets with errors. The <i>value</i> variable specifies the number of packets that experienced multi-access collisions. <p>These parameters are not currently supported and will always display 0.</p>
Network Processor transmitted <i>value</i> packets	The <i>value</i> variable specifies the number of packets transmitted from the Network Processor.
Received from Traffic Manager <i>value</i> packets	The <i>value</i> variable specifies the number of packets received by the Network Processor from the Traffic Manager.

Displaying statistics information for an Ethernet port

You can view statistical information about the traffic passing through a specified Ethernet port in one of two ways. The **monitor** commands allow you to monitor traffic statistics in real time, while the **show statistics** command provides a snapshot of the most recent traffic statistics.

Monitoring Ethernet port statistics in real time

You can monitor Ethernet traffic statistics in real time for a single port or traffic counters for all Ethernet ports using the **monitor** commands. When you execute a **monitor** command it retrieves and displays traffic statistics once per polling interval (2 seconds by default) until you pause or stop the display. The terminal window is fully occupied by the real-time display, and the command prompt is replaced by a footer listing options for pausing, canceling or modifying the display. When real-time monitoring is canceled, the command prompt is restored and the CLI resumes normal operation.

The following considerations affect the use of the **monitor** commands:

- Real-time monitor commands can be executed via Telnet, SSH, or a console session. Because of the slower communication rate in a console session, Extreme recommends executing the **monitor** commands *only* from a Telnet or SSH session. The default poll interval for telnet and SSH is 2 seconds, but the default polling interval for a console session is 8 seconds. If you execute **monitor** commands from a console session, flickering of the display may occur.
- If the **monitor** command is executed in a console session, console debug messages will not be displayed on the console screen.
- When the **monitor** command is executed via telnet or SSH, debug messages will not be displayed during execution of the command even with a **debug destination telnet session** configuration present.
- monitor** commands, in general, display two kinds of statistics: aggregated (counted since system startup or since last cleared using a **clear** command) and delta (counted since start of this **monitor** command or since last cleared using the *c* footer option on the monitor screen).
- Resizing of the terminal window is not supported during real-time statistics display. You must stop the execution of the command before resizing the terminal window.

- Terminal display size must be at least 80 characters wide by 24 lines in order to avoid garbled or truncated display.
- Execution of the **monitor** commands is unaffected by Telnet or SSH idle timeouts; as long as the **monitor** command is running, the terminal is not idle.
- There can be a noticeable impact on CPU utilization if the polling interval (monitor refresh interval) is short and multiple sessions are simultaneously executing **monitor** commands. When monitoring takes place by way of multiple simultaneous sessions, increase the polling interval to minimize impact on the CPU. (The polling interval/refresh rate ranges from 2 to 30 seconds, with a default value of 2 seconds for SSH or telnet connections and 8 seconds for a console session.)
- When you quit the **monitor** command, the CLI command prompt will usually be displayed at the bottom of the screen. If it appears instead in the middle of the screen, clear the screen using the command *c/s* before executing further commands.

Real-time monitoring of traffic statistics for a specific Ethernet port

To monitor traffic statistics for a specific Ethernet port, enter the following command at the Privileged EXEC level of the CLI.

```
device# monitor statistics ethernet 1/2
```

Syntax: **monitor statistics ethernet slot/port**

The *slot/port* variable specifies the port for which you want to display statistics.

The **monitor statistics** command uses page mode display to show a detailed, port-specific traffic statistics screen which is updated every poll interval. (In the XMR Series and MLX Series, this command also shows a second screen displaying network processor statistics.) You can modify the display using the commands shown in the footer. (Note that when you enter footer options they are not executed immediately but will be interpreted by the monitor engine during the next polling cycle.) The footer commands and their effects are described in the following table.

TABLE 31 Footer commands for monitor statistics display

Footer command	Description
t	Displays the transmit/output statistics (the default) and continues the execution of the original command.
r	Displays the receive/input statistics and continues the execution of the original command.
n	Continues the execution of the command for the next available Ethernet interface. If there is no next interface available, the monitor continues to display statistics for the current interface.
p	Continues the execution of the command for the previous Ethernet interface. If there is no next interface available, the monitor continues to display statistics for the current interface.
c	Clears the current delta counters and continues the execution of the original command. To clear the aggregate counters, use the appropriate clear command.
f	Freezes the execution of the command; pauses retrieval and display of the statistics. While display is frozen, the only valid commands are s and q : you can restart or quit the monitor, but any other command will be ignored.
s	Restarts the execution of the command; resumes retrieval and display of the statistics.
F	Decreases the polling interval (monitor refresh interval) by one second and continues the execution of the original command with the new refresh interval. This option will decrease the refresh interval until it is equal to 2 seconds, the minimum supported refresh interval value. The default value is 2 seconds. This command is not displayed in the footer of the statistics screen.
S	Increases the polling interval (monitor refresh interval) by one second and continues the execution of the original command with the new refresh interval. This option will increase the refresh interval until it is equal to 30 seconds, the maximum supported refresh interval value. The default value is 2 seconds. This command is not displayed in the footer of the statistics screen.
q or escape or ^c	Quits the execution of the command and returns to the command prompt.
u	XMR Series and MLX Series only: Displays the first page of the multi-page display (page-up operation).
d	XMR Series and MLX Series only: Displays the second page of the multi-page display (page-down operation).

XMR Series and MLX Series example

```
device# monitor statistics ethernet
4/1

                                Seconds: 8      poll: 8   Time: Aug 19 16:10:59
Page 1 of 2   Interface Tx Statistics      Current      Delta
Ethernet 4/1 Tx interface statistics
Traffic statistics:
  Out Packets          17083660926          533508
  Out Octets          1093354299264        34144512
  Out Unicast Packets  17083660926          533508
  Out Multicast Packets      0              0
  Out Broadcast Packets      0              0

Error statistics:
  Out Errors          0              0
  Out Discards        0              0
```

```
Tx/Rx=t/r, Page1/2=u/d, Next/Prev=n/p, Clear=c :Freeze=f/s Quit=q
                                Seconds: 40     poll: 8   Time: Aug 19 16:11:31
Page 2 of 2   NP      Tx Statistics      Current      Delta
Ethernet 4/1 Tx NP statistics
  Sent to MAC Packet    17085805774          2670758
  Raw Good Packet       17085805774          2670758
  IPX HW Forwarded Packet      0              0
  Receive from TM        17085805775          2670759
  Unicast Packet         17085805774          2670758
  Broadcast Packet        0              0
  Multicast Packet        0              0
Error statistics :
  Bad Packet Count      0              0

  ACL Drop              0              0
  Source Port Supress Drop      0              0
  IPv4 Packet           0              0
  IPv6 Packet           0              0
  IPv4 Byte             0              0
  IPv6 Byte             0              0
```

```
Tx/Rx=t/r, Page1/2=u/d, Next/Prev=n/p, Clear=c :Freeze=f/s Quit=q
```

The previous output shows the first and second pages of the detailed traffic statistics display for Ethernet port 4/1 from a XMR Series or MLX Series, displaying transmit counters (the default).

CES 2000 Series and CER 2000 Series example

```
device# monitor statistics ethernet
1/2

                                Seconds: 26     poll: 2   Time: Aug 19 16:01:41
Ethernet 1/2 Tx interface statistics      Current      Delta
Traffic statistics:
  In Packets          24847720          7738201
  In Octets          1590253440        495244864
  In Unicast Packets  24847720          7738201
  In Multicast Packets      0              0
  In Broadcast Packets      0              0

Error statistics:
  In Errors          0              0
  In Discards        0              0
```

Tx/Rx=t/r, Next/Prev=n/p, Clear=c :Freeze=f/s Quit=q

The previous output shows the detailed traffic statistics display for Ethernet port 1/2 from a CES 2000 Series or CER 2000 Series, displaying transmit counters (the default).

Real-time monitoring of traffic statistics for all Ethernet ports

To monitor summary traffic data (total packets or bytes sent and received) for all Ethernet ports (displaying up to 16 ports per screen), enter the following command at the Privileged EXEC level of the CLI.

```
device# monitor interface traffic
Seconds: 248          Time: Mar 11 20:12:08
Interface traffic statistics:
      InPackets      Delta      OutPackets      Delta
e1/1      24615      4004      24308      3986
e1/2          0          0          0          0
e1/3          0          0          0          0
e1/4          0          0          0          0
e1/5          0          0          0          0
e1/6          0          0          0          0
e1/7          0          0          0          0
e1/8          0          0          1          1
e1/9          0          0          0          0
e1/10         0          0          0          0
e1/11         0          0          0          0
e1/12         0          0          0          0
e1/13         0          0          0          0
e1/14         0          0          0          0
e1/15         0          0          0          0
e1/16         0          0          0          0
Packets=p or Bytes=b, Delta=d or Rate=r, Clear=c, Next=n :Freeze=f/s Quit=q
```

Syntax: monitor interface traffic [ethernet slot/port]

The **monitor interface traffic** command uses page mode display to produce an updating statistics screen which is updated every poll interval and which can be modified using the commands shown at the bottom of the display. (Note that when you enter footer options they are not executed immediately but will be interpreted by the monitor engine during the next polling cycle.) Normally the display begins with the lowest numbered Ethernet port; the **ethernet slot/port** option starts the display instead with the specified port.

The footer commands and their effects are described in the following table.

TABLE 32 Footer commands for monitor interface traffic display

Footer command	Description
p	Displays input/output packets instead of bytes and continues the execution of the original command.
b	Displays input/output bytes instead of packets and continues the execution of the original command.
d	Displays delta counters instead of rate counters and continues the execution of the original command.
r	Displays rate counters instead of delta counters and continues the execution of the original command.
c	Clears the current delta counters and continues the execution of the original command. To clear the aggregate counters, use the appropriate clear command.
n	Moves to the next group of interfaces and continues the execution of the original command.
f	Freezes the execution of the command; pauses retrieval and display of the statistics. While display is frozen, the only valid commands are s and q : you can restart or quit the monitor, but any other command will be ignored.
s	Restarts the execution of the command; resumes retrieval and display of the statistics.

TABLE 32 Footer commands for monitor interface traffic display (continued)

Footer command	Description
F	Decreases the polling interval (monitor refresh interval) by one second and continues the execution of the original command with the new refresh interval. This option will decrease the refresh interval until it is equal to 2 seconds, the minimum supported refresh interval value. The default value is 2 seconds. This command is not displayed in the footer of the statistics screen.
S	Increases the polling interval (monitor refresh interval) by one second and continues the execution of the original command with the new refresh interval. This option will increase the refresh interval until it is equal to 30 seconds, the maximum supported refresh interval value. The default value is 2 seconds. This command is not displayed in the footer of the statistics screen.
q or escape or ^c	Quits the execution of the command and returns to the command prompt.

Displaying recent traffic statistics for an Ethernet port

To display information from the **show statistics** command for an Ethernet port, enter the following command at any CLI level.

```

device# show statistics ethernet 9/1
PORT 9/1 Counters:

210753550720
      InPkts      1646511726      OutPkts      1646512119
InBroadcastPkts      0      OutBroadcastPkts      0
InMulticastPkts      0      OutMulticastPkts      0
InUnicastPkts      1646511726      OutUnicastPkts      1646512142
InDiscards      0      OutDiscards      0
InErrors      0      OutErrors      0
InCollisions      0      OutCollisions      0
      OutLateCollisions      0
      Alignment      0      FCS      0
InFlowCtrlPkts      0      OutFlowCtrlPkts      0
GiantPkts      0      ShortPkts      0
InBitsPerSec      3440829770      OutBitsPerSec      3440686411
InPktsPerSec      3360185      OutPktsPerSec      3360085
InUtilization      39.78%      OutUtilization      39.78%

```

Syntax: **show statistics ethernet** *slot/port*

The *slot/port* variable specifies the port for which you want to display statistics.

TABLE 33 Display of show statistics

Field	Description
InOctets	The total number of good octets and bad octets received.
OutOctets	The total number of good octets and bad octets transmitted.
InPkts	The total number of packets received. The count includes rejected and local packets that are not transmitted to the switching core for transmission.
OutPkts	The number of good packets received. The count includes unicast, multicast, and broadcasts packets.
InBroadcastPkts	The total number of good broadcast packets received.
OutBroadcastPkts	The total number of good broadcast packets transmitted.
InMulticastPkts	The total number of good multicast packets received.
OutMulticastPkts	The total number of good multicast packets transmitted.
InUnicastPkts	The total number of good unicast packets received.
OutUnicastPkts	The total number of good unicast packets transmitted.
InDiscards	The total number of packets that were received and then dropped due to a lack of receive buffers.

TABLE 33 Display of show statistics (continued)

Field	Description
OutDiscards	The total number of packets that were transmitted and then dropped due to a lack of transmit buffers.
InErrors	The total number of packets received that had Alignment errors or phy errors.
OutErrors	The total number of packets transmitted that had Alignment errors or phy errors.
InCollisions	The total number of packets received in which a Collision event was detected.
OutCollisions	The total number of packets transmitted in which a Collision event was detected.
OutLateCollisions	The total number of packets transmitted in which a Collision event was detected, but for which a receive error (Rx Error) event was not detected.
Alignment	The total number of packets received that were from 64 - 1518 octets long, but had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a non-integral number of octets (Alignment Error).
FCS	The Frame Checksum error.
InFlowCtrlPkts	The total number of ingress flow control packets. "N/A" indicates that the interface module does not support flow control statistics.
OutFlowCtrlPkts	The total number of egress flow control packets. "N/A" indicates that interface module does not support flow control statistics.
GiantPkts	<p>The total number of packets for which all of the following was true:</p> <ul style="list-style-type: none"> • The data length was longer than the maximum allowable frame size. • No Rx Error was detected. <p>This counter is only for 10GbE interfaces.</p>
ShortPkts	<p>The total number of packets received for which all of the following was true:</p> <ul style="list-style-type: none"> • The data length was less than 64 bytes. • No Rx Error was detected. • No Collision or late Collision was detected.
InBitsPerSec	The number of bits received per second.
OutBitsPerSec	The number of bits transmitted per second.
InPktsPerSec	The number of packets received per second.
OutPktsPerSec	The number of packets transmitted per second.
InUtilization	The percentage of the port's bandwidth used by received traffic.
OutUtilization	The percentage of the port's bandwidth used by transmitted traffic.

Displaying and modifying default settings for system parameters

The Multi-Service IronWare has default table sizes for the following parameters. The table sizes determine the maximum number of entries the tables can hold. You can adjust individual table sizes to accommodate your configuration needs:

- MAC address entries
- VLANs supported on a system
- Virtual interfaces
- Spanning tree instances
- RSTP instances
- IP cache size

- ARP entries
- IP routes
- IP ACL filter entries
- L2 ACL entries per ACL table
- Size for management port ACL
- IP subnets per port and per device
- IPv6 Multicast routes
- IPv6 PIM mcache
- Layer 4 sessions supported
- Number of VPLS's
- VPLS MAC entries
- VRF routes
- IPv6 cache
- IPv6 routes
- Number of tunnels
- Number of LAGs
- Configuration file size

The tables you can configure as well the defaults and valid ranges for each table differ depending on the Extreme device you are configuring.

NOTE

If you increase the number of subnet addresses you can configure on each port to a higher amount, you might also need to increase the total number of subnets that you can configure on the device.

NOTE

Changing the table size for a parameter reconfigures the device's memory. Whenever you reconfigure the memory on a Extreme device, you must save the change to the startup configuration file, then reload the software to place the change into effect.

To display the configurable tables and their defaults and maximum values, enter the following command at any level of the CLI.

Output for the XMR Series and Extreme MLX Series.

```
device#show default values
sys log buffers:50          mac age time:300 sec      telnet sessions:5
ip arp age:10 min          bootp relay max hops:4    ip ttl:64 hops
ip addr per intf:24
when multicast enabled :
igmp group memb.: 260 sec  igmp query:          125 sec
when ospf enabled :
ospf dead:40 sec          ospf hello:10 sec        ospf retrans:5 sec
ospf transit delay:1 sec
when bgp enabled :
bgp local pref.:100       bgp keep alive:60 sec    bgp hold:180 sec
bgp metric:10            bgp local as:1          bgp cluster id:0
bgp ext. distance:20      bgp int. distance:200    bgp local distance:200
when IS-IS enabled :
isis hello interval:10 sec isis hello multiplier:3
isis port metric:10       isis priority:64
isis csnp-interval:10 sec isis default-metric:10
isis distance:115         isis lsp-gen-interval:10 sec
isis lsp-interval:33 msec isis lsp-refresh-interval:900 sec
isis max-lsp-lifetime:1200 sec isis maximum-paths:4
```



```

isis retransmit-interval:5 sec      isis spf-interval:5 sec
filter change update delay:10 sec
System Parameters   Default      Maximum      Current      Actual      Bootup Revertible
mac                 131072     2097152     2097152     2097152     2097152     Yes
vlan                512        4095        4095        4095        4095        No
spanning-tree       32         128         128         128         128         No
rstp                32         128         128         128         128         No
ip-arp              8192       65536       65536       65536       65536       No
multicast-route (IPv6) 8192       153600     153600     153600     153600     Yes
pim-mcache (IPv6)   4096       4096        4096        4096        4096       Yes
ip-cache            204800     1048576    1048576    1048576    1048576     Yes
ip-route            204800     1048576    1048576    1048576    1048576     Yes
ip-subnet-port      24         128         128         128         128         No
virtual-interface   255        4095        4095        4095        4095        No
vpls-mac            8192       1000000     1000000     1000000     1000000     Yes
vpls-num            2048       16384       16384       16384       16384        No
session-limit       32768      163840     163840     163840     163840     Yes
ip-filter-sys       4096       40960       40960       40960       40960        No
mgmt-port-acl-size  20         100         100         100         100         No
l2-acl-table-entries 64         256         256         256         256         No
ipv6-cache          65536      245760     245760     245760     245760     Yes
ipv6-route          65536      245760     245760     245760     245760     Yes
vrf-route           5120       262143     262143     262143     262143     Yes
receive-cam         1024       8192        8192        8192        8192        No
ip-tunnels          256        8192        8192        8192        8192        No
lsp-out-acl-cam     0          16384       16384       16384       16384        No
trunk-num           128        256         256         256         256         No
config-file-size    8388608    16777216   16777216   16777216   16777216    No
ifl-cam             0          81920       49152       49152       49152        No
ip-source-guard-cam 0          131072     30000       30000       30000        No
ipv4-mcast-cam      8192       65536       10000       10000       10000        No
ipv6-mcast-cam      2048       16384        3500        3500        3500         No

```

Output for the CES 2000 Series

```

device#show default values
sys log buffers:50      mac age time:300 sec      telnet sessions:5
ip arp age:10 min      bootp relay max hops:4    ip ttl:64 hops
ip addr per intf:24
when multicast enabled :
igmp group memb.:140 sec  igmp query:60 sec
when ospf enabled :
ospf dead:40 sec          ospf hello:10 sec          ospf retrans:5 sec
ospf transit delay:1 sec
when bgp enabled :
bgp local pref.:100      bgp keep alive:60 sec      bgp hold:180 sec
bgp metric:10            bgp local as:1            bgp cluster id:0
bgp ext. distance:20     bgp int. distance:200     bgp local distance:200
when IS-IS enabled :
isis hello interval:10 sec      isis hello multiplier:3
isis port metric:10            isis priority:64
isis csnp-interval:10 sec      isis default-metric:10
isis distance:115             isis lsp-gen-interval:10 sec
isis lsp-interval:33 msec       isis lsp-refresh-interval:900 sec
isis max-lsp-lifetime:1200 sec  isis maximum-paths:4
isis retransmit-interval:5 sec  isis spf-interval:5 sec
filter change update delay:10 sec
System Parameters   Default      Maximum      Current
mac                 56320       131072      56320
vlan                512         4095        512
spanning-tree       32          128         32
rstp                32          128         32
ip-arp              4096        16384       4096
multicast-route (IPv6) 1024        2048        1024
pim-mcache (IPv6)   1024        2048        1024
ip-cache            16384       32768       16384
ip-route            16384       32768       16384
ip-subnet-port      24          128         24
virtual-interface   255         1024        255
vpls-mac            512         1024        512
vpls-num            512         1024        512

```

session-limit	32768	32768	32768
ip-filter-sys	4096	8192	8192
mgmt-port-acl-size	20	100	20
l2-acl-table-entries	64	256	256
ipv6-cache	1024	8192	1024
ipv6-route	1024	8192	1024
vrf-route	1024	32768	1024
receive-cam	1	1	1
ip-tunnels	32	128	32
lsp-out-acl-cam	1	1	1
trunk-num	128	255	128

Output for the CER 2000 Series device

```

device#show default values
sys log buffers:50          mac age time:300 sec          telnet sessions:5
ip arp age:10 min          bootp relay max hops:4      ip ttl:64 hops
ip addr per intf:24
when multicast enabled :
igmp group memb.: 260 sec  igmp query:          125 sec
when ospf enabled :
ospf dead:40 sec          ospf hello:10 sec          ospf retrans:5 sec
ospf transit delay:1 sec
when bgp enabled :
bgp local pref.:100       bgp keep alive:60 sec      bgp hold:180 sec
bgp metric:10             bgp local as:1            bgp cluster id:0
bgp ext. distance:20      bgp int. distance:200     bgp local distance:200
when IS-IS enabled :
isis hello interval:10 sec isis hello multiplier:3
isis port metric:10       isis priority:64
isis csnp-interval:10 sec isis default-metric:10
isis distance:115         isis lsp-gen-interval:10 sec
isis lsp-interval:33 msec isis lsp-refresh-interval:900 sec
isis max-lsp-lifetime:1200 sec isis maximum-paths:4
isis retransmit-interval:5 sec isis spf-interval:5 sec
filter change update delay:10 sec
System Parameters      Default      Maximum      Current      Actual      Bootup      Revertible
mac                    65536       131072       65536       65536       65536       Yes
vlan                   512         8192        512         512         512         Yes
spanning-tree         32          128         32          32          32         Yes
rstp                   32          128         32          32          32         Yes
ip-arp                 4096        16384       4096        4096        4096       Yes
ip-cache               290816      524288     290816      290816      290816     Yes
ip-route               290816      524288     290816      290816      290816     Yes
ip-subnet-port         24          128         24          24          24         Yes
virtual-interface      255         4095        255         255         255         Yes
vpls-mac               2048        131072      2048        2048        2048       Yes
vpls-num               128         1024        128         128         128         Yes
session-limit          32768       32768       32768       32768       32768       Yes
ip-filter-sys          4096        32768       4096        4096        4096       Yes
mgmt-port-acl-size     20          100         20          20          20         Yes
l2-acl-table-entries   64          256         64          64          64         Yes
ipv6-cache             8192        131072      8192        8192        8192       Yes
ipv6-route             8192        131072      8192        8192        8192       Yes
ip-vrf-route           1024        524288     1024        1024        1024       Yes
ip-tunnels             32          256         32          32          32         Yes
config-file-size       8388608     16777216   8388608     8388608     8388608    Yes
ip-source-guard-cam    0           131072      0           0           0          No
ip-vrf                 16          128         16          16          16         Yes
ipv6-vrf-route         128         16384       128         128         128        Yes
openflow-pvlan-entries 0           2048        0           0           0          Yes
No

```

Syntax: show default values

The following table describes the system-max values of the **show default values** command for XMR Series and MLX Series.

TABLE 34 Display of show default values for system parameters

Field	Description
Default	The default value for the system-max element. This value is used in the following conditions: a) There is no system-max configured for the corresponding element. b) If the system-max element configuration is reverted at bootup time (if it is a revertible element).
Maximum	The maximum value that this element can be configured at.
Current	The most current configured value for the system-max element. If the system-max element is configured in the running system, then the value under this column will change to reflect this value. NOTE The new value does not take affect until the node is reloaded.
Actual	The system-max value that is used by the target application of the running system. If system-max elements are reverted at bootup, then only the Actual column is affected. The Application is now using default values and will be displayed in the Actual column. Please refer to the example on the next page for more information. The Current and Bootup values are still configured on the system, and are not affected by the reversion of system-max elements at bootup.
Bootup	The system-max value that was read from the configuration when the system was booting up. If the read values are found to be acceptable, and not reverted, then the values in this column, and in the "Actual" column will have the same values. However, if the values were reverted during bootup, then the values are different for the "Revertible" elements.
Revertible	This column displays which corresponding system-max element is revertible or not. If "Yes" is displayed then the value is changed to a default value. If "No" is displayed then there no change to the value.

If system-max elements are reverted at bootup time, then the following message will display on the CLI.

```
device#show default values
...
NOTE: All the Revertible Elements were Reverted During System Bringup.
System Parameters  Default      Maximum      Current      Actual      Bootup      Revertible
mac                131072      2097152      2097152      131072      2097152      Yes
vlan               512         4095         512          512         512         No
spanning-tree      32          128          32           32          32         No
rstp               32          128          32           32          32         No
ip-arp             8192        65536        65536        65536        65536      No
multicast-route (IPv6) 8192        153600       8192         8192        8192      Yes
pim-mcache (IPv6)  4096        4096         4096         4096        4096      Yes
ip-cache           204800      1048576      524288       204800      524288     Yes
...
```

Information for the configurable tables appears under the columns shown in bold type. To simplify configuration, the command parameter you enter to configure the table is used for the table name.

To increase the size of the IP route table, enter the following commands.

```
device(config)# system-max ip-route 120000
device(config)# write memory
device(config)# exit
device# reload
```

NOTE

If you enter a value that is not within the valid range of values, the CLI will display the valid range for you.

To increase the number of IP subnet interfaces you can configure on each port on a Extreme device to 64, enter the following commands.

```
device(config)# system-max ip-subnet-port 64
device(config)# write memory
device(config)# exit
device# reload
```

Syntax: [no] **system-max ip-subnet-port** *num*

The *num* parameter specifies the maximum number of subnet addresses per port. The minimum, maximum and default values for this parameter are described in [Configuring system max values](#) on page 34.

NOTE

You must reload the software for the change to take effect.

Network Time Protocol

• Network Time Protocol overview.....	109
• How NTP works.....	111
• Configuring NTP.....	113
• Show commands.....	117
• Packet timestamping.....	122
• Management VRF for NTP.....	122

Network Time Protocol overview

Network Time Protocol (NTP) maintains uniform time across all devices in a network. The NTP commands support the configuration of an external time server to maintain synchronization between all local clocks in a network.

NTP has a hierarchical structure. At the highest level, or stratum, are precise hardware clocks, which can synchronize to highly accurate external time reference. These hardware clock devices are known as stratum 0 devices. A stratum 1 time server obtains time directly from a hardware clock and is the most accurate reference in the NTP hierarchy. All lower stratum devices obtain time from the stratum above over a network. As the network introduces timing discrepancies, lower stratum devices are a factor less accurate.

A hierarchical structure allows the overhead of providing time to many clients to be shared among many time servers. Not all clients need to obtain time directly from a stratum 1 reference, but can utilize stratum 2 or 3 references.

NTP operates on a client-server basis. A network time client periodically requests time from a time server. The time server responds with a packet of information containing a time stamp. The time stamp is then used by the client to synchronize its system time.

The NTP client maintains the server and peer state information as an association. The server and peer association is mobilized at startup when a new NTP peer connection is established, or when a user configures an NTP server or peer. The symmetric passive association is mobilized upon arrival of the NTP packet from the peer, which is not statically configured. A syslog message is generated when a new association is mobilized. The statically configured server or peer associations are not demobilized unless the user removes the configuration. If the NTP packet from the symmetric passive peer results in an error or timeout, then the symmetric passive peer is demobilized. A syslog message is generated when an association is demobilized. For more information about the generated syslog messages, refer to *Extreme NetIron Monitoring Configuration Guide*.

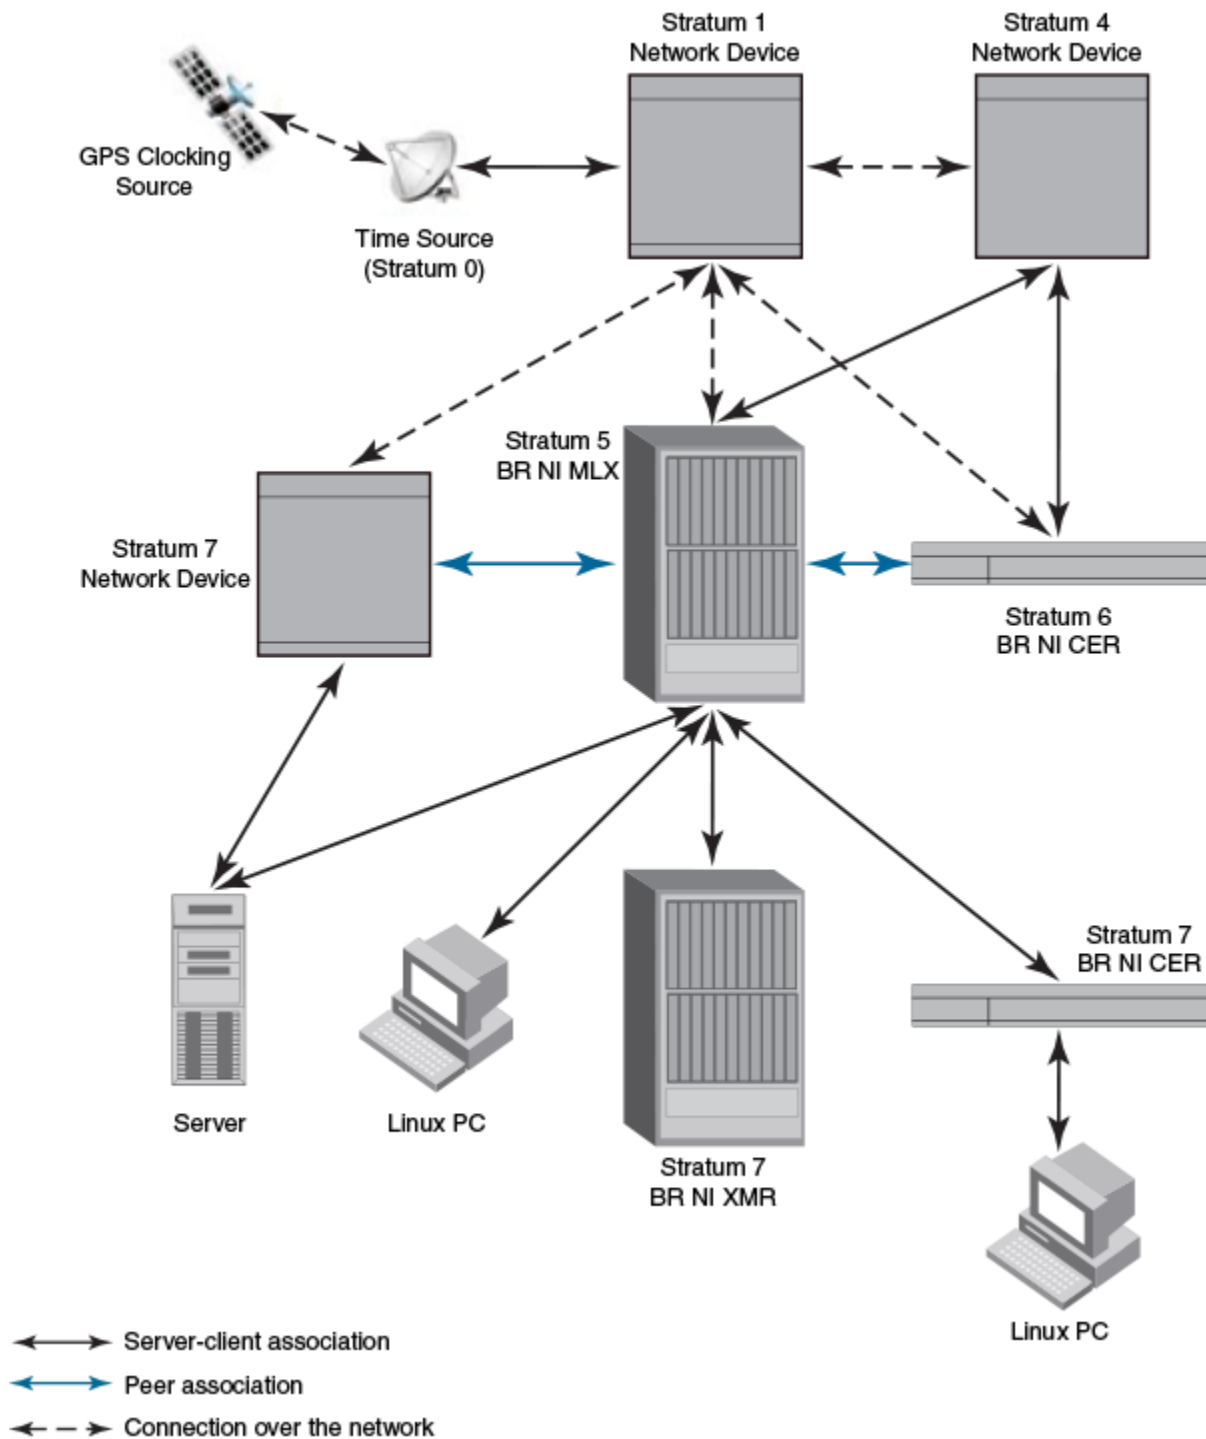
NTP uses UTC (Universal Time Coordinated) time, which is similar to GMT time. It knows nothing of local time zones or daylight-saving time. It is a function of the time client to apply an offset to the supplied time to adjust for local time. In this manner, a time server located anywhere in the world can provide synchronisation to a client located anywhere else in the world. It allows clients to utilize different time zone and daylight-saving properties.

To keep the time in your network current, it is recommended that each device have its time synchronized with at least four external NTP servers. External NTP servers should be synchronized among themselves in order to maintain time synchronization.

NOTE

Network Time Protocol (NTP) commands must be configured on each individual device.

FIGURE 1 NTP sample network configuration



Network Time Protocol leap second

A leap second is a second added to Coordinated Universal Time (UTC) in order to keep it synchronized with astronomical time (UT1).

There are two main reasons that cause leap seconds to occur. The first is that the atomic second defined by comparing cesium clocks to the Ephemeris Time (ET) scale was incorrect, as the duration of the ephemeris second was slightly shorter than the mean solar second and this characteristic was passed along to the atomic second. The second reason for leap seconds is that the speed of the Earth's rotation is not constant. It sometimes speeds up, and sometimes slows down, but when averaged over long intervals the trend indicates that it is gradually slowing. This gradual decrease in the rotational rate is causing the duration of the mean solar second to gradually increase with respect to the atomic second.

Leap seconds are added in order to keep the difference between UTC and astronomical time (UT1) to less than 0.9 seconds. The International Earth Rotation and Reference Systems Service (IERS), measures Earth's rotation and publishes the difference between UT1 and UTC. Usually leap seconds are added when UTC is ahead of UT1 by 0.4 seconds or more.

How Extreme supports leap second handling for NTP

The obvious question raised is what happens during the NTP leap second itself.

Specifically, a positive leap second is inserted between second 23:59:59 of a chosen UTC calendar date (the last day of a month, usually June 30 or December 31) and second 00:00:00 of the following date. This extra second is displayed on UTC clocks as 23:59:60. On clocks that display local time tied to UTC, the leap second may be inserted at the end of some other hour (or half-hour or quarter-hour), depending on the local time zone. When ever there is a leap second the NTP server notifies by setting the NTP leap second bits.

On Extreme devices when ever there is a negative leap second, the clock is set once second backward of the following date as described here. On positive leap second the clock suppress second 23:59:59 of the last day of a chosen month, so that second 23:59:58 of that date would be followed immediately by second 00:00:00 of the following date.

How NTP works

NTP server

A **NTP server** will provide the correct network time on your device using the Network time protocol (NTP). Network Time Protocol can be used to synchronize the time on devices across a network. A NTP time server is used to obtain the correct time from a time source and adjust the local time in each connecting device.

The NTP server can operate in master mode to serve time using the local clock, when it has lost synchronization.

NTP client

An NTP client gets time responses from an NTP server or servers, and uses the information to calibrate its clock. This consists of the client determining how far its clock is off and adjusting its time to match that of the server. The maximum error is determined based on the round-trip time for the packet to be received.

NTP peer

NTP peer mode is intended for configurations where a group of devices operate as mutual backups for each other. If one of the devices lose a reference source, the time values can flow from the surviving peers to all the others.

The NTP peer can operate in:

Symmetric Active - When the peer is configured using the peer command.

Symmetric Passive - Dynamically learnt upon arrival of a NTP packet from the peer which is not configured. The symmetric passive association is removed on timeout or error.

NTP broadcast server

An NTP server can also operate in a broadcast or multicast mode. Both work similarly; broadcast servers send periodic time updates to a broadcast address, while multicast servers send periodic updates to a multicast address. Using broadcast packets can greatly reduce the NTP traffic on a network, especially for a network with many NTP clients.

The interfaces should be enabled with NTP broadcasting. The NTP broadcast server broadcasts the NTP packets periodically (every 64 sec) to subnet broadcast IP address of the configured interface.

NTP broadcast client

An NTP broadcast or multicast client listens for NTP packets on a broadcast or multicast address. When the first packet is received, it attempts to quantify the delay to the server in order to better quantify the correct time from later broadcasts. This is accomplished by a series of brief interchanges where the client and server act as a regular (non-broadcast) NTP client and server. Once these interchanges occur, the client has an idea of the network delay and thereafter can estimate the time based only on broadcast packets.

Synchronizing time

After the system peer is chosen, the system time is synchronized using one of the following ways based on the time difference with system peer:

- < 128 msec - The system clock is adjusted slowly towards the system peer time reference time.
- > 128 msec and < 1000 sec - The system clock is stepped to the system peer reference time and the NTP state information is cleared.
- > 1000 sec - NTP is operationally disabled. The admin should set the time to the current UTC time.

Configuration considerations of NTP

- NTP multicast server, client, and manycast client functionalities are not supported.
- In a scaled network, Extreme recommends configuring the NetIron device to one external NTP server (at minimum), or a dedicated internal NTP server.
- While upgrading from R05.2.00 or lower versions to R05.3.00, the SNTP configuration will be ignored.
- On reboot or MP switchover, all the NTP state information will be lost and time synchronization will start from fresh. The time synchronized to real time clock is retained across reboot and MP switchover.
- The following SNTP MIB objects are not supported.
 - snNTPPollInterval
 - snNTPSync
 - All the objects in snNTPServerTable
- The web management support for SNTP is removed

The following optional features are not supported

- NTP version 4 Extension fields
- The NTP packets having control (6) or private (7) packet modes
- Autokey public key authentication
- NTP version 1 and 2
- Hostnames

Configuring NTP

Before you begin to configure NTP, you must use the **clock set** command to set the time on your device to within 1000 seconds of the coordinated Universal Time (UTC).

Changing to the NTP mode

Use the **ntp** command to enable the NTP client and server mode.

```
device(config)# ntp
```

Syntax: ntp

Enabling NTP authentication

To enable Network Time Protocol (NTP) strict authentication, use the **authenticate** command. To disable the function, use the **no** form of this command.

```
device(config-ntp)# authenticate
```

Syntax: [no] authenticate

Defining an authentication key

To define an authentication key for Network Time Protocol (NTP), use the **authentication-key** command. To remove the authentication key for NTP, use the **no** form of this command.

```
device(config-ntp)# authentication-key key-id 1 md5 moof
```

Syntax: [no] authentication-key *key-id* [**md5** | **sha1**] *keystring*

The valid *key-id* parameter is 1 to 65535.

The **md5** keyword specifies the message authentication support that is provided using the Message Digest 5 Algorithm.

The **sha1** keyword specifies that the SHA1 keyed hash algorithm is used for NTP authentication.

NOTE

In JITC mode, MD5 authentication scheme is disabled for NTP. For more information on JITC, refer to *Extreme NetIron Security Configuration Guide*.

The *keystring* parameter is the value of the MD5 key or SHA1 key. The maximum length of the key string may be defined up to 16 characters. Up to 32 keys may be defined.

Specifying a source interface

To use a particular source interface in Network Time Protocol (NTP) packets, use the **source-interface** command. To remove the specified source address, use the **no** form of this command.

NOTE

If the **source-interface** is not configured, then the lowest IP address in the outgoing interface will be used in the NTP packets.

```
device(config-ntp)# source-interface ethernet 3/1
```

Syntax: **[no] source-interface ethernet slot/port || loopback num | ve num**

The **ethernet slot/port** parameter specifies the ethernet port number.

The **loopback num** parameter specifies the loopback interface number.

The **ve num** parameter specifies the virtual port number.

Enabling or disabling the VLAN containment for NTP

To enable or disable the VLAN containment for NTP, use the **access-control vlan** command. To remove the specified NTP VLAN configuration, use the **no** form of this command.

NOTE

The management interface is not part of any VLAN. When configuring the VLAN containment for NTP, it will not use the management interface to send or receive the NTP packets.

```
device(config-ntp)# access-control vlan 100
```

Syntax: **[no] access-control vlan vlan-id**

The **vlan-id** parameter specifies the VLAN ID number.

Configuring the NTP client

To configure the device in client mode and specify the NTP servers to synchronize the system clock, use the **server** command. A maximum 8 NTP servers can be configured. To configured NTP server, use the **no** form of this command.

```
device(config-ntp)#server 10.2.3.4 key 1234
```

Syntax: **# [no] server ipv4address | ipv6address [version 3 | 4] [key keyid] [minpoll interval] [maxpoll interval]**

The **ipv4 address / ipv6 address** parameter is the IP address of the server providing the clock synchronization.

The **version 3/4** option defines the Network Time Protocol (NTP) version number. If this option is not specified, then it defaults to 4.

The **key key id** option defines the authentication key. By default, no authentication key is configured.

The **minpoll interval** option is the shortest polling interval. The range is from 4 through 17. Default is 6. The interval argument is power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s etc.).

The **maxpoll interval** option is the longest polling interval. The range is 4 through 17. Default is 10. The interval argument is calculated by the power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s etc.).

Configuring the NTP peer

To configure the software clock to synchronize a peer or to be synchronized by a peer, use the **peer** command. A maximum of 8 NTP peers can be configured. To disable this capability, use the **no** form of this command.

```
device(config-ntp)# peer 10.2.3.4 key 1234
```

Syntax: **[no] peer** *ipv4address* | *ipv6address* [**version** 3 | 4] [**key** *keyid*] [**minpoll** *interval*] [**maxpoll** *interval*]

The *ipv4 address* / *ipv6 address* parameter is the IP address of the peer providing the clock synchronization.

The **version** 3/4 option defines the Network Time Protocol (NTP) version number. If this option is not specified, then it defaults to 4.

The **key** *key id* option defines the authentication key. By default, no authentication key is configured.

The **minpoll** *interval* option is the shortest polling interval. The range is from 4 through 17. Default is 6. The interval argument is power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s etc.).

The **maxpoll** *interval* option is the longest polling interval. The range is 4 through 17. Default is 10. The interval argument is calculated by the power of 2 (4=16s, 5=32s, 6=64s, 7=128s, 8=256s, 9=512s etc.).

Disable Symmetric Passive Association

A router's association mode changes to *Symmetric Passive* after it receives NTP packets from Active peers within a network. A peer entry is also created on the Active peer for this Symmetric Passive peer router.

Use the **[no] disable symmetric passive** command to prevent this device from receiving NTP packets from active peers. NTP packets from Active peers are silently dropped. When this command is executed, all existing Active peer entries are also removed.

The default state for the device is to accept any NTP packet from an Active peer.

```
device(config-ntp)# disable symmetric-passive
device(config-ntp)# no disable symmetric-passive
```

Syntax: **[no] disable symmetric-passive**

To view the current status for **disable symmetric-passive**, use the **show ntp status** command.

```
device(config-ntp)# show ntp status

Clock is synchronized, stratum 2, reference clock is 10.6.25.32
precision is 2**-16
reference time is E2D6553A.80000000 (10:08:57.2750081523 GMT+00 Thu Aug 06 2020)
clock offset is 4.4461 msec, root delay is 72.3396 msec
root dispersion is 15.2223 msec, peer dispersion is 0.2244 msec
system poll interval is 1024, last clock update was 665 sec ago
NTP server mode is disabled, NTP client mode is enabled
NTP symmetric passive mode is disabled
NTP master mode is disabled, NTP master stratum is 8
NTP is not in panic mode
```

Configuring NTP on an interface

To configure the NTP interface context, use the **ntp-interface** command. The broadcast server or client is configured on selected interfaces. To remove the NTP broadcast configurations on the specified interface, use the **no** form of this command.

NOTE

The **ntp-interface** command is a mode change command, and will not be included in to the **show run output** unless there is configuration below that interface.

```
device(config-ntp)# ntp-interface ethernet 2/13
device(config-ntp-if-e1000-2/13)#

device(config-ntp)# ntp-interface management 1
(config-ntp-mgmt-1)#

device(config-ntp)# ntp-interface ve 100
device(config-ntp-ve-100)#
```

Syntax: **[no] ntp-interface [management 1 | ethernet slot/port | ve id]**

The **management 1** parameter is the management port 1.

The **ethernet slot/port** parameter specifies the ethernet port number.

The **ve id** parameter specifies the virtual port number.

Configuring the broadcast client

To configure a device to receive Network Time Protocol (NTP) broadcast messages on a specified interface, use the **broadcast client** command. NTP broadcast client can be enabled on maximum of 16 ethernet interfaces. If the interface is operationally down or NTP is disabled, then NTP broadcast server packets are not received. To disable this capability, use the **no** form of this command.

```
device(config-ntp mgmt-1)# broadcast client
```

Syntax: **[no] broadcast client**

Configuring the broadcast destination

To configure the options for broadcasting Network Time Protocol (NTP) traffic, use the **ntp broadcast destination** command. NTP broadcast server can be enabled on maximum 16 Ethernet interfaces and four subnet addresses per interface. If the interface is operationally down or there is no ip address configured for the subnet address, then NTP broadcast server packets are not sent. To disable this capability, use the **no** form of this command.

NOTE

This command is not effective, if the NTP server is disabled.

```
device(config)#int m1
device(config-if-mgmt-1)#ip address 10.20.99.173/24
device(config-if-mgmt-1)#ntp
device(config-ntp)#ntp-interface m1
device(config-ntp-mgmt-1)# broadcast destination 10.20.99.0 key 2
```

Syntax: **[no] broadcast destination ip-address [key key-id] [version 3 | 4]**

The **IP-address** parameter is the IPv4 subnet address of the device to send NTP broadcast messages to.

The **key key id** option defines the authentication key. By default, no authentication key is configured.

The **version 3 | 4** option defines the Network Time Protocol (NTP) version number. If this option is not specified, then it defaults to 4.

Disabling NTP

To disable the NTP server and client, use the **disable** command. Disabling the NTP server or client mode will not remove the configurations. To enable receipt of NTP packets, use the **no** form of this command.

```
device(config-ntp)# disable
```

Syntax: [no] **disable** [**serve**]

If the **serve** keyword is specified, then NTP will not serve the time to downstream devices. This keyword disables the NTP server mode functionalities.

If this keyword is not specified, then both NTP client mode and NTP server mode functionalities will be disabled.

Configuring the master

To configure the Multi-Service IronWare as a Network Time Protocol (NTP) master clock to which peers synchronize themselves when an external NTP source is not available, use the **master** command. The master clock is disabled by default. To disable the master clock function, use the **no** form of this command.

NOTE

This command is not effective, if the NTP is enabled in client-only mode.

```
device(config-ntp)# master stratum 5
```

Syntax: [no] **master** [**stratum number**]

Stratum *number* is the number from 2 to 15. It indicates the NTP stratum number that the system will claim.

Enabling or disabling NTP logging

To enable or disable Network Time Protocol (NTP) message logging, use the **logging enable ntp** command. By default, the logging is enabled for NTP. To disable NTP logging, use the **no** form of this command.

```
device(config)# logging enable ntp
```

Syntax: [no] **logging enable ntp**

Show commands

The **show ntp** command allows you to display NTP status and association information of the NTP server or peers.

Displaying NTP status

Use the **show ntp status** command to display the NTP status

```
device#show ntp status

Clock is synchronized, stratum 4, reference clock is 10.20.99.174
precision is 2**-16
reference time is D281713A.80000000 (03:21:29.3653007907 GMT+00 Thu Dec 01 2011)
clock offset is -2.3307 msec, root delay is 24.6646 msec
root dispersion is 130.3376 msec, peer dispersion is 84.3335 msec
system poll interval is 64, last clock update was 26 sec ago
NTP server mode is enabled, NTP client mode is enabled
NTP symmetric passive mode is disabled
```

```
NTP master mode is disabled, NTP master stratum is 8
NTP is not in panic mode
```

TABLE 35 show ntp status command output descriptions

Field	Description
synchronized	Indicates the system clock is synchronized to NTP server or peer.
stratum	Indicates the stratum number that this system is operating. Range 2..15.
reference	IPv4 address or first 32 bits of the MD5 hash of the IPv6 address of the peer to which clock is synchronized.
precision	Precision of the clock of this system in Hz.
reference time	Reference time stamp.
clock offset	Offset of clock (in milliseconds) to synchronized peer.
root delay	Total delay (in milliseconds) along path to root clock.
root dispersion	Dispersion of root path.
peer dispersion	Dispersion of root path.
system poll interval	Poll interval of the local system.
last update	Time the router last updated its NTP information
server mode	Status of the NTP server mode for this device.
symmetric passive mode	Status of Symmetric Passive Association mode for this device.
client mode	Status of the NTP client mode for this device.
master	Status of the master mode
master stratum	Stratum number that will be used by this device when master is enabled and no upstream time servers are accessible.
panic mode	Status of the panic mode. If the clock offset is more than 1000 seconds with the current time, then panic mode will be on.

Displaying NTP associations

Use the **show ntp associations** command to display detailed association information of the NTP server or peers.

```
device# show ntp associations
address      ref clock      st  when poll reach delay  offset  disp
*~172.19.69.1 172.24.114.33 3   25   64   3     2.89  0.234  39377
~2001:db8::234
      INIT      16   -    64   0     0.00  0.000  15937
* synced, # selected, + candidate, - outlayer, x falseticker, ~ configured
```

TABLE 36 show ntp associations command output descriptions

Field	Description
*	The peer has been declared the system peer and lends its variables to the system variables.
#	This peer is a survivor in the selection algorithm.
+	This peer is a candidate in the combine algorithm.
-	This peer is discarded as outlier in the clustering algorithm.
x	This peer is discarded as falseticker in the selection algorithm.
~	The server or peer is statically configured.
address	IPv4 or IPv6 address of the peer.
ref clock	IPv4 address or first 32 bits of the MD5 hash of the IPv6 address of the peer to which clock is synchronized.

TABLE 36 show ntp associations command output descriptions (continued)

Field	Description
St	Stratum setting for the peer.
when	Time, in seconds, since last NTP packet was received from peer.
poll	Polling interval (seconds).
reach	Peer reachability (bit string, in octal).
delay	Round-trip delay to peer, in milliseconds.
offset	Relative time difference between a peer clock and a local clock, in milliseconds.
disp	Dispersion.

Displaying NTP associations details

Use the **show ntp associations detail** command to display all the NTP servers and peers association information.

```
device# show ntp association detail
2001:db8:99:30::1 configured server, sys peer, stratum 3
ref ID 10.235.61.9, time d288dc3b.f2a17891 (10:23:55.4070668433 Pacific Tue Dec 06 2011)
our mode client, peer mode server, our poll intvl 10, peer poll intvl 10,
root delay 0.08551025 msec, root disp 0.09309387, reach 17, root dist 0.17668502
delay 0.69961487 msec, offset -13.49459670 msec, dispersion 17.31550718,
precision 2**-16, version 4
org time d288df70.a91de561 (10:37:36.2837308769 Pacific Tue Dec 06 2011)
rcv time d288df70.a0c8d19e (10:37:36.2697515422 Pacific Tue Dec 06 2011)
xmt time d288df70.a086e4de (10:37:36.2693194974 Pacific Tue Dec 06 2011)
filter delay      1.7736      0.9933      0.8873      0.6699      0.7709      0.7712      0.7734      6.7741
filter offset    -17.9936     33.0014    -13.6604    -13.4494    -14.4481    -16.4453    -18.4423    -22.0025
filter disp      15.6660      0.0030     17.7730     17.7700     17.6670     17.6640     17.6610     16.6635
filter epoch      55824      56866      55686      55688      55690      55692      55694      55759
```

Use the **show ntp associations detail IPv4 address / IPv6 address** command to display the NTP servers and peers association information for a specific ip address.

```
device# show ntp association detail 10.99.40.1
10.99.40.1 configured server, candidate, stratum 3
ref ID 10.45.57.38, time d288de7d.690ca5c7 (10:33:33.1762436551 Pacific Tue Dec 06 2011)
our mode client, peer mode server, our poll intvl 10, peer poll intvl 10,
root delay 0.02618408 msec, root disp 0.10108947, reach 3, root dist 0.23610585
delay 0.92163588 msec, offset 60.77749188 msec, dispersion 70.33842156,
precision 2**-16, version 4
org time d288defa.b260a71f (10:35:38.2992678687 Pacific Tue Dec 06 2011)
rcv time d288defa.a2efbd41 (10:35:38.2733620545 Pacific Tue Dec 06 2011)
xmt time d288defa.a2ae54f8 (10:35:38.2729334008 Pacific Tue Dec 06 2011)
filter delay      0.000      6.7770      6.7773      6.7711      6.7720      6.7736      6.7700      0.9921
filter offset      0.000     19.0047     19.1145     19.2245     19.3313     17.4410     15.4463     60.7777
filter disp    16000.000     16.0005     15.9975     15.9945     15.9915     15.8885     15.8855      0.0030
filter epoch      55683      55683      55685      55687      55689      55691      55693      56748
```

Syntax: **show ntp association detail IPv4 address | IPv6 address**

The IPv4 or IPv6 address of the NTP peer

TABLE 37 show ntp associations detail command output descriptions

Field	Description
server	Indicates server is statically configured.
symmetric active peer	Indicates peer is statically configured.
symmetric passive peer	Indicates peer is dynamically configured.

TABLE 37 show ntp associations detail command output descriptions (continued)

Field	Description
sys_peer	This peer is the system peer
candidate	This peer is chosen as candidate in the combine algorithm.
reject	This peer is rejected by the selection algorithm
false tick	This peer is dropped as false tick by the selection algorithm
outlier	This peer is dropped as outlier by the clustering algorithm
Stratum	Stratum number
ref ID	IPv4 address or hash of IPv6 address of the upstream time server to which the peer is synchronized.
Time	Last time stamp that the peer received from its master.
our mode	This system's mode relative to peer (active /passive /client /server /bdcast /bdcast client).
peer mode	Mode of peer relative to this system
our poll intvl	This system's poll interval to this peer.
peer poll intvl	Poll interval of peer to this system
root delay	The delay along path to root (the final stratum 1 time source).
root disp	Dispersion of path to root.
root distance	The distance from the server or peer to the client
reach peer	The peer reachability (bit string in octal).
Delay	Round-trip delay to peer.
offset	Offset of peer clock relative to this clock
Dispersion	Dispersion of peer clock
precision	Precision of peer clock
version	Peer NTP version number
org time	Originate time stamp of the last packet.
rcv time	Receive time stamp of the last packet.
xmt time	Transmit time stamp of the last packet.
filter delay	Round-trip delay in milliseconds of last 8 samples.
filter offset	Clock offset in milliseconds of last 8 samples.
filter error	Approximate error of last 8 samples

Configuration examples

The following sections lists configuration examples to configure the Extreme device.

NTP server and client mode configuration

Sample CLI commands to configure the NI device in NTP server and client modes.

```

device(config-ntp)# server 10.1.2.3 minpoll 5 maxpoll 10
device(config-ntp)# server 2001:db8::1/64
device(config-ntp)# peer 10.100.12.18
device(config-ntp)# peer 10.100.12.20
device(config-ntp)# peer 10.100.12.67
device(config-ntp)# peer 10.100.12.83

```


NTP client mode configuration

Sample CLI commands to configure the Extreme device in NTP client mode.

```
device(config-ntp)# server 10.1.2.3 minpoll 5 maxpoll 10
device(config-ntp)# server 2001:db8::1/24
device(config-ntp)# peer 10.100.12.83
device(config-ntp)# disable serve
```

NTP strict authentication configuration

Sample CLI commands to configure the Extreme device in strict authentication mode.

```
device(config-ntp)# authenticate
device(config-ntp)# authentication-key key-id1 md5 key123
device(config-ntp)# server 10.1.2.4 key 1
```

NTP loose authentication configuration

Sample CLI commands to configure the NI device in loose authentication mode. This allows some of the servers or clients to use the authentication keys.

```
device(config-ntp)# authentication-key-id key-id1 md5 key123
device(config-ntp)# server 10.1.2.4 key 1
device(config-ntp)# server 10.1.2.7
```

NTP interface context for broadcast server or client mode

Sample CLI command enter the NTP interface context.

```
device(config)#int management 1
device(config-if-mgmt-1)#ip address 10.20.99.173/24
device(config-if-mgmt-1)#ntp
device(config-ntp)# ntp-interface management 1
device(config-ntp-mgmt-1)# broadcast destination 10.23.45.128
device(config-ntp)# ntp-interface ethernet 1/3
device(config-ntp-if-e1000-1/3)# broadcast destination 10.1.1.0 key 1
device(config-ntp)# ntp-interface ve 100
device(config-ntp-ve-100)# broadcast destination 10.2.2.0 key 23
```

NTP broadcast client configuration

Sample CLI commands to configure the NTP broadcast client

```
device(config-ntp)# ntp-interface management 1
device(config-ntp-mgmt-1)# broadcast client
device(config-ntp)# ntp-interface ethernet 1/5
device(config-ntp-if-e1000-1/5)# broadcast client
device(config-ntp)# ntp-interface ve 100
device(config-ntp-ve-100)# broadcast client
```

Packet timestamping

Packet timestamping appends a timestamp to the incoming packets to a port so that the egressing packets have the timestamp. The timestamp enables the analytical servers connected on the egress port to perform various types of performance analysis.

Packet timestamping is configured per packet processor. By default, packet timestamping is disabled. When packet timestamping is enabled, a timestamp of 8 bytes in length is added after the payload, followed by Cyclic Redundancy Check (CRC) of 4 bytes. The CRC is recalculated after inserting the timestamp in the incoming packet. The Frame Check Sequence (FCS) in the packet is recalculated and updated. All other fields in the packet remain unchanged. This behavior will be consistent on all packet types (tagged, untagged, ipv4, ipv6, and so on).

If Network Time Protocol (NTP) is configured, the NTP time is stamped on the packets. Otherwise, the internal clock time of the line card is stamped on the packets. If both packet timestamping and source port labeling are enabled, the timestamp bytes will precede the port label bytes, followed by the recalculated CRC.

NOTE

Because the MLXe does not support Precision Time Protocol (PTP), nanosecond accuracy is not supported.

Supported hardware

Packet timestamping feature is supported on the 20X10G, 4X40G, and 2X100G line cards. On 20X10G and 2X100G line cards, only the NPB FPGA supports packet timestamping. Main-FPGA does not support packet timestamping on the 20X10G and 2X100G line cards.

Configuring packet timestamping

The **packet-timestamp** command configures the timestamping functionality.

1. Enter the global configuration mode.

```
device# configure terminal
```

2. Enter the **packet-timestamp** command to configure the timestamping functionality.

```
device(config)# packet-timestamp slot 1 device-id 1
```

3. Enter the **show packet-timestamp** command to display the packet timestamp configuration.

```
device(config)# show packet-timestamp
```

Management VRF for NTP

NTP supports the management-VRF to isolate management traffic from network data traffic.

Network Time Protocol (NTP) maintains uniform time across all devices in a network. The NTP commands support the configuration of an external time server to maintain synchronization between all local clocks in a network. NTP operates on a client-server basis. A network time client periodically requests time from a time server. The time server responds with a packet of information containing a time stamp. The time stamp is then used by the client to synchronize its system time. Network Time Protocol (NTP) commands must be configured on each individual device.

The management VRF is used to provide secure management access to the device by sending inbound and outbound management traffic through the VRF specified as a global management VRF and through the out-of-band (OOB) management port, thereby isolating

management traffic from the network data traffic. Any VRF, except the default VRF, can be configured as a management VRF. When a management VRF is configured, the management traffic is allowed through the ports belonging to the specified VRF and the out-of-band management port. The management traffic through the ports belonging to the other VRFs and the default VRF are dropped and the rejection statistics are incremented.

Enabling management-VRF support for NTP causes the incoming and outgoing traffic to travel through the management VRF or an out-of-bound (OOB) port. To support management-VRF support for NTP, all interfaces must be configured to be part of the global management VRF.

NOTE

The IPv6 management VRF is not supported on Extreme NetIron CES Series and NetIron CER Series devices.

Restrictions for NTP support of the management VRF

If the **source-interface** command is configured, but the source interface is not part of the management VRF, the response packets are dropped. Traffic to any interface that is not part of the management VRF is dropped. If source interface for NTP support is configured, the command is accepted and a warning message is displayed: `No packets will be received since the source-interface for NTP is not part of the management-vrf.`

All NTP interfaces, including the source interface and broadcast interfaces must be part of the management VRF.

Enabling Management VRF for NTP

To isolate management traffic from network data traffic, configure a management VRF and enable NTP for the management VRF.

NTP server and client configuration

While the main reason for configuring this task is to enable the separation of management traffic from data traffic, the steps here show how to configure a VRF and use it as the management VRF. Both these items must be configured before enabling the NTP support of management VRF.

If management VRF for NTP is configured before a management VRF is configured the following warning is shown when the **management-vrf enable** command is entered: `Warning - Global Management VRF is not configured. Management VRF will not be enforced for NTP Tx and Rx packets.`

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **vrf** command to create a VRF instance, "corporate" in this example.

```
device(config)# vrf corporate
```

3. Configure the VRF, corporate, as the management VRF.

```
device(config-vrf-corporate)# management-vrf corporate
VRF corporate has been configured as management-vrf
```

4. Return to global configuration mode.

```
device(config-vrf-corporate)# exit
```

5. Enable NTP.

```
device(config)# ntp
```

6. Enable NTP support for the management VRF.

```
device(config-ntp)# management-vrf enable
```

The following example configures a VRF, uses it as the management VRF, and enables the NTP support of the management VRF. Extra configuration steps are shown where the Ethernet interface 1/2 is part of the management VRF and is configured as a broadcast client under NTP.

```
device# configure terminal
device(config)# vrf corporate
device(config-vrf-corporate)# management-vrf corporate
VRF corporate has been configured as management-vrf
device(config-vrf-corporate)# exit
device(config)# ntp
device(config-ntp)# management-vrf enable
!
device(config-ntp)# ntp-interface ethernet 1/2
device(config-ntp-if-e1000-1/2)# broadcast client
```

After NTP support has been configured and traffic is being isolated, perform the Displaying NTP support of the management VRF task.

Displaying NTP support of the management VRF

To display management VRF packet and session rejection statistics including dropped packets due to failure in management VRF validation.

To verify that the NTP support for management VRF is enabled, use the **show management-vrf** command.

```
device# show management-vrf

Management VRF name : corporate
Management Application  Rx Drop Pkts      Tx Drop Pkts
SNMP Engine             36      0
RADIUS Client            0      8
TFTP Client              0      4
Traps                   -      55
SysLogs                 -      78
NTP Server               0      0
NTP Client               0      0
TCP Connection rejects:
Telnet                   :      1
SSH                      :      1
TACACS+ Client           :      8
```

Cisco Discovery Protocol

• Cisco Discovery Protocol overview.....	125
• Enabling CDP packet interception.....	125
• Displaying CDP packet information.....	126
• Clearing CDP statistics and neighbor information.....	127

Cisco Discovery Protocol overview

Using multicast announcements to share information about Cisco devices, Cisco Discovery Protocol (CDP) is a proprietary Layer 2 protocol that is equivalent to the Extreme protocol Foundry Discovery Protocol (FDP).

Cisco Discovery Protocol (CDP) packets are used by Cisco devices to advertise themselves to other Cisco devices. By default, Extreme devices forward these packets without examining their contents. You can configure a Extreme device to intercept and display the contents of CDP packets. This feature is useful for learning device and interface information for Cisco devices in the network.

Extreme devices support intercepting and interpreting CDP version 1 and CDP version 2 packets.

NOTE

The Extreme device can interpret only the information fields that are common to both CDP version 1 and CDP version 2.

NOTE

When you enable interception of CDP packets, the Extreme device drops the packets. As a result, Cisco devices will no longer receive the packets.

CDP support was replaced with the IEEE 802.1AB standard Link Layer Discovery Protocol (LLDP) that is implemented by multiple vendors and is functionally similar to CDP. It is used to share information about other directly connected Cisco equipment, such as the operating system version and IP address. CDP can also be used for On-Demand Routing, which is a method of including routing information in CDP announcements so that dynamic routing protocols do not need to be used in simple networks.

Enabling CDP packet interception

A Extreme device can be enabled to intercept and display Cisco Discovery Protocol (CDP) packets.

CDP packet interception is disabled by default on all interfaces. CDP packet interception can be enabled globally to apply to all interfaces. If CDP packet interception is to be disabled for an individual interface, the configuration is applied in interface configuration mode. This task shows how to enable CDP globally, disable CDP on one interface and reenable CDP on the interface.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Globally enable CDP packet interception.

```
device(config)# cdp run
```

3. Enter interface configuration mode.

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

4. Disable CDP packet interception on Ethernet interface 1/2.

```
device(config-if-e1000-1/2)# no cdp enable
```

5. Reenable CDP packet interception on Ethernet interface 1/2.

```
device(config-if-e1000-1/2)# cdp enable
```

The following example enables CDP packet interception globally and disables CDP packet interception on Ethernet interface 1/2.

```
device# configure terminal
device(config)# cdp run
device(config)# interface ethernet 1/2
device(config-if-e1000-1/2)# no cdp enable
```

Displaying CDP packet information

After enabling CDP packet interception, you can view CDP packet information.

Ensure that CDP has been enabled.

You can display the following CDP information:

- Cisco neighbors
- CDP entries for all Cisco neighbors or a specific neighbor
- CDP packet statistics

NOTE

The commands used to display CDP information are the same as those used to display FDP information. In the following steps we are only displaying CDP information that a Extreme device has intercepted. You will normally see Foundry Discovery Protocol (FDP) information in addition to CDP information.

1. To display CDP entries for all neighbors, enter the following command:

```
device# show fdp entry *

Device ID: Router
Entry address(es):
IP address: 10.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet 5/0/0
Holdtime : 124 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1, RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

2. To display CDP entries for a specific device, specify the device ID.

```
device# show fdp neighbors ethernet 1/1

Device ID: Router
Entry address(es):
IP address: 10.95.6.143
Platform: cisco RSP4, Capabilities: Router
Interface: Eth 1/1, Port ID (outgoing port): FastEthernet5/0/0
Holdtime : 127 seconds
Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-JSV-M), Version 12.0(5)T1, RELEASE SOFTWARE
(fc1)
Copyright (c) 1986-1999 by cisco Systems, Inc.
Compiled Thu 19-Aug-99 04:12 by cmong
```

3. To display CDP packet statistics, enter the following command:

```
device# show fdp traffic

CDP counters:
Total packets output: 0, Input: 3
Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
No memory: 0, Invalid packet: 0, Fragmented: 0
```

Clearing CDP statistics and neighbor information

Cisco Discovery Protocol (CDP) update information and statistics can be cleared.

Before clearing CDP information ensure that CDP is enabled.

You can clear the following CDP information:

- Information received in CDP updates
- CDP statistics

NOTE

The same commands clear information for both FDP and CDP.

1. To clear the information received in CDP updates from neighboring devices, enter the following command:

```
device# clear fdp table
```

2. To clear CDP statistics, enter the following command:

```
device# clear fdp counters
```


Network Configuration Protocol

• NETCONF protocol introduction.....	129
• NETCONF in client/server architecture.....	130
• Basic NETCONF operations.....	133
• Clients establishing a NETCONF session with NetIron devices.....	151
• Data models and mapping.....	155

NETCONF protocol introduction

The Network Configuration protocol (NETCONF) uses Extensible Markup Language (XML) for automated configuration management. The NETCONF protocol runs on top of a secure transport, such as Secure Shell version 2 (SSHv2). Only one NETCONF session is supported at a time and any new NETCONF connection requests are rejected after the first session has been established.

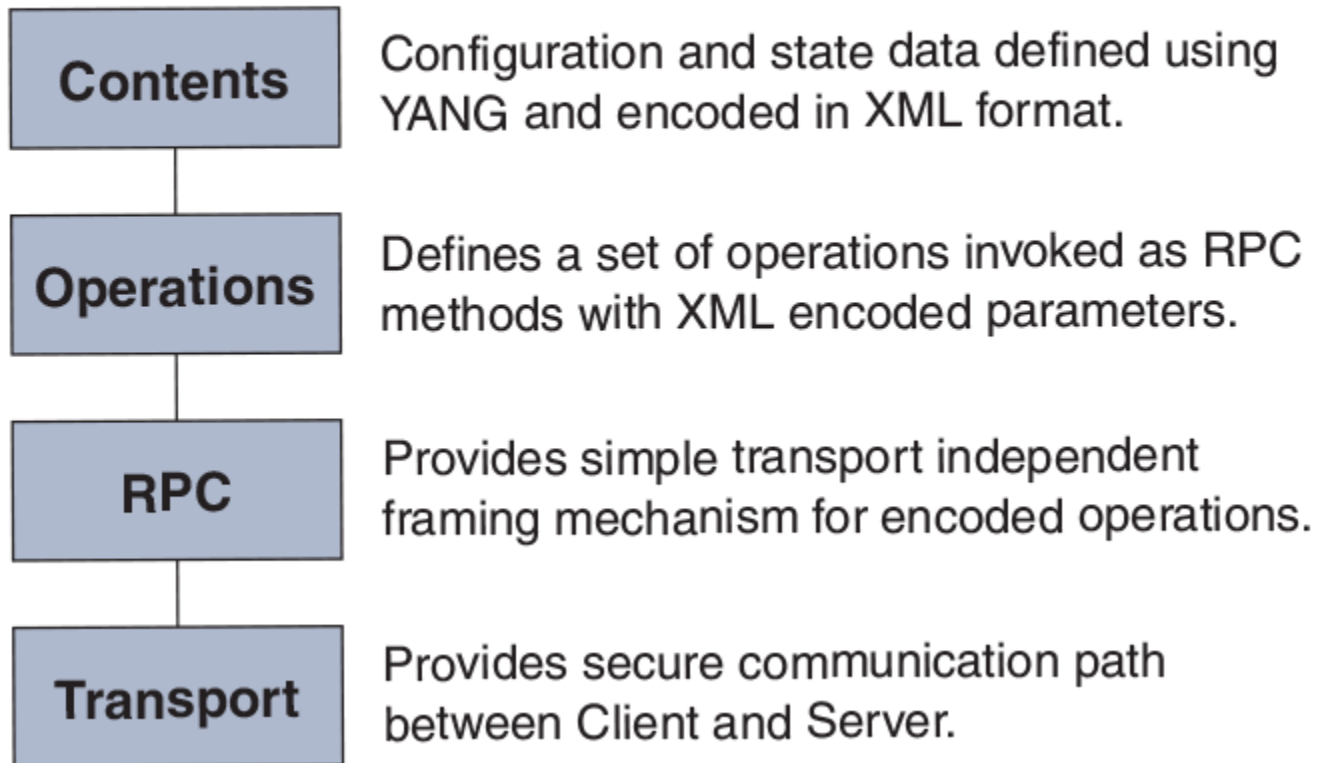
NETCONF provides mechanisms through which you can do the following:

- Manage multiple network devices
- Retrieve full or partial configuration and state data
- Upload and manipulate new configurations

Figure 2 illustrates NETCONF conceptually partitioned into four layers.

FIGURE 2 Four layers of NETCONF

Four layers of NETCONF



Platforms

NETCONF is supported on the MLX Series, XMR Series, CER 2000 Series, and CES 2000 Series devices.

Related documentation

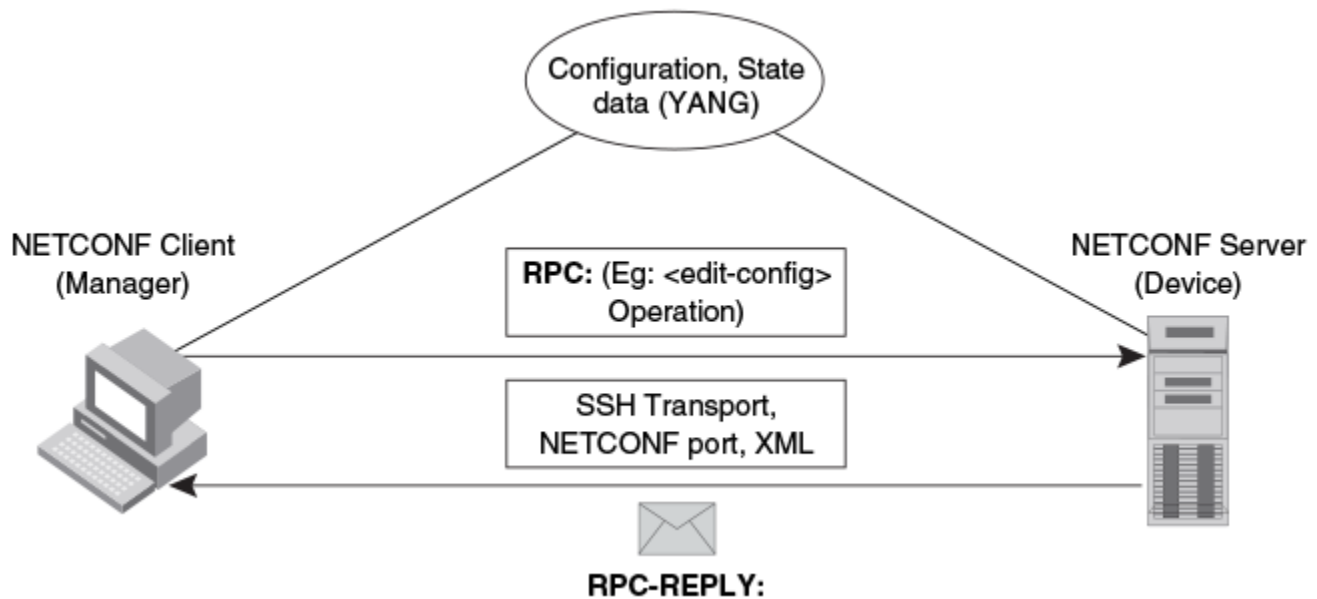
For detailed information about NETCONF, refer to RFC 4741.

For detailed information about using the NETCONF protocol over the Secure Shell (SSH), refer to RFC 4742.

NETCONF in client/server architecture

The NETCONF protocol uses a Remote Procedure Call (RPC) paradigm to facilitate communication between the client (NETCONF Manager or application) and the server (NETCONF Agent or device). A client encodes an RPC in XML and sends it to a server using a secure, connection-oriented session. The server responds with a reply encoded in XML. [Figure 3](#) illustrates the NETCONF communication between a client and a server.

FIGURE 3 NETCONF communication



The communication between the client and server consists of a series of alternating request and reply messages. The NETCONF peers use `<rpc>` and `<rpc-reply>` elements to provide transport protocol-independent framing of NETCONF requests and responses. The NETCONF server processes the RPC requests sequentially in the order in which they are received.

RPC request

The `<rpc>` element is used for enclosing a NETCONF request sent from the client to the server. Every `<rpc>` element contains a mandatory attribute, the `message-id`. This attribute has a unique value for every RPC request, and is used to associate every RPC request with the corresponding response. The `message-id` value is a monotonically increasing integer string. The maximum length of the string is 4095 characters. If the `message-id` is not present in the RPC request, the server rejects the request by returning an `<rpc-error>` with the `error-tag` element set to the `missing-attribute`.

If there are any additional attributes present in the RPC request, the NETCONF server returns them unmodified in the corresponding RPC reply.

RPC reply

An `<rpc-reply>` element is sent in response to every RPC request. The `<rpc-reply>` element contains the mandatory attribute `message-id` copied from the corresponding RPC request, along with any additional attributes that are present in the RPC request.

For successfully processed `get` or `get-config` requests, the response data is encoded as the content of the `<rpc-reply>` element.

For successfully processed `edit-config` or `close-session` requests, the `<ok>` element is encoded as the content of the `<rpc-reply>` element.

For unsuccessful RPC requests, one or more `<rpc-error>` elements are encoded inside the `<rpc-reply>` element.

RPC and error handling

If the RPC request fails, an `<rpc-error>` element, the first detected error, is encoded inside the `<rpc-reply>` element and sent to the client. The server is not required to detect or report multiple errors. If the server detects multiple errors then the order of the error detection and reporting is at the discretion of the server.

CLI and SSH subsystem

The NETCONF client must use Secure Shell Version 2 (SSHv2) as the network transport to connect to the NETCONF server. Only the SSHv2 protocol is supported as the NETCONF transport protocol.

To run NETCONF over SSHv2, the client establishes an SSH transport connection using the SSH transport protocol to the NETCONF port. The default NETCONF port is 830. The underlying SSH client and server exchange keys for message integrity and encryption.

The SSHv2 client invokes the `ssh-userauth` service to authenticate the user. All currently supported SSH user authentication methods such as the public-key, password, and keyboard-interactive authentications are supported for a NETCONF session also. If the SSH user authentication is disabled, the user is allowed full access.

On successful user authentication, the client invokes the `ssh-connection` service, also known as the SSH connection protocol. After the SSH session is established, the NETCONF client invokes NETCONF as an SSH subsystem called `netconf`.

NETCONF user privileges

Every NETCONF session has a corresponding authentication, authorization, and accounting (AAA) session. The AAA attributes apply to the NETCONF session. Only authentication and EXEC authorization are supported. Other forms of accounting and command authorization are not supported.

The privilege level of the user (read-only(5), read-write(0)) is obtained from the AAA server, if it is provided. If the privilege level is not provided by the AAA server, the default privilege level applies for the NETCONF session.

Table 38 provides the mapping between the NETCONF privilege levels and the AAA privilege levels.

TABLE 38 Privilege levels

AAA privilege level	NETCONF privilege level
0	NETCONF_PRIVILEGE_LEVEL_0
1-5	NETCONF_PRIVILEGE_LEVEL_5

Table 39 provides the mapping between the NETCONF privilege levels and the supported NETCONF operations.

TABLE 39 NETCONF operations and privilege levels

Operations	NETCONF_PRIVILEGE_LEVEL_0	NETCONF_PRIVILEGE_LEVEL_5
<code><get></code>	Yes	Yes
<code><get-config></code>	Yes	Yes
<code><edit-config></code>	Yes	No
<code><close-session></code>	Yes	Yes

Recommendations for NETCONF

- Use an authentication method to secure the underlying SSH session and to prevent any unauthorized access.

- Use a NETCONF client to generate the RPCs. If you have manually written the XML requests, recycling the XML from a successful request is recommended.
- Refer to the *MLX Series and Extreme Netlron Family YANG guide* for XML and data verification.
- Plan the configuration or state information that must be sent or retrieved to avoid sending and receiving large RPC messages.
- Use of a scripting language or other custom interface is recommended.

Basic NETCONF operations

The NETCONF protocol provides a small set of low-level operations to manage device configurations and retrieve device state information. The base protocol provides operations to retrieve, configure, copy, and delete configuration data stores. Additional operations are provided based on the capabilities advertised by the device.

The following base protocol operations are supported:

- `get`
- `get-config`
- `edit-config`
- `close-session`

NOTE

Other operations, including `copy-config`, `delete-config`, `lock`, `unlock`, and `kill-session` are not supported.

Initial connection

Each NETCONF session begins with a handshake in which the NETCONF server and the client specify the NETCONF capabilities they support. The following sections describe how to start a NETCONF session.

Hello messages

After establishing a secure transport connection, both the NETCONF server and client send a `<hello>` element simultaneously to announce their capabilities and session identifier.

After sending the hello message, the server starts the hello timer (default is 600 seconds) and waits for the hello message from the client. If no hello message is received by the server before the hello timer expires, the server aborts the NETCONF session by closing the underlying SSH session.

The NETCONF server must include the `<session-id>` element, which contains the unique session value for the NETCONF session, in the `<hello>` element. If the client receives the `<hello>` element without the `<session-id>`, the client aborts the NETCONF session by closing the underlying SSH session.

The NETCONF client must not include the `<session-id>` element in the `<hello>` element. If the server receives the `<hello>` element with the `<session-id>`, the server aborts the NETCONF session by closing the underlying SSH session.

The NETCONF client must include a valid `xmlns` attribute in the `<hello>` element. If the server receives the `<hello>` element without a valid `xmlns` attribute, the server aborts the NETCONF session by closing the underlying SSH session.

The NETCONF client must include a base capability. The server receiving the `<hello>` element without a NETCONF base capability aborts the NETCONF session by closing the underlying SSH session.

The server receiving the `<rpc>` element without receiving the `<hello>` element aborts the NETCONF session by closing the underlying SSH session.

The NETCONF client may send arbitrary data before sending a valid hello message. The server discards the data until a valid <hello> element is received from the client.

The following is an example for a <hello> element from the NETCONF server.

```
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <capabilities>
    <capability>
      urn:ietf:params:netconf:base:1.0
    </capability>
    <capability>
      urn:ietf:params:netconf:capability:writable-running:1.0
    </capability>
  </capabilities>
  <session-id>4</session-id>
</hello>
```

Capabilities

A NETCONF capability is a set of protocol extensions that supplements the base NETCONF specification. A NETCONF capability is identified with a Uniform Resource Identifier (URI). Capabilities augment the base operations of the NETCONF server, describing both the additional operations and the contents allowed inside the operations. To support a capability, the NETCONF server must support all the dependent capabilities.

The following capabilities are supported on the NetIron platforms:

- Base capability: The base capability is the set of operations and contents that any NETCONF implementation must support. The URI for the base capability is `urn:ietf:params:xml:ns:netconf:base:1.0`. Both the NETCONF client and server must support the base capability.
- Writable-running capability: The writable-running capability indicates that the device supports `edit-config` and `copy-config` operations where the `<running>` configuration is the target. The URI is `urn:ietf:params:netconf:capability:writable-running:1.0`.

NOTE

Other capabilities, including Candidate Configuration Capability, Confirmed Commit Capability, and Validate Capability, are not supported.

get operation

The NETCONF <get> operation retrieves the devices and the state data, or a filtered subset of the data.

If the device can satisfy the request, the server sends an <rpc-reply> element containing a <data> element with the results of the query. If the request cannot be completed, an <rpc-error> element is included in the <rpc-reply> element.

Parameter

The <get> operation uses the `filter` parameter. The `filter` parameter specifies the portion of the system data to retrieve. If this parameter is not present, show version information is returned.

Examples

The following is an example of a <get> operation:

```
<nc:rpc message-id="1" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
  <nc:get>
</nc:get>
```

```
</nc:rpc>
]]>]]>
```

The following is an example of a <get> operation for MPLS state data.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/
netiron-config/" message-id="25">
  <nc:get>
  <nc:filter >
    <brcd:netiron-statedata>
    <brcd:mpls-statedata/>
  </brcd:netiron-statedata>
</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>
```

The following is an example of a <get> operation for a specific LSP.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/
netiron-config/" message-id="25">
  <nc:get>
  <nc:filter nc:type="subtree" >
    <brcd:netiron-statedata>
    <brcd:mpls-statedata>
    <brcd:mpls-lsp-statedata>
    <brcd:name>scriptlsp1001</brcd:name>
  </brcd:mpls-lsp-statedata>
</brcd:mpls-statedata>
</brcd:netiron-statedata>
</nc:filter>
</nc:get>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/
config/netiron-config/" message-id="25">
  <nc:data>
  <netiron-statedata xmlns="http://brocade.com/ns/netconf/config/netiron-config/">
    <brcd:mpls-statedata>
    <brcd:mpls-lsp-statedata>
      <brcd:name>scriptlsp1001</brcd:name>
      <brcd:to>10.0.0.1</brcd:to>
      <brcd:admin-state>
        <brcd:up></brcd:up>
      </brcd:admin-state>
      <brcd:oper-state>
        <brcd:down></brcd:down>
      </brcd:oper-state>
      <brcd:tunnel-intf>tnl0</brcd:tunnel-intf>
      <brcd:up-dn-times>0</brcd:up-dn-times>
      <brcd:retry-no>273</brcd:retry-no>
    </brcd:mpls-lsp-statedata>
  </brcd:mpls-statedata>
</netiron-statedata>
</nc:data>
</nc:rpc-reply>
]]>]]>
```

The following is an example of a <get> operation for VLAN state data.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/
netiron-config/" message-id="25">
  <nc:get>
  <nc:filter >
    <brcd:netiron-statedata>
    <brcd:vlan-statedata/>
  </brcd:netiron-statedata>
</nc:filter>
</nc:get>
```

```
</nc:rpc>
]]>]]>
```

The following is an example of a <get> operation for VLAN 1001.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/" message-id="25">
  <nc:get>
    <nc:filter nc:type="subtree" >
      <brcd:netiron-statedata>
        <brcd:vlan-statedata>
          <brcd:vlan>
            <brcd:vlan-id>1001</brcd:vlan-id>
          </brcd:vlan>
        </brcd:vlan-statedata>
      </brcd:netiron-statedata>
    </nc:filter>
  </nc:get>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/" message-id="25">
  <nc:data>
    <netiron-statedata xmlns="http://brocade.com/ns/netconf/config/netiron-config/">
      <brcd:vlan-statedata>
        <brcd:vlan>
          <brcd:vlan-id>1001</brcd:vlan-id>
          <brcd:vlan-name>scriptVlan1001</brcd:vlan-name>
          <brcd:topo-hw-idx>65535</brcd:topo-hw-idx>
          <brcd:topo-sw-idx>257</brcd:topo-sw-idx>
          <brcd:topo-next-vlan>0</brcd:topo-next-vlan>
          <brcd:port>
            <brcd:port-id>ethernet 1/20 </brcd:port-id>
            <brcd:tag-mode>TAGGED</brcd:tag-mode>
            <brcd:state>DISABLED</brcd:state>
          </brcd:port>
          <brcd:bytes-received>0</brcd:bytes-received>
        </brcd:vlan>
        <brcd:dual-mode>ethernet 1/1 to 1/20 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 2/1 to 2/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 5/1 to 5/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 7/1 to 7/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 8/1 to 8/48 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 9/1 to 9/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 10/1 to 10/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 11/1 to 11/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 12/1 to 12/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 13/1 to 13/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 14/1 to 14/8 </brcd:dual-mode>
        <brcd:dual-mode>ethernet 15/1 to 15/2 </brcd:dual-mode>
        <brcd:default-vlan-id>1</brcd:default-vlan-id>
        <brcd:control-vlan-id>4095</brcd:control-vlan-id>
        <brcd:maximum-port-vlan-entries>4095</brcd:maximum-port-vlan-entries>
      </brcd:vlan-statedata>
    </netiron-statedata>
  </nc:data>
</nc:rpc-reply>
]]>]]>
```

The following is an example of a <get> operation for Interface state data.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/" message-id="25">
  <nc:get>
    <nc:filter >
      <brcd:netiron-statedata>
        <brcd:interface-statedata/>
      </brcd:netiron-statedata>
    </nc:filter>
  </nc:get>
</nc:rpc>
]]>]]>
```



```
</nc:rpc>
]]>]]>
```

The following is an example of a `<get>` operation for a specific Interface state data.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/" message-id="25">
  <nc:get>
    <nc:filter nc:type="subtree" >
      <brcd:netiron-statedata>
        <brcd:interface-statedata>
          <brcd:interface>
            <brcd:interface-id>ethernet 1/20</brcd:interface-id>
          </brcd:interface>
        </brcd:interface-statedata>
      </brcd:netiron-statedata>
    </nc:filter>
  </nc:get>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/" message-id="25">
  <nc:data>
    <netiron-statedata xmlns="http://brocade.com/ns/netconf/config/netiron-config/">
      <brcd:interface-statedata>
        <brcd:interface>
          <brcd:interface-id>ethernet 1/20</brcd:interface-id>
          <brcd:link-state>
            <brcd:down></brcd:down>
          </brcd:link-state>
          <brcd:l2-state>
            <brcd:disabled></brcd:disabled>
          </brcd:l2-state>
          <brcd:duplex>
            <brcd:none></brcd:none>
          </brcd:duplex>
          <brcd:speed></brcd:speed>
          <brcd:tag-mode>
            <brcd:yes></brcd:yes>
          </brcd:tag-mode>
          <brcd:priority-level>
            <brcd:level0></brcd:level0>
          </brcd:priority-level>
          <brcd:mac-address>0000.0085.2d00</brcd:mac-address>
        </brcd:interface>
      </brcd:interface-statedata>
    </netiron-statedata>
  </nc:data>
</nc:rpc-reply>
]]>]]>
```

get-config operation

The NETCONF `<get-config>` operation retrieves all or part of a configuration from the source data store. The `<get-config>` operation is similar to the **show running-config** command.

If the device can satisfy the request, the server sends an `<rpc-reply>` element containing a `<data>` element with the results of the query. If the request cannot be completed, an `<rpc-error>` element is included in the `<rpc-reply>` element.

Parameters

The parameters used for `<get-config>` are as follows:

- `source`: Name of the configuration data store being queried, such as `<running/>`. Only running configuration data store is supported.

- `filter`: Specifies the portions of the device configuration to retrieve. If this parameter is not present, no configuration is returned. The `filter` parameter must contain a `type` attribute. This attribute indicates the type of filtering syntax used within the `filter` parameter. The subtree filtering is the default filtering mechanism used in NETCONF.

NOTE

xpath filtering is not supported.

Examples

The following is an example of a `<get-config>` operation for MPLS configuration.

```
<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
  <nc:get-config>
    <nc:source>
      <nc:running/>
    </nc:source>
    <nc:filter nc:type="subtree">
      <brcd:netiron-config>
      <brcd:mpls-config/>
    </brcd:netiron-config>
    </nc:filter>
  </nc:get-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
message-id="1">
  <nc:data>
    <brcd:netiron-config>
    <brcd:mpls-config>
    <brcd:path>
      <brcd:name>example</brcd:name>
      <brcd:strict>10.99.161.1</brcd:strict>
    </brcd:path>
    <brcd:path>
      <brcd:name>example2</brcd:name>
      <brcd:strict>10.99.145.1</brcd:strict>
    </brcd:path>
    <brcd:lsp>
      <brcd:name>examplelsp1</brcd:name>
      <brcd:adaptive></brcd:adaptive>
      <brcd:from>10.99.10.1</brcd:from>
      <brcd:to>10.99.161.1</brcd:to>
      <brcd:enable></brcd:enable>
      <brcd:hop-limit>10</brcd:hop-limit>
      <brcd:ipmtu>1526</brcd:ipmtu>
      <brcd:ldp-tunneling></brcd:ldp-tunneling>
      <brcd:metric>600</brcd:metric>
      <brcd:primary-path>example</brcd:primary-path>
      <brcd:record></brcd:record>
      <brcd:reoptimize-timer>3600</brcd:reoptimize-timer>
      <brcd:revert-timer>43200</brcd:revert-timer>
      <brcd:mpls-traffic-eng>
        <brcd:max-burst>44736</brcd:max-burst>
        <brcd:max-rate>6312</brcd:max-rate>
        <brcd:mean-rate>1544</brcd:mean-rate>
      </brcd:mpls-traffic-eng>
      <brcd:secondary-path>
        <brcd:name>example2</brcd:name>
      </brcd:secondary-path>
    </brcd:lsp>
    <brcd:lsp>
      <brcd:name>examplelsp2</brcd:name>
      <brcd:adaptive></brcd:adaptive>
      <brcd:from>10.99.10.1</brcd:from>
```

```

<brcd:to>10.99.145.1</brcd:to>
<brcd:enable></brcd:enable>
<brcd:hop-limit>10</brcd:hop-limit>
<brcd:ipmtu>1526</brcd:ipmtu>
<brcd:ldp-tunneling></brcd:ldp-tunneling>
<brcd:metric>600</brcd:metric>
<brcd:primary-path>example2</brcd:primary-path>
<brcd:record></brcd:record>
<brcd:reoptimize-timer>3600</brcd:reoptimize-timer>
<brcd:revert-timer>43200</brcd:revert-timer>
<brcd:mpls-traffic-eng>
  <brcd:max-burst>44736</brcd:max-burst>
  <brcd:max-rate>6312</brcd:max-rate>
  <brcd:mean-rate>1544</brcd:mean-rate>
</brcd:mpls-traffic-eng>
<brcd:secondary-path>
  <brcd:name>example</brcd:name>
</brcd:secondary-path>
</brcd:lsp>
<brcd:lsp>
  <brcd:name>exmaplelsp</brcd:name>
  <brcd:from>10.99.10.1</brcd:from>
  <brcd:to>10.99.161.1</brcd:to>
  <brcd:disable></brcd:disable>
  <brcd:record></brcd:record>
</brcd:lsp>
</brcd:mpls-config>
</brcd:netiron-config>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

The following is an example of a <get-config> operation for VLAN configuration.

```

<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/
netiron-config/"
message-id="1">
  <nc:get-config>
    <nc:source>
      <nc:running/>
    </nc:source>
    <nc:filter nc:type="subtree">
      <brcd:netiron-config>
      <brcd:vlan-config/>
    </brcd:netiron-config>
    </nc:filter>
  </nc:get-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/
config/netiron-config/"
message-id="1">
  <nc:data>
    <brcd:netiron-config>
    <brcd:vlan-config>
    <brcd:vlan>
      <brcd:vlan-id>1</brcd:vlan-id>
      <brcd:vlan-name>DEFAULT-VLAN</brcd:vlan-name>
      <brcd:untagged>ethernet 1/1 </brcd:untagged>
      <brcd:untagged>ethernet 1/3 to 1/24 </brcd:untagged>
    </brcd:vlan>
    <brcd:vlan>
      <brcd:vlan-id>100</brcd:vlan-id>
      <brcd:vlan-name></brcd:vlan-name>
      <brcd:untagged>ethernet 2/1 to 2/2 </brcd:untagged>
      <brcd:router-interface>ve 100</brcd:router-interface>
    </brcd:vlan>
    <brcd:vlan>
      <brcd:vlan-id>200</brcd:vlan-id>
      <brcd:vlan-name></brcd:vlan-name>
      <brcd:tagged>ethernet 1/1 </brcd:tagged>

```

```

    <brcd:tagged>ethernet 1/3 </brcd:tagged>
    <brcd:tagged>ethernet 1/5 </brcd:tagged>
  </brcd:vlan>
  <brcd:vlan>
    <brcd:vlan-id>300</brcd:vlan-id>
    <brcd:vlan-name></brcd:vlan-name>
    <brcd:untagged>ethernet 1/2 </brcd:untagged>
  </brcd:vlan>
  <brcd:vlan>
    <brcd:vlan-id>4095</brcd:vlan-id>
    <brcd:vlan-name>CONTROL-VLAN</brcd:vlan-name>
    <brcd:tagged>ethernet 1/1 to 1/24 </brcd:tagged>
    <brcd:tagged>ethernet 2/1 to 2/2 </brcd:tagged>
  </brcd:vlan>
  <brcd:default-vlan-id>1</brcd:default-vlan-id>
</brcd:vlan-config>
</brcd:netiron-config>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

The following is an example of a <get-config> operation for Interface configuration.

```

<nc:rpc xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/config/
netiron-config/"
message-id="1">
  <nc:get-config>
  <nc:source>
  <nc:running/>
</nc:source>
  <nc:filter nc:type="subtree">
    <brcd:netiron-config>
    <brcd:interface-config/>
  </brcd:netiron-config>
</nc:filter>
</nc:get-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/
config/netiron-config/"
message-id="1">
  <nc:data>
    <brcd:netiron-config>
    <brcd:interface-config>
    <brcd:interface>
      <brcd:interface-id>management 1</brcd:interface-id>
      <brcd:enable></brcd:enable>
      <brcd:ip>
        <brcd:address>10.20.99.187/20</brcd:address>
      </brcd:ip>
      <brcd:ipv6>
        <brcd:address>2001:db8::10:20:99:187/64</brcd:address>
      </brcd:ipv6>
      <brcd:priority>
      </brcd:priority>
    </brcd:interface>
    <brcd:interface>
      <brcd:interface-id>ethernet 1/1</brcd:interface-id>
      <brcd:disable></brcd:disable>
      <brcd:loop-detection>
      </brcd:loop-detection>
      <brcd:flow-control></brcd:flow-control>
      <brcd:speed-duplex>auto</brcd:speed-duplex>
      <brcd:priority>
      </brcd:priority>
    </brcd:interface>
    <brcd:interface>
      <brcd:interface-id>ethernet 1/2</brcd:interface-id>
      <brcd:disable></brcd:disable>
      <brcd:loop-detection>
      </brcd:loop-detection>

```

```

    <brcd:flow-control></brcd:flow-control>
    <brcd:speed-duplex>auto</brcd:speed-duplex>
    <brcd:priority>
  </brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/3</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/4</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/5</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/6</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/7</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/8</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/9</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>

```

```

    <brcd:priority>
  </brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/10</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/11</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/12</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/13</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/14</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/15</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/16</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>

```

```

</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/17</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/18</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/19</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/20</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/21</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/22</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 1/23</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:speed-duplex>auto</brcd:speed-duplex>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>

```

```

    <brcd:interface-id>ethernet 1/24</brcd:interface-id>
    <brcd:disable></brcd:disable>
    <brcd:loop-detection>
  </brcd:loop-detection>
    <brcd:flow-control></brcd:flow-control>
    <brcd:speed-duplex>auto</brcd:speed-duplex>
    <brcd:priority>
  </brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 2/1</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ethernet 2/2</brcd:interface-id>
  <brcd:disable></brcd:disable>
  <brcd:loop-detection>
</brcd:loop-detection>
  <brcd:flow-control></brcd:flow-control>
  <brcd:priority>
</brcd:priority>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>ve 100</brcd:interface-id>
  <brcd:enable></brcd:enable>
  <brcd:ip>
    <brcd:address>10.2.2.2/24</brcd:address>
  </brcd:ip>
</brcd:interface>
<brcd:interface>
  <brcd:interface-id>loopback 1</brcd:interface-id>
  <brcd:enable></brcd:enable>
  <brcd:ip>
    <brcd:address>10.13.32.1/32</brcd:address>
  </brcd:ip>
</brcd:interface>
</brcd:interface-config>
</brcd:netiron-config>
</nc:data>
</nc:rpc-reply>
]]>]]>

```

Error handling for get or get-config operations

A 32K response size limit is supported in releases prior to NetIron 6.0.00. In NetIron 6.0.00 and later, 512K response size limit is supported. An error is returned, if the response size limit is exceeded.

The following error response is generated when a client makes a NETCONF RPC request, resulting in a response size limit that exceeds 512 kilobytes.

```

<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <rpc-error>
    <error-type>rpc</error-type>
    <error-tag>too-big</error-tag>
    <error-severity>error</error-severity>
    <error-message>Response buffer exceeded</error-message>
  </rpc-error>
</rpc-reply>

```

NOTE

Refer to the examples in get or get-config operations for a specific interface state data information.

edit-config operation

The NETCONF `<edit-config>` operation loads all the configurations into the specified target configuration.

Elements in the `<config>` subtree may contain an `operation` attribute. The attribute identifies the point in the configuration to perform the operation and might appear on multiple elements throughout the `<config>` subtree.

The operation attribute contains any one of the following values: `merge`, `replace`, `create`, `delete`.

The values `merge`, `replace`, or `create` is enforced by the behavior of the individual CLI, so these options are ignored. The `delete` operation alone is supported.

Parameters

The parameters used for `<edit-config>` are as follows:

- `target`: Name of the configuration data store being edited, such as `<nc:running/>`.
- `test-option`: This option is not supported.
- `default-operation`: Only the `none` value is supported. The other values such as `merge`, `replace`, `create`, and `delete` are ignored because the behaviors are enforced by the individual CLI.
- `error-option`: Only the `stop-on-error` option is supported. The other values such as `continue-on-error` and `rollback-on-error` are ignored.
- `config`: A hierarchy of configuration data as defined by the data models of the device. The new configuration must be inline configuration and other configuration options such as local file, remote file, and URL are not supported.

Examples

The following is an example for an `<edit-config>` operation for MPLS configuration.

```
<nc:rpc message-id="1" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://
brocade.com/ns/netconf/config/netiron-config/">
  <nc:edit-config>
    <nc:target>
      <nc:running/>
    </nc:target>
    <nc:default-operation>merge</nc:default-operation>
    <nc:config>
      <brcd:netiron-config>
      <brcd:mpls-config>
      <brcd:lsp nc:operation="delete">
        <brcd:name>examplelsp2</brcd:name>
      </brcd:lsp>
    </brcd:mpls-config>
  </brcd:netiron-config>
</nc:config>
</nc:edit-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/
config/netiron-config/"
message-id="1">
  <nc:ok></nc:ok>
</nc:rpc-reply>
]]>]]>
```

The following is an example for an `<edit-config>` operation for VLAN configuration.

```
<nc:rpc message-id="1" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://
brocade.com/ns/netconf/config/netiron-config/">
  <nc:edit-config>
```

```

<nc:target>
<nc:running/>
</nc:target>
<nc:default-operation>merge</nc:default-operation>
<nc:config>
<brcd:netiron-config>
<brcd:vlan-config>
<brcd:vlan nc:operation="delete">
<brcd:vlan-id>200</brcd:vlan-id>
</brcd:vlan>
</brcd:vlan-config>
</brcd:netiron-config>
</nc:config>
</nc:edit-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/
config/netiron-config/"
message-id="1">
<nc:ok></nc:ok>
</nc:rpc-reply>
]]>]]>

```

The following is an example for an `<edit-config>` operation to configure interface ethernet 1/1 with the IP address of 10.1.1.1/24 and enable the interface.

```

<nc:rpc message-id="1" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/
netconf/config/netiron-config/">
<nc:edit-config>
<nc:target>
<nc:running/>
</nc:target>
<nc:config>
<brcd:netiron-config>
<brcd:interface-config>
<brcd:interface>
<brcd:interface-id>ethernet 1/1</brcd:interface-id>
<brcd:enable></brcd:enable>
<brcd:ip>
<brcd:address>10.1.1.1/24</brcd:address>
</brcd:ip>
</brcd:interface>
</brcd:interface-config>
</brcd:netiron-config>
</nc:config>
</nc:edit-config>
</nc:rpc>
]]>]]>
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0" xmlns:brcd="http://brocade.com/ns/netconf/
config/netiron-config/"
message-id="1">
<nc:ok></nc:ok>
</nc:rpc-reply>
]]>]]>

```

Error handling

The `error-option` element contains the `stop-on-error` value. The `stop-on-error` value aborts the `edit-config` operation on the first error. All the configuration items before the error are already applied on the system. This is the default error option.

After receiving the complete `edit-config` RPC, the configuration items specified in the XML are applied sequentially as per the order specified in the YANG. If all the configuration items are successfully applied, an `<ok>` element is sent in the `<rpc-reply>` element. Otherwise, an `<rpc-error>` element with the details of the error is sent in the `<rpc-reply>` element.

Closing sessions

The NETCONF `<close-session>` operation is used for gracefully closing the current NETCONF session. The `<close-session>` operation uses no additional parameters.

When a NETCONF server receives a `<close-session>` request, the server releases resources associated with the session and closes the underlying SSH connection. Any NETCONF requests received after a `<close-session>` request are ignored. If the device is able to close the connection, an `<rpc-reply>` element is sent that includes an `<ok>` element. Otherwise, an `<rpc-error>` element with the details of the error is sent in the `<rpc-reply>` element.

NETCONF commands and specifications

The following sections describe the configuration of NETCONF using the CLI and the associated show commands, the syslog messages, and the system limitations of NETCONF.

Configuring NETCONF server

To enable the NETCONF server on a device, enter the following command.

```
device(config)# netconf server
device# netconf server?
server    Enable NETCONF server functionality
```

When no port number is specified, the command applies to the default port (830).

To enable the NETCONF server for a specific port, enter the following command.

```
device(config)# netconf server port 2001
```

Syntax: `[no] netconf server [port port-number]`

The **port** option allows you to enable NETCONF on a non-default port.

The *port-number* variable specifies the port number of the device. The range is from 1 through 65535.

Both the SSH server and the NETCONF server must be enabled to establish a NETCONF session. The **netconf server** command displays the following warning message if the SSH server configuration is disabled.

```
Warning: SSH server is disabled. Please enable the SSH server.
```

Configuring session hello-timeout

A NETCONF session hello-timeout indicates the number of seconds a session waits before the hello message is received from the NETCONF client. A session is dropped if no hello message is received before the specified number of seconds elapses. If this parameter is set to zero, the server never drops a session.

NOTE

Setting the NETCONF session hello-timeout value to zero permits denial of service attacks.

To configure a NETCONF session hello-timeout, enter the following command.

```
device(config)# netconf hello-timeout 300
```

Syntax: `[no] netconf hello-timeout [seconds]`

The *seconds* variable specifies the number of seconds the server waits to receive a hello message. The range is from 1 through 3600 seconds. The default value is 600 seconds.

Configuring session idle-timeout

A NETCONF session idle-timeout indicates the number of seconds that a session may remain idle without issuing any RPC requests. A session is dropped if it is idle for an interval longer than the specified number of seconds. If this parameter is set to zero, the server never drops a session because it is idle.

To configure a NETCONF session idle-timeout, enter the following command.

```
device(config)# netconf idle-timeout 86400
```

Syntax: [no] netconf idle-timeout [seconds]

The *seconds* variable specifies the number of seconds a session remains idle. The range is from 1 through 360000 seconds. The default value is 3600 seconds.

Displaying NETCONF statistics

To display the NETCONF server level information and statistics, enter the following command.

```
device# show netconf server
NETCONF server status: Enabled, Port: 830, Transport: SSH
Start Time: Feb  4 19:20:31
Max allowed sessions: 1, Active sessions: 1
Hello timeout: 600 seconds, Idle timeout: 3600 seconds
Server statistics:
  In sessions      : 1           In bad hellos   : 0
  Dropped sessions : 0           In too big rpcs : 0
  In rpcs          : 1           In bad rpcs    : 0
  Out rpcs         : 1           Out rpc errors  : 0
  Out too big rpcs : 0
```

Syntax: show netconf server

Table 40 describes the output of the **show netconf server** command.

TABLE 40 NETCONF server parameters

Field	Description
server status	The admin status (enabled or disabled) of the NETCONF server. Also displays the SSH status, when SSH is not enabled.
Port	The NETCONF server port number.
Transport	The NETCONF transport (currently only SSH is supported).
Start Time	The time at which the NETCONF subsystem is started.
Max allowed sessions	The maximum number of simultaneous NETCONF sessions supported by the server.
Active sessions	The number of active NETCONF sessions.
Hello timeout	The NETCONF session hello message timeout in seconds.
Idle timeout	The NETCONF session idle message timeout in seconds.
In sessions	The number of sessions started.
In bad hellos	The number of sessions silently dropped because an invalid hello message was received.
Dropped sessions	The number of sessions that were abnormally terminated (for example, due to transport close).
In too big rpcs	The total number of RPC requests received by the server that are larger than the supported maximum RPC request size.
In rpcs	The total number of correct RPC requests received by the server.
In bad rpcs	The total number of incorrect RPC messages received by the server. This includes XML parse errors and errors on the RPC layer.

TABLE 40 NETCONF server parameters (continued)

Field	Description
Out rpcs	The total number of RPC reply messages sent by the server containing an <code><rpc-ok></code> element or <code><data></code> element.
Out rpc errors	The total number of RPC reply messages sent by the server containing an <code><rpc-error></code> element.
Out too big rpcs	The total number of RPC reply messages sent by the server containing an <code><rpc-error></code> element with <code>too-big</code> as the error tag.

To display the NETCONF session level statistics, enter the following command.

```
device# show netconf session
Session Id: 1  SSH session Id: 1
Username: lab  Login time: Feb  7 21:28:47
Client Ip Address: 10.120.73.112
Privilege Level: <edit-config> <get-config> <get> <close-session>
Session Statistics:
    In rpcs      : 1          In bad rpcs      : 0
    Out rpcs     : 1          Out rpc errors : 0
    Edit-Config : 0          Get-Config   : 0
```

Get : 1 Un-supported : 0

Syntax: show netconf sessions

Table 41 describes the output of the **show netconf sessions** command.

TABLE 41 NETCONF session parameters

Field	Description
Session Id	The unique identification value for the NETCONF session.
SSH session Id	The unique identification value for the SSH session.
Username	The authenticated SSH user name. The value is <code><none></code> for public key authentication.
Login time	The time at which the session is established.
Client Ip Address	The IP address of the NETCONF client.
Privilege Level	The supported NETCONF privilege level operations for a session, where privilege is derived from the SSH user privilege.
In rpcs	The number of correct RPC requests received.
In bad rpcs	The total number of incorrect RPC messages received by the server. This includes XML parse errors and errors on the RPC layer.
Out rpcs	The total number of RPC reply messages sent by the server containing an <code><rpc-ok></code> element or <code><data></code> element.
Out rpc errors	The total number of RPC reply messages sent by the server containing an <code><rpc-error></code> element.
Edit-Config	The number of well-formed <code><edit-config></code> operations received.
Get-Config	The number of well-formed <code><get-config></code> operations received.
Get	The number of well-formed <code><get></code> operations received.
Unsupported	The number of unsupported operations received.

Syslog messages for NETCONF

The following syslog message is generated when the NETCONF session is established.

```
SYSLOG: <14>Feb  8 01:03:00 NETCONF session [1] from 10.20.99.130 user ncradsuper has been established.
```

Syntax: NETCONF session id from IPaddress user username has been established

The following syslog message is generated when the NETCONF session is disconnected.

```
SYSLOG: <14>Feb  8 01:03:00 NETCONF session [1] from 10.20.99.130 user ncradsuper has been disconnected.
```

Syntax: NETCONF session *id* from *IPaddress* user *username* has been disconnected

Clearing NETCONF statistics

To clear the NETCONF server level statistics, enter the following command.

```
device# clear netconf server-stats
```

Syntax: clear netconf server-stats

To clear the NETCONF session level statistics, enter the following command.

```
device# clear netconf session-stats
```

Syntax: clear netconf session-stats

System limitations for NETCONF

The following are the system limitations for NETCONF.

- Only one NETCONF session is supported at a time. Any new NETCONF connection requests are rejected after the first session is established.
- Only the <running> configuration data store is supported.
- The <running> configuration data store displays the commands that are currently supported by NETCONF.
- The NETCONF notifications are not supported.
- A partial set of configuration and state display commands are supported.
- The XPATH filtering is not supported.
- A 32K response size limit is supported in releases prior to NetIron 6.0.00. In NetIron 6.0.00 and later, 512K response size limit is supported. An error is returned, if the response size limit is exceeded.
- A 16K request buffer limit is supported. An error is returned, if the request size limit is exceeded.
- Only a subset of the subtree filtering is supported.

Clients establishing a NETCONF session with Netlon devices

PuTTY Link

The PuTTY Link client allows a user to establish a NETCONF session with the Netlon device on the network.

- Configure Secure Shell (SSH). Refer to the *Extreme Netlon Security Configuration Guide* for more information.
 - Configure NETCONF server.
1. Connect to the device using Putty PLINK and provide the user name and password credentials.

```
D:\>D:\putty\plink -s -P 830 lab@192.0.2.0 netconf
Using keyboard-interactive authentication.
Password:
<?xml version="1.0" encoding="UTF-8"?> <hello xmlns="urn:ietf:params:xml:ns:netc
onf:base:1.0"> <capabilities> <capability>urn:ietf:params:netconf:base:1.0</capa
bility> <capability>urn:ietf:params:netconf:base:1.1</capability> <capability>ur
n:ietf:params:netconf:capability:writeable-running:1.0</capability> </capabiliti
es> <session-id>1</session-id> </hello>]]>]]>
```

2. Copy and paste the following hello message to complete session negotiation.

NOTE

To avoid any special formatting characters getting pasted, copy and paste the message to a plain text editor such as Microsoft® Notepad first. Later, copy the message from Notepad and paste the text to the session.

```
<?xml version="1.0" encoding="UTF-8"?>
  <hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
    <capabilities>
      <capability>urn:ietf:params:netconf:base:1.0</capability>
    </capabilities>
  </hello>
]]>]]>
```

3. Confirm that the session has been negotiated successfully.

```
device# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)

...

Dynamic Log Buffer (5000 lines):
Mar  8 15:53:39:I:NETCONF session [1] from 192.0.2.0 user lab has been established.

...
```

If you have debugging turned on (#debug ip netconf), then you will see the following message:

```
"NETCONF[0]: Hello message exchanged successfully."
```

4. Send the NETCONF Remote Procedure Call (RPC) on the newly established NetCONF session. To obtain the default state data, copy and paste the following code.

NOTE

To avoid any special formatting characters getting pasted, please copy and paste the message to a plain text editor such as Microsoft® Notepad first. Later, copy the message from Notepad and paste the text to the session.

```
<nc:rpc message-id="25" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
  <nc:get xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/">
    </nc:get>
  </nc:rpc>
]>]]>
```

5. Observe the output which is the response to the request issued in step 4.

```
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/" message-id="25">
  <nc:data>
    <netiron-statedata xmlns="http://brocade.com/ns/netconf/config/netiron-config/">

      <brcd:version-statedata>
        <brcd:system>NetIron CER</brcd:system>
        <brcd:system-id>
          <brcd:serial>P00524F075</brcd:serial>
          <brcd:part>40-1000347-04</brcd:part>
        </brcd:system-id>
        <brcd:license>
          <brcd:software-packaging-type>ADV_SVCS_PREM</brcd:software-packaging-type>
          <brcd:license-id>rFFKHJhFMK</brcd:license-id>
        </brcd:license>
        <brcd:cpld-version>16</brcd:cpld-version>
        <brcd:micro-controller-version>13</brcd:micro-controller-version>
      ...
    </netiron-statedata>
  </nc:data>
</nc:rpc-reply>
```

6. Enter **Ctrl-C** command to disconnect the client connection.

Netopeer

The Netopeer client allows users to establish a NETCONF session with the NetIron device on a network.

- Configure Secure Shell (SSH). Refer to the *Extreme NetIron Security Configuration Guide* for more information.

- Configure NETCONF server.

See the following example log that has a detailed debugging information for a NETCONF session with the device to obtain the VLAN configuration.

```
$ cat get-config-filter.txt
<base10:get-config xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
  xmlns:base10="urn:ietf:params:xml:ns:netconf:base:1.0">
  <base10:source>
    <base10:running/>
  </base10:source>
  <base10:filter>
    <brcd:netiron-config>
    <brcd:vlan-config>
    </brcd:vlan-config>
    </brcd:netiron-config>
  </base10:filter>
</base10:get-config>
$

$ netopeer-cli

netconf> debug
Verbose level set to DEBUG
netconf> connect --login lab 192.0.2.0
libnetconf DEBUG: child proces with PID 25163 forked
libnetconf DEBUG: waiting for a password request
libnetconf DEBUG: writing the password to ssh
Password:
libnetconf DEBUG: XML message begin found, waiting for the password finished
libnetconf DEBUG: Writing message (session ): <?xml version="1.0" encoding="UTF-8"?>
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
....

netconf> user-rpc --file get-config-filter.txt
libnetconf DEBUG: Writing message (session 1): <?xml version="1.0" encoding="UTF-8"?>
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="2">
  <base10:get-config xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/"
    xmlns:base10="urn:ietf:params:xml:ns:netconf:base:1.0">
    <base10:source>
      <base10:running/>
    </base10:source>
    <base10:filter>
      <brcd:netiron-config>
      <brcd:vlan-config>
      </brcd:vlan-config>
      </brcd:netiron-config>
    </base10:filter>
    </base10:get-config>
  </rpc>

libnetconf DEBUG: Received message (session 1): <nc:rpc-reply
xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
  xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/" message-id="2">
  <nc:data>
    <brcd:netiron-config>
    <brcd:vlan-config>
    <brcd:vlan>
      <brcd:vlan-id>1</brcd:vlan-id>
      <brcd:vlan-name>DEFAULT-VLAN</brcd:vlan-name>
      <brcd:untagged>ethernet 1/1 to 1/24 </brcd:untagged>
    </brcd:vlan>
    <brcd:vlan>
      <brcd:vlan-id>4095</brcd:vlan-id>
      <brcd:vlan-name>CONTROL-VLAN</brcd:vlan-name>
    </brcd:vlan>
    <brcd:default-vlan-id>1</brcd:default-vlan-id>
    </brcd:vlan-config>
    </brcd:netiron-config>
  </nc:data>
</nc:rpc-reply>
```

```

Result:
<brcd:netiron-config>
  <brcd:vlan-config>
    <brcd:vlan>
      <brcd:vlan-id>1</brcd:vlan-id>
      <brcd:vlan-name>DEFAULT-VLAN</brcd:vlan-name>
      <brcd:untagged>ethernet 1/1 to 1/24 </brcd:untagged>
    </brcd:vlan>
    <brcd:vlan>
      <brcd:vlan-id>4095</brcd:vlan-id>
      <brcd:vlan-name>CONTROL-VLAN</brcd:vlan-name>
    </brcd:vlan>
    <brcd:default-vlan-id>1</brcd:default-vlan-id>
  </brcd:vlan-config>
</brcd:netiron-config>
netconf> quit
$

```

Linux OpenSSH

The Linux OpenSSH client allows a user to establish a NETCONF session with the NetIron device on the network.

- Configure Secure Shell (SSH). Refer to the *Extreme NetIron Security Configuration Guide* for more information.
 - Configure NETCONF server.
1. Connect to the device using Linux OpenSSH and provide the user name and password credentials.

```

lab@monroe{347}: ssh -p 830 -s lab@10.24.12.69 netconf
Password:
<?xml version="1.0" encoding="UTF-8"?>
<hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
<capabilities> <capability>urn:ietf:params:netconf:base:1.0</capability>
<capability>urn:ietf:params:netconf:base:1.1</capability>
<capability>urn:ietf:params:netconf:capability:writeable-running:1.0</capability> </capabilities>
<session-id>1</session-id> </hello>]]>]]>

```

2. Copy and paste the following hello message to complete session negotiation.

NOTE

To avoid any special formatting characters getting pasted, copy and paste the message to a plain text editor such as Microsoft® Notepad first. Later, copy the message from Notepad and paste the text to the session.

```

<?xml version="1.0" encoding="UTF-8"?>
  <hello xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
    <capabilities>
      <capability>urn:ietf:params:netconf:base:1.0</capability>
    </capabilities>
  </hello>
]]>]]>

```

3. Confirm that the session has been negotiated successfully.

```
device# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)

...

Dynamic Log Buffer (5000 lines):
Mar  8 15:53:39:I:NETCONF session [1] from 192.0.2.0 user lab has been established.

...
```

If you have debugging turned on (#debug ip netconf), then you will see the following message:

```
"NETCONF[0]: Hello message exchanged successfully."
```

4. Send the NETCONF Remote Procedure Call (RPC) on the newly established NetCONF session. To obtain the default state data, copy and paste the following code.

NOTE

To avoid any special formatting characters getting pasted, copy and paste the message to a plain text editor such as Microsoft® Notepad first. Later, copy the message from Notepad and paste the text to the session.

```
<nc:rpc message-id="25" xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0">
  <nc:get xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/">
    </nc:get>
  </nc:rpc>
]>]]>
```

5. Observe the output which is the response to the request issued in step 4.

```
<nc:rpc-reply xmlns:nc="urn:ietf:params:xml:ns:netconf:base:1.0"
xmlns:brcd="http://brocade.com/ns/netconf/config/netiron-config/" message-id="25">
  <nc:data>
    <netiron-statedata xmlns="http://brocade.com/ns/netconf/config/netiron-config/">

    <brcd:version-statedata>
      <brcd:system>NetIron CER</brcd:system>
      <brcd:system-id>
        <brcd:serial>P00524F075</brcd:serial>
        <brcd:part>40-1000347-04</brcd:part>
      </brcd:system-id>
      <brcd:license>
        <brcd:software-packaging-type>ADV_SVCS_PREM</brcd:software-packaging-type>
        <brcd:license-id>rFFKHJhFMK</brcd:license-id>
      </brcd:license>
      <brcd:cpld-version>16</brcd:cpld-version>
      <brcd:micro-controller-version>13</brcd:micro-controller-version>
    ...
```

6. Enter **Ctrl-C** command to disconnect the client connection.

Data models and mapping

YANG is a data modeling language used to model configuration and state data manipulated by the NETCONF protocol, NETCONF remote procedure calls, and NETCONF notifications. In order for NETCONF to be an interoperable protocol, models must be defined in a vendor-neutral language. YANG provides the language and rules for defining models for use with NETCONF.

The YANG language is currently being developed by the IETF NETCONF Data Modeling Language Working Group (NETMOD) and is defined in RFC 6020.

Each block of YANG data is encapsulated as a module, containing a header statement, linkage information, meta information, and revision history. Modules can contain one or more submodules with the same structure.

The following code example shows the structure of a header statement, along with linkage and meta information, which contains contact information and a high-level description of the module.

```
module netiron-config
{
  namespace "http://brocade.com/ns/netconf/config/netiron-config/";
  prefix "brcd";
  include common-defs;
  include vlan-config;
  include interface-config;
  include mpls-config;
  organization
    "Brocade Communications Inc.";
  contact
    "Technical Support Center"+
    "130 Holger Way,"+
    "San Jose, CA 95134"+
    "Email: ipsupport@brocade.com"+
    "Phone: 1-800-752-8061"+
    "URL: www.brocade.com";
  description
    "NetIron Config module. VERSION: ";
  revision 2011-04-20
  {
    description "Initial revision";
  }
}
```

Example in YANG, XML, and CLI

Table 42 provides an example to describe the VLAN name in the YANG model and the equivalent XML and CLI.

TABLE 42 Example in YANG, XML, and CLI

YANG	XML	CLI
leaf vlan-name { type string { length "1..31"; } description "VLAN Name"; }	<brcd:vlan-name >example</brcd:vlan-name >	[no] vlan <i>vlan-id</i> [<i>name</i> <i>vlan-name</i>]

For further examples and information on the YANG model, refer to the *MLX Series and Extreme NetIron Family YANG Guide*.

Foundry Discovery Protocol

• Foundry Discovery Protocol overview.....	157
• Enabling FDP.....	157
• Advertising IPv4 or IPv6 management addresses to FDP neighbors.....	158
• Verifying FDP.....	158
• Clearing FDP statistics and neighbor information.....	160

Foundry Discovery Protocol overview

The Foundry Discovery Protocol (FDP) enables Extreme devices to advertise themselves to other Extreme devices on the network. When you enable FDP on a Extreme device, the device periodically advertises information including the following:

- Hostname (device ID)
- Product platform and capability
- Software version
- VLAN and Layer 3 protocol address information for the port sending the update. IP information is supported.

A Extreme device running FDP sends FDP updates on Layer 2 to MAC address 00-00-00-CC-CC-CC. Other Extreme devices listening on that address receive the updates and can display the information in the updates. Extreme devices can send and receive FDP updates on ethernet interfaces.

FDP is disabled by default.

NOTE

If FDP is not enabled on a Extreme device that receives an FDP update or the device is running a software release that does not support FDP, the update passes through the device at Layer 2.

Enabling FDP

A Extreme device can be enabled to send FDP packets.

FDP is disabled by default on all interfaces. FDP can be enabled globally to apply to all interfaces. If FDP is to be disabled for an individual interface, the configuration is applied in interface configuration mode. This task shows how to enable FDP globally, set some optional FDP parameters, disable FDP on one interface and reenable FDP on the interface.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Globally enable FDP.

```
device(config)# fdp run
```

3. Change the FDP update timer to send an FDP update every 120 seconds.

```
device(config)# fdp timer 120
```

By default, FDP sends an update every 60 seconds.

4. Change the FDP hold time to 360 seconds.

```
device(config)# fdp holdtime 360
```

By default, the FDP hold time is 180 seconds.

5. Enter interface configuration mode.

```
device(config)# interface ethernet 1/4
```

6. Disable FDP on Ethernet interface 1/4.

```
device(config-if-e1000-1/4)# no fdp enable
```

7. Reenable FDP on Ethernet interface 1/4.

```
device(config-if-e1000-1/4)# fdp enable
```

The following example enables FDP globally and sets the FDP timer and hold time. FDP is disabled on Ethernet interface 1/4.

```
device# configure terminal
device(config)# fdp run
device(config)# fdp timer 120
device(config)# fdp holdtime 360
device(config)# interface ethernet 1/4
device(config-if-e1000-1/4)# no fdp enable
```

Advertising IPv4 or IPv6 management addresses to FDP neighbors

When FDP is enabled, by default, the Extreme device advertises one IPv4 address and one IPv6 address to its FDP neighbors. You can configure the device to advertise only the IPv4 management address or only the IPv6 management address.

Ensure that FDP is enabled.

You can set the advertising IPv4 or IPv6 addresses to FDP neighbors configuration globally on a Layer 2 switch, or on an interface on a Layer 3 switch.

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

```
device# configure terminal
```

2. To configure a Layer 2 switch to advertise the IPv4 address, enter the following command in global configuration mode:

```
device(config)# fdp advertise ipv4
```

The following example configures a Layer 3 switch to advertise the IPv6 address.

```
device# configure terminal
device(config)# interface ethernet 1/1/4
device(config-if-e1000-1/1/4) fdp advertise ipv6
```

Verifying FDP

After enabling FDP you can verify the configuration and view FDP information.

Ensure that FDP has been enabled.

You can display the following Foundry Discovery Protocol (FDP) information:

- FDP entries for Extreme neighbors
- Individual FDP entries
- FDP information for an interface on the device you are managing
- FDP packet statistics

NOTE

NOTE

If the Extreme device has intercepted CDP updates, then the CDP information is also displayed.

1. To display a summary list of all the Extreme neighbors that have sent FDP updates to this Extreme device enter the following command:

```
device# show fdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater
(*) indicates a CDP device
Device ID      Local Int   Holdtm Capability Platform   Port ID
-----
deviceB        Eth 2/9     178    Router    NetIron Rou  Eth 2/9
```

2. To display detailed information about all the Extreme neighbors that have sent FDP updates to this Extreme device enter the following command:

```
device# show fdp neighbors detail

Device ID: deviceB configured as default VLAN1, tag-type8100
Entry address(es):
  IP address: 192.168.0.13
  IPv6 address (Global): c:a:f:e:c:a:f:e
Platform: FastIron Router, Capabilities: Router
Interface: Eth 2/9
Port ID (outgoing port): Eth 2/9 is TAGGED in following VLAN(s):
  9 10 11
Holdtime : 176 seconds
Version :
Foundry, Inc. Router, IronWare Version 07.6.01b1T53 Compiled on Aug 29
2002 at 10:35:21 labeled as B2R07601b1
```

3. To display detailed FDP entry information for a specific Extreme neighbor device, enter the following command:

```
device# show fdp entry FastIronB

Device ID: FastIronB configured as default VLAN1, tag-type8100
Entry address(es):
Platform: NetIron Router, Capabilities: Router
Interface: Eth 2/9
Port ID (outgoing port): Eth 2/9 is TAGGED in following VLAN(s):
  9 10 11
Holdtime : 176 seconds
Version :
Foundry, Inc. Router, IronWare Version 07.6.01b1T53 Compiled on Aug 29
2002 at 10:35:21 labeled as B2R07601b1
```

4. To display FDP information for a specific Ethernet interface, enter the following:

```
device# show fdp interface ethernet 2/3

FastEthernet 2/3 is up, line protocol is up
Encapsulation ethernet
Sending FDP packets every 5 seconds
Holdtime is 180 seconds
```

This example shows information for a specific Ethernet interface indicating how often the port sends FDP updates and how long neighbors that receive the updates, can hold them before discarding them.

5. To display FDP and CDP packet statistics, enter the following command:

```
device# show fdp traffic

CDP/FDP counters:
Total packets output: 6, Input: 5
Hdr syntax: 0, Chksum error: 0, Encaps failed: 0
No memory: 0, Invalid packet: 0, Fragmented: 0
Internal errors: 0
```

Clearing FDP statistics and neighbor information

FDP update information and statistics can be cleared.

Before clearing FDP information ensure that FDP is enabled.

You can clear the following FDP and CDP information:

- Information received in FDP and CDP updates
- FDP and CDP statistics

NOTE

The same commands clear information for both FDP and CDP.

1. To clear the information received in FDP updates from neighboring devices, enter the following command:

```
device# clear fdp table
```

2. To clear FDP and CDP statistics, enter the following command:

```
device# clear fdp counters
```


High Availability

• High availability overview.....	161
• Management module redundancy configuration.....	161
• Managing management module redundancy.....	162
• Monitoring management module redundancy.....	165
• Displaying switchover information.....	166
• Flash memory and auxiliary flash card file management commands.....	167
• Verifying available flash space on the management module before an image is copied.....	168

High availability overview

When redundant management modules are installed, one module becomes the active module and the other module is assigned the status of standby. Files are synchronized between them. If the active module goes down, the standby module becomes the active module providing high availability with a minimum of interruption to traffic forwarding.

When you apply power to a Extreme device with two management modules installed, by default, the management module in slot M1 becomes the active module and the module in slot M2 becomes the standby module. You can change the default active slot from M1 to M2 using the **active-management** command.

After the active and standby modules are determined, both modules boot from the source specified for the active module. The active module can boot from the following sources:

- The flash memory on the active management module
- An Auxiliary Flash card in an Auxiliary Flash slot on the active management module.

Once the modules boot, the system compares the flash code and system-config files on the standby module to the files on the active module. If the files are not the same, the files on the standby module are synchronized with those on the active module.

During normal operation, the active module handles tasks such as obtaining network topology and reachability information and determining the best paths to known destinations. The active module also monitors the standby module.

The standby module functions in an active standby mode. Configuration changes made from the CLI to the active management module are also written to the standby management module even if they are not written to flash memory. Synchronizing the system-config and running-config files on both modules allows the standby module to assume the role of active module seamlessly, if necessary.

The interface modules are not reset, and continue to forward traffic while the standby management module takes over operation of the system. The new now-active management module receives updates from the interface modules and sends verification information to the interface modules to ensure that they are synchronized. If the new active management module becomes out of sync with an interface module, information on the interface module may be overwritten, which can cause an interruption of traffic forwarding. An out of sync state should only occur if there is a layer 3 topology change elsewhere in the network during the management failover. Extreme devices support Layer 3 hitless failover with restart for high-availability routing in protocols such as BGP and OSPF. With these high-availability features enabled, when a device experiences a failover or restart, forwarding disruptions are minimized, and route flapping diminished to provide continuous service.

Management module redundancy configuration

Configuring management module redundancy consists of performing one optional task (changing the default active chassis slot) as described in the following section.

Changing the default active chassis slot

By default, the Extreme system considers the module installed in slot M1 to be the active management module. However, you can change the default active chassis slot to M2 using the **active-management** command.

The **active-management** command determines which management module will become active after a power cycle. By default, the top management module of the Extreme XMR 16000 and Extreme MLX-16 or the left management module of the Extreme XMR 4000, Extreme XMR 8000, Extreme MLX-4 and Extreme MLX-8 become active after a power cycle. This information is stored in the chassis's backplane EPROM and not in the configuration file.

To change the default active chassis slot from the default state of M1 to M2, enter the following commands.

```
device(config)# redundancy
device(config-redundancy)# active-management mgmt-2
```

Syntax: active-management mgt-module

The **mgt-module** parameter specifies the management module, either mgmt-1 or mgmt-2.

NOTE

This configuration has no effect on the **reload** and **boot** commands. It only applies to the power cycle when both management modules are installed in a chassis.

Managing management module redundancy

You can perform the following management tasks related to management module redundancy for Extreme devices:

- Perform immediate synchronization of files
- Perform a manual switchover to the standby module
- Reboot the standby module

File synchronization between active and standby management modules

Each active and standby management module contains the following files that can be synchronized between the two modules:

- **Flash code** - The flash code can include the following files:
 - monitor, which contains the Real Time Operating System (RTOS) for the management module
 - primary, which contains the primary Multi-Service IronWare image for the management module
 - secondary, which contains the secondary Multi-Service IronWare image for the management module

A Extreme Multi-Service IronWare image contains layer 1 - 3 software used by the management module.

During startup or switchover, the flash code on the active module is compared to the flash code on the standby module. If the files differ, the files on the standby module are synchronized to the files on the active module. If you update the flash code on the active module, the flash code on the standby module is automatically synchronized (without comparison) to the new file on the active module.

- **System-config file** - The flash code includes the system-config file. During startup or switchover, the system-config file on the active module is compared to the system-config file on the standby module. If the files are different, the system-config file on the standby module is synchronized with that of the active module. When you save changes to the system-config file on the active module, the system-config file on the standby module is automatically (without comparison) synchronized to match the system-config file on the active module.


- **Running-config** - The running-config file resides in the Extreme system memory, and is automatically synchronized (without comparison) between the active and the standby module at regular intervals. The default interval is 7 seconds.
- **Boot code** - Each active and standby management module also includes boot code that is run when a module boots. The boot code resides in the boot flash of each module. Boot code is synchronized between the active and standby modules, which allows the system to use an older version of boot code on the standby module if desired.

NOTE

However, when the standby module is inserted to the standby slot, the images get synchronized to the standby image.

Figure 4 shows how the files are synchronized between the active module and the standby module.

FIGURE 4 Active and standby management module file synchronization



The Extreme system allows you to perform the following file synchronization tasks:

- Compare files on the active module with files on the standby module and immediately synchronize any files that are different.
- Immediately synchronize all files between the active and standby modules.

The following sections explain how to perform these tasks.

Comparing and synchronizing files

You can initiate a comparison of the flash code, system-config, and running-config files on the active management module with these files on the standby module and synchronize the files immediately if differences exist. When you synchronize the files, the active module files are copied to the standby module, replacing the standby module files.

To compare and immediately synchronize files between the active and standby modules, enter the following command at the Privileged EXEC level.

```
device# sync-standby
```

Synchronizing files without comparison

You can synchronize the flash code, system-config file, and running-config file immediately without comparison. When you synchronize the files, active module files are copied to the standby module, replacing the files on the standby module.

To immediately synchronize the files between the active and standby modules, enter the following command at the Privileged EXEC level.

```
device# force-sync-standby
```

Manually switching over to the standby management module

You can cause the Extreme system to switch over to the standby module (and thus make it the active module). Enter the **switchover** command at the Privileged EXEC level.

```
device# switchover
```

In prior versions of the Multi-Service IronWare, typing the **switchover** command caused the Extreme device to switch control over to the redundant management module immediately without confirmation. Currently, you are presented with the question "**Are you sure?** " after the switchover command is executed. At this question, you can either type **y** to proceed with the switchover or type **n** to abort the switchover.

The following is an example of the new switchover procedure.

```
device#switchover
Are you sure? (enter 'y' or 'n'): y
```

NOTE

The switchover command should not be used immediately after downloading new code to the Extreme systems with redundant management modules.

Rebooting the active and standby management modules

You can reboot management modules, while maintaining the active and standby roles, using the **boot system** or **reload** commands. You can also reboot the standby module only, maintaining the standby role, using the **reboot-standby** command.

For example, to reboot the active and standby management modules from the primary Extreme Multi-Service IronWare image in the management module flash memory, enter the following command at the Privileged EXEC level.

```
device# boot system flash primary
device# Are you sure? (enter 'y' or 'n'): y
```

Syntax: `[no] boot system bootp | [flash primary | flash secondary] | slot number filename | tftp ip-address filename`

The **flash primary** keyword specifies the primary Extreme Multi-Service IronWare image in the management module flash memory. The **flash secondary** keyword specifies the secondary Extreme Multi-Service IronWare image in the flash memory.

For the *number* parameter, specify 1 for Auxiliary Flash slot 1 on the active management module and 2 for Auxiliary Flash slot 2 on the active management module. For the *filename* parameter, specify the name of the image on the Auxiliary flash card.

The **tftp** keyword directs the Extreme device to boot from an Extreme Multi-Service IronWare image on a TFTP server located at *ip-address* with the specified *filename*.

For example, to reboot the active and standby management modules, enter the following command at the Privileged EXEC level.

```
device# reload
```

To reboot the standby module only, enter the following command at the Privileged EXEC level.

```
device# reboot-standby
```

Monitoring management module redundancy

You can monitor the following aspects of management module redundancy:

- The status of the management modules (if a module is in active or standby mode)
- The switchover history for the management modules

The following sections explain how to monitor the management modules.

Determining management module status

You can determine the status of a management module in the following ways:

- **LEDs** - LEDs on the management module indicate whether a module is active or standby, and if the module has power.
- **Module information in software** - The module information displayed by the software indicates whether a module is active or standby.

Status LED

You can determine which management module is currently active and which is standby by observing the Active LED on each module. If this LED is on (green), the module is the active module. If this LED is off, the module is the standby module.

You can also observe the Pwr LED on each module. If this LED is on (green), the module is receiving power. If the LED is off, the module is not receiving power. (A module without power will not function as either the active or standby module.)

For information about what to do if these LED indicators are not what you expect, refer to the appropriate hardware installation guide.

Software

To display the status of the management modules using the software, enter the following command at any level.

```
device# show module
      Module
M1 (left): NI-XMR-MR Management Module   Status      Ports  Starting MAC
M2 (right): NI-XMR-MR Management Module   Standby (Ready)
...

```

The Status column indicates the module status. The management module status can be one of the following:

- **ACTIVE** - Current active management module
- **STANDBY** - Current standby management module.

The status of the standby module can be one of the following:

- **Init** - Currently initializing as the standby module
- **Ready** - Ready to take over as the active module, if necessary
- **Wait** - Waiting for boot information from the active management module
- **Sync** - Active module is currently synchronizing files on the standby module

Monitoring the status change of a module

The Extreme system now logs the status change of a module. The status change of a module is logged when the module becomes:

- **Up or Ready** - The module is running or ready to run.
- **Down** - The module is not running normally.

Upon the status change of a module, a message is logged in the syslog memory. At the CLI level, type the **show log** command to view the logged messages.

The following example displays a syslog message on an Interface Module in the Down state.

```
Feb  5 12:16:17:N:System: Module down in slot 1, reason REBOOTED. Error Code 0
```

The following example displays a syslog message on a Standby Management Module in the Down state.

```
Feb  5 14:38:58:N:System: Standby Management Module was down, reason Heartbeat Loss. Error Code 5
```

Displaying temperature information

All management, interface and switch fabric modules contain temperature sensors. By default, the Extreme system polls module temperature every 60 seconds. You can display the current temperature of the modules by entering either of the following commands:

- `show chassis`
- `show temperature`

For information about these commands, refer to the *appropriate hardware installation guide*.

Displaying switchover information

You can display the following information about a switchover:

- Redundancy parameter settings and statistics, including the number of switchovers that have occurred
- System log or traps logged on an SNMP trap receiver, including Information about whether a switchover has occurred.

To view the redundancy parameter settings and statistics, enter the following command at any level of the CLI.

```
device# show redundancy
=== MP Redundancy Settings ===
Default Active Slot = M1 (upper)
Running-Config Sync Period = 7 seconds
=== MP Redundancy Statistics ===
Current Active Session:
Active Slot=M2(lower),Standby Slot=M1(upper)(Ready State), Switchover Cause = No Switchover
Start Time = 1900-0-0  0:6:21 (Monday)
```

```

Previous Active Session #1:
Active Slot=M1(upper), Standby Slot=M2(lower), Switchover Cause = MP Upgrade to Ver3.7.0T163
Start Time = 1900-0-0 0:3:4 (Monday), End Time = 1900-0-0 0:6:21 (Monday)
Previous Active Session #2:
Active Slot = M2 (lower), Standby Slot = M1(upper), Switchover Cause = Active Rebooted
Start Time = 1900-0-0 0:1:1 (Monday), End Time = 1900-0-0 0:3:4 (Monday)
Previous Active Session #3:
Active Slot = M1 (upper), Standby Slot = M2(lower), Switchover Cause = MP Upgrade to Ver3.7.0T163
Start Time = 2036-2-6 6:43:54 (Wednesday), End Time = 1900-0-0 0:1:1 (Monday)

```

This output displays that the default active chassis slot is configured as slot M1 and the automatic synchronization interval is configured for 7 seconds. It also displays that in the current active session, the module installed in M2 is the active module, the module installed in M1 is the standby module, which is in Ready state, and no switchovers have occurred.

However, in three previous sessions, switchovers occurred. In sessions #1 and #3, the switchovers occurred because the software was upgraded to "Ver3.7.0T163". In session #2 the switchover occurred because the active module was rebooted. In sessions #1 and #3, the modules installed in M1 were the active modules, while the modules installed in M2 were the standby modules. In session #2, the module installed in M2 was the active module, while the module installed in M1 was the standby module.

To view the system log or traps logged on an SNMP trap receiver, enter the following command at any level.

```

device# show log
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 24 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
               I=informational N=notification W=warning
Static Log Buffer:
Sep 28 11:31:25:A:Power Supply 1, 1st left, not installed
Sep 28 11:31:25:A:Power Supply 3, middle left, not installed
Sep 28 11:31:25:A:Power Supply 4, middle right, failed
Sep 28 11:31:25:A:Power Supply 5, 2nd right, not installed
Dynamic Log Buffer (50 lines):
Sep 27 18:06:58:I:Interface ethernet6/2, state up
Sep 27 18:06:57:I:Interface ethernet3/2, state up
Sep 27 15:39:42:I:Interface ethernet3/2, state up
Sep 27 15:39:42:I:Interface ethernet6/2, state up
...
Sep 27 14:23:45:N:Module up in slot 6
Sep 27 14:23:45:N:Module up in slot 3
Sep 27 14:23:27:A:Management module at slot 9 state changed from standby to active

```

This output indicates that one switchover occurred.

Flash memory and auxiliary flash card file management commands

The Extreme system supports file systems in the following locations:

- Flash memory on the management module
- An Auxiliary flash card inserted in management module slots 1 or 2

[Table 43](#) outlines the root directory for each file system.

TABLE 43 Extreme file system root directories

File system	Root directory
Flash memory	/flash/
Auxiliary flash card in slot 1	/slot1/
Auxiliary flash card in slot 2	/slot2/

This section describes commands that manage the files in flash memory and on the flash cards. Use the file management commands to perform the following tasks:

- Format a flash card
- Determine the current management focus
- Switch the management focus
- Display a directory of files
- Display the contents of a file
- Display the hexadecimal output of a file
- Create a subdirectory
- Remove a subdirectory
- Rename a file
- Change the read-write attribute of a file
- Delete a file
- Recover (undelete) a file
- Append one file to another (join two files)
- Perform copy operations using the **copy** command
- Perform copy operations using the **cp** command
- Load the system software from flash memory, a flash card, or other sources during system reboot
- Change the save location of the startup-config file from the default location (flash memory) to a flash card in slot 1 or 2

You can access all file management commands at the Privileged EXEC level of the CLI.



CAUTION

Do not add or remove a flash card while a file operation involving the slot where the flash card is installed is in progress. Doing so can result in corruption of the flash card. If this occurs, you may need to reformat the flash card to make it usable again. Reformatting erases all data stored on the card.

Verifying available flash space on the management module before an image is copied

The Management Module of the Extreme system accommodates 32 MB of flash space. However, as the size of the Interface Module, Management Module, and FPGA images increase, the Management Module flash may not have enough space to accommodate these images. The space in the Management Module flash is too small to hold more than two images (primary and secondary) and hence, downloading a new image is not possible without deleting one of the images that is already present in the flash.

Before an image is copied onto the Management Module or Interface Module, the software now checks to refer to if there is enough space available in the Management Module flash to support the copy operation. If there is not enough free space available on the Management Module flash, the following error message will display on the user interface.

The 32 MB flash space is capable of holding two CES 2000 Series or CER 2000 Series images (image size is about 11 MB). However, during the TFTP copy operation, it needs more buffer space. It is not possible to copy or update an existing image to a 32 MB flash, if there are two images in the flash already. If you try to copy or update an image, the following error message is displayed.

For TFTP copy operation, the following error message is displayed.

```
device#copy tftp flash 10.20.10.62 xmr04001b1.bin primary
There is not enough space on MP flash. Please clean up MP flash and retry, or use "delete-first" option.
TFTP: Download to primary flash failed - Flash is full
```

For SCP copy operation, the following error message is displayed.

```
C:\>scp xmr04001b1.bin lab@10.22.2.21:image:primary
There is not enough space on MP flash. Please clean up MP flash and retry, or use "delete-first" option.
C:\>
```

In the example above the copy procedure is cancelled because there is not enough space on Management Module flash to copy the image. To make space for an image to be copied, you must clean up the flash space on the Management Module, and then retry copying the image again. You may also use the delete-first option, along with the CLI copy command, to make space for an image to be copied. The delete-first option allows you to delete existing target files on the Management Module flash.

The example below displays how the delete-first option is used. In this example, the existing secondary file image is removed from the flash to make space for a new image to be copied. The TFTP copy operation is able to successfully download the new image to the secondary flash.

```
device#copy tftp flash 10.53.1.82 xmr04001b1.bin secondary delete-first
Removing secondary from flash.
.....TFTP: Download to
secondary flash done.
```

When the delete-first option is used, the existing target files are deleted only if there is enough free space to accommodate the copy operation. If, after the delete-first option is used and there is still a shortage of free space then the following error message will display.

```
device#copy tftp flash 10.53.1.82 xmr04001b1.bin secondary delete-first
There will not be enough space on MP flash even after deleting the target files. Please clean up MP flash
and retry.
```

Management focus

The **management focus** determines the default file system (flash memory or the flash card inserted in slot 1 or 2) to which a file management operation applies. When you power on or reload a Extreme system, by default, the management focus is on flash memory.

You can change the management focus from flash memory to a slot and subdirectory using the **cd** or **chdir** command. (For more information, refer to [Switching the management focus](#) on page 173.)

To determine the slot and subdirectory that have the current management focus, enter the **pwd** command. (For more information about this command, refer to [Determining the current management focus](#) on page 172.)

Most file management commands provide the option of specifying the file system to which the command applies. If you want the command to apply to the file system that has the current management focus, you do not need to specify the file system. If you want the operation to apply to the file system that does not have the current management focus, you must specify one of the following keywords:

- **flash** - indicates flash memory
- **slot1** - indicates the flash card inserted in slot 1
- **slot2** - indicates the flash card inserted in slot 2

For example, if you want to display a directory of files in flash memory and flash memory has the current management focus, you do not need to specify the **flash** keyword. However, if you want to display a directory of files for slot 1 and flash memory has the current focus, you must specify the **slot1** keyword.

Flash memory file system

The flash memory file system is flat, which means that it does not support subdirectories. As a result, you cannot create or delete subdirectories in this file system using the **md /mkdir** and **rd /rmdir** commands, respectively. Also, when specifying the syntax for the various file management commands, you will not need to specify a pathname to a subdirectory because it is not possible for a subdirectory to exist.

File naming conventions

A file name in the flash memory file system can contain a maximum of 31 characters. File names are case sensitive. The flash memory file system does not accept spaces as part of a file name.

The following characters are valid in file names:

- All upper and lowercase letters
- All digits
- Any of the following special characters:
 - \$
 - %
 - '
 - -
 - _
 - @
 - ~
 - `
 - !
 - (
 -)
 - {
 - }
 - ^
 - #
 - &

Auxiliary flash card file system

The Auxiliary flash card file system is hierarchical, which means that it supports subdirectories. Therefore, you can create or delete subdirectories in this file system using the **md /mkdir** and **rd /rmdir** commands, respectively. Also, when specifying the syntax for the various file management commands, you may need to specify a pathname to a subdirectory as appropriate to manipulate a file in a subdirectory.

Auxiliary flash card subdirectories

The full path name for the location of a file can be a maximum of 256 characters. You can nest subdirectories as deep as you want as long as the full path name is 256 characters or less.

When you include a subdirectory path in a file management command, use a slash between each level. For example, to create a subdirectory for flash code and copy a flash image file to the subdirectory, enter commands such as the following.

```
device# mkdir slot1 /switchCode/initial-release
```

These commands create two levels of subdirectories on the flash card in Auxiliary flash slot 1.

File and subdirectory naming conventions

The Auxiliary flash slots supports file names of up to 32 characters. File names are not case sensitive. Thus, the software considers the name "test.cfg" and "TEST.CFG" to be the same.

Files and subdirectory names can be up to 32 characters long, including spaces and the special characters listed. The following characters are valid in file and subdirectory names:

- All upper and lowercase letters
- All digits
- Spaces
- Any of the following special characters:
 - \$
 - %
 - '
 - -
 - _
 - @
 - ~
 - `
 - !
 - (
 -)
 - {
 - }
 - ^
 - #
 - &

You can use spaces in a file or subdirectory name if you enclose the name in double quotes. For example, to specify a subdirectory name that contains spaces, enter a string such as the following: "a long subdirectory name".

A subdirectory or file name can be a maximum of 256 characters long. A complete subdirectory path name cannot contain more than 256 characters.

There is no maximum file size. A file can be as large as the available flash card space.

NOTE

Auxiliary flash card file system applies to the XMR Series and MLX Series only.

Wildcards

Commands to display a directory of files, to change the read-write attribute of a file, or to delete files accept wildcards in the file name (*file-name*). With these commands, you can use "*" (asterisk) as a wildcard for any part of the name. For example, all the following values are valid for *file-name*:

- teststartup.cfg
- test*.cfg

- nmb02200.bin
- *.bin
- m*.bin
- m*.*

Formatting a flash card

The flash cards shipped with a management module are pre-formatted for the 16 FAT file system used by the modules. If you want to use a flash card that is not formatted for the 16 FAT file system, you need to reformat the flash card before you can store files on it.



CAUTION

Make sure the flash card is empty or does not contain files you want to keep. Formatting a flash card completely erases all files on the card.



CAUTION

Once you start the formatting process, you cannot stop it. Even if you enter CTRL-C to stop the CLI output and a new prompt appears, the formatting continues. Make sure you want to format the card before you enter the command.

To reformat a flash card in slot 2 on the management module, for example, enter the following command.

```
device# format slot2
.....
.....
.....
.....
80809984 bytes total card space.
80809984 bytes available on card.
  2048 bytes in each allocation unit.
 39458 allocation units available on card.
```

Syntax: format slot1 | slot2

The **slot1** | **slot2** keyword specifies the Auxiliary flash slot that contains the flash card you are formatting.

Determining the current management focus

For conceptual information about management focus, refer to [Management focus](#) on page 169.

To determine which file system has the current management focus, enter the following command.

```
device# pwd
Flash /flash/
```

In this example, the management focus is the flash memory.

In the following example, the management focus is the root directory of the flash card in slot 1.

```
device# pwd
/slot1/
```

In the following example, the management focus is a subdirectory called "test" on the flash card in slot 1.

```
device# pwd
/slot1/test/
```

Switching the management focus

The effect of file management commands depends on the file system that has the current management focus. For example, if you enter a command to delete a file and do not specify the location of the file, the software attempts to delete the file from the location that currently has the management focus.

By default, the management focus is on the flash memory on the management module. You can switch the focus from flash memory to flash cards in slot 1 or slot 2 on the management module using the **cd** or **chdir** commands, which have the same syntax and function exactly the same.

For example, to switch the focus from flash memory to the flash card in slot 2, enter the following command.

```
device# cd /slot2
device#
```

When you enter this command, the software changes the management focus to slot 2 then displays a new command prompt. If a slot you specify does not contain a flash card, the software displays the message shown in the following example.

```
device# cd /slot2
Device not present
```

Syntax: **cd** *directory-pathname*

Syntax: **chdir** *directory-pathname*

For the *directory-pathname* parameter for both **cd** and **chdir** commands, specify */slot1* or */slot2* to switch the focus to slot 1 or slot 2, respectively. Specify */flash* to switch the focus to flash memory.

After you have switched the focus to slot 2, you can specify the *directory-pathname* parameter to switch the focus to a subdirectory on a flash card inserted in slot 2. For example, to switch the focus from the root directory level (/) of slot 2 to the subdirectory named "PLOOK," enter the following command.

```
device# cd /PLOOK
```

If you specify an invalid subdirectory path, the CLI displays a message such as the following.

```
device# cd /PLOOK
Path not found
```

If you are certain the path you specified exists, make sure you are at the correct level to reach the path. For example, if you are already at the PLOOK level, the CLI cannot find the subdirectory *"/PLOOK"* because it is not a subdirectory from the level that currently has the management focus.

To change the management focus back to flash memory, enter the following command.

```
device# cd /flash
device#
```

Displaying a directory of the files

You can display a directory of the files in the flash memory on the management module, or on a flash card inserted in management module slot 1 or slot 2 using the **dir** or **ls** commands.

The software displays the directory of the file system that has the current management focus. By default, flash memory has the management focus. However, you do not need to change the focus to list the files on the file system that does not currently have management focus. In this case, you can specify the */path-name/* parameter with the **dir** or **ls** commands to display the directory of the desired file system.

For example, to display a directory of the files in flash memory, if flash memory has the management focus, enter the following command.

```
device# dir
Directory of /flash/
07/28/2003 15:57:45          3,077,697 1060.tmp
07/28/2003 15:56:10          3,077,697 14082.tmp
07/28/2003 16:00:08          3,077,697 2084.tmp
07/25/2003 18:00:23          292,701 boot
00/00/00 00:00:00           12 boot.ini
07/28/2003 14:40:19          840,007 lp-primary-0
07/28/2003 15:18:18          840,007 lp-secondary-0
07/28/2003 09:56:16          391,524 monitor
07/28/2003 15:08:12          3,077,697 primary
07/28/2003 16:02:23           1,757 startup-config
07/25/2003 18:02:14           1,178 startup.sj2
07/28/2003 14:28:47           1,662 startup.spa
07/26/2003 12:16:29           1,141 startup.vso
07/25/2003 18:11:01           1,008 startup.vsr
07/28/2003 09:40:54           1,554 startup.vsrp.ospf
          15 File(s)          14,683,339 bytes
           0 Dir(s)          15,990,784 bytes free
```

Syntax: `dir ls` [*path-name*]

You can enter either **dir** or **ls** for the command name.

Specify the *path-name* parameter to display the following:

- The files that match the value for a flash memory directory, or flash card directory/subdirectory you specify
- The files that match the value for a name you specify

For example, to list only files that contain a *.tmp suffix in flash memory, if flash memory is the current management focus, enter a command such as the following.

```
device# dir *.tmp
Directory of /flash/
07/28/2003 15:57:45          3,077,697 1060.tmp
07/28/2003 15:56:10          3,077,697 14082.tmp
07/28/2003 16:00:08          3,077,697 2084.tmp
          3 File(s)          9,292,701 bytes
           0 Dir(s)          15,990,784 bytes free
```

For example, to display a directory of the files on the flash card in slot 2, if flash memory has the management focus, enter the following command.

```
device# dir /slot2/
Directory of /slot2/
/
08/01/2003 18:25:28          3,092,508 PRIMARY
08/01/2003 18:28:06          3,092,508 primary.1234
08/01/2003 18:28:24          389,696 MONITOR
08/01/2003 18:28:30          389,696 MONITOR1
08/01/2003 18:28:01          389,696 MONITOR2
08/01/2003 18:28:03          389,696 MONITOR3
08/01/2003 18:29:04          389,696 MONITOR4
08/01/2003 18:29:12    <DIR>          DIR1
08/01/2003 18:32:03          389,696 1234567890.12345
08/01/2003 18:32:08          389,696 123456.123
08/01/2003 18:32:11          389,696 123456.123
08/01/2003 18:32:14          389,696 123456.123
08/01/2003 18:32:17          389,696 123456.123
          12 File(s)          10,081,976 bytes
           1 Dir(s)          114,577,408 bytes free
```

The following information is displayed for each file.

TABLE 44 CLI display of directory information

Field	Description
File date	The date on which the file was placed in the flash memory or card, if the device system clock is set.
Time of day	The time of day at which the file was placed in the flash memory or card, if the device system clock is set.
File size	The number of bytes in the file.
Read-write attribute	If you have set the read-write attribute of the file to read-only, "R" appears before the file name. If the read-write attribute of the file is read-write (the default), no value appears in this column. For information, refer to Changing the read-write attribute of a file on page 178.
File name	The file name.
Long file name	This field applies to files on a flash card only. The longer file name applies if the file was created on a PC and the name is longer than the 8.3 format.

The directory also lists the total number of files that match the parameters you specified, the total number of bytes used by all the files, and the number of bytes still free.

Displaying the contents of a file

You can display the contents of a file in the flash memory on the management module or on a flash card inserted in management module slot 1 or slot 2.

The software displays the specified file in the file system that has the current management focus (flash memory by default). However, you do not need to change the focus to display the file in a file system that does not currently have management focus. In this case, you can specify the */directory/ path-name* parameter with the **more** command to display the file in the desired file system.

For example, to display the contents of a file in flash memory, if flash memory has the current management focus, enter a command such as the following.

```
device# more cfg.cfg
```

Syntax: **more** [*directory*] *file-name*

Use the *directory* parameter to specify a directory in a file system that does not have current management focus.

Use the *file-name* parameter to specify the file you want to display.

For example, to display the contents of a file on the flash card in slot 2, if flash memory has the current management focus, enter a command such as the following.

```
device# more /slot2/cfg.cfg
```

Displaying the hexadecimal output of a file

You can display the hexadecimal output of a file in flash memory on the management module or on a flash card inserted in management module slot 1 or slot 2.

The software displays the hexadecimal output of a specified file in the file system that has the current management focus (flash memory by default). However, you do not need to change the focus to display the hexadecimal output of a file in a file system that does not currently have management focus. In this case, you can specify the */directory/file-name* parameter with the **hd** command to display the output of the file in the desired file system.

For example, to display the hexadecimal output of a file in flash memory, if flash memory has the current management focus, enter the following command.

```
device# hd cfg.cfg
```

Syntax: [no] **hd** [/directory/] **file-name**

Use the *directory* parameter to specify a directory in a file system that does not have current management focus.

Use the *file-name* parameter to specify a file for which you want to display the hexadecimal output.

For example, to display the hexadecimal output of a file in a flash card inserted in slot 2, if flash memory has the current management focus, enter the following command.

```
device# hd /slot2/cfg.cfg
```

Creating a subdirectory

Create a subdirectory in the flash card file system using the **md** and **mkdir** commands, which have the same syntax and function exactly the same.

NOTE

You cannot create subdirectories in the flash memory file system. Therefore, the **md** and **mkdir** commands do not apply to the flash memory file system.

The software creates a subdirectory in the file system that has the current management focus (flash memory by default). However, you do not need to change the focus to create a subdirectory in a file system that does not currently have management focus. In this case, you can specify the **slot1** or **slot2** keyword with the **md** or **mkdir** command to create the subdirectory in the desired file system.

For example, to create a subdirectory on the flash card inserted in slot 2, if the flash memory has current management focus, enter a command such as the following.

```
device# mkdir slot2 TEST
```

Syntax: [no] **md** | **mkdir** [slot1 | slot2] **dir-name**

You can enter either **md** or **mkdir** for the command name.

Specify the **slot1** or **slot2** keyword to create a subdirectory on the flash card in slot 1 or slot 2, respectively. If you do not specify one of these parameters, the command applies to the file system that currently has the management focus.

The *dir-name* parameter specifies the subdirectory name. You can enter a name that contains any combination of the following characters. Do not enter a slash " / " in front of the name. Remember, a file name preceded by a slash represents the absolute path name (/flash, /slot1, or /slot2).

- All upper and lowercase letters
- All digits
- Spaces
- Any of the following special characters:
 - \$
 - %
 - '
 - -
 - _
 - @
 - ~


```

- `
- !
- (
- )
- {
- }
- ^
- #
- &

```

You can use spaces in a subdirectory name if you enclose the name in double quotes. For example, to specify a subdirectory name that contains spaces, enter a string such as the following: "a long subdirectory name".

A subdirectory name can be a maximum of 256 characters long. A complete subdirectory path name cannot contain more than 260 characters.

The name is not case sensitive. You can enter upper- or lowercase letters, however the CLI displays the name using uppercase letters.

To verify successful creation of the subdirectory, enter a command such as the following to change to the new subdirectory level.

```

device# chdir /slot2/TEST
Current directory of slot2 is: /TEST

```

For information about changing the directory using the **cd** and **chdir** commands, refer to [Switching the management focus](#) on page 173.

Removing a subdirectory

You can remove a subdirectory from the flash card file system using the **rd** and **rmdir** commands, which have the same syntax and function exactly the same.

NOTE

You cannot remove subdirectories from the flash memory file system. Therefore, the **rd** and **rmdir** commands do not apply to the flash memory file system.

NOTE

You can remove a subdirectory only if the subdirectory does not contain files or other subdirectories.

The software will remove a subdirectory from the file system that has the current management focus (flash memory by default). However, you do not need to change the focus to remove a subdirectory from a file system that does not currently have management focus. In this case, you can specify the **slot1** or **slot2** keyword with the **rd** or **rmdir** command to remove the subdirectory from the desired file system.

For example, to remove a subdirectory from the flash card inserted in slot 2, if the flash memory has current management focus, enter a command such as the following.

```

device# rmdir slot2 TEST

```

Syntax: **[no] rd rmdir** | **[slot1 | slot2] dir-name**

You can enter either **rd** or **rmdir** for the command name.

Specify the **slot1** or **slot2** keyword to remove a subdirectory on the flash card in slot 1 or slot 2, respectively. If you do not specify one of these parameters, the command applies to the file system that currently has the management focus.

The *dir-name* parameter specifies the subdirectory you want to delete. You can enter a path name if the subdirectory is not in the current directory.

If you receive a message such as the following, enter the **pwd** command to verify that the management focus is at the appropriate level of the directory tree.

```
device# rmdir TEST
rmdir /slot1/test/dir1/temp failed - File not found
```

For information about using the **pwd** command, refer to [Determining the current management focus](#) on page 172.

Renaming a file

You can rename a file in the flash memory on the management module or on a flash card inserted in management module slot 1 or slot 2 using the **rename** or **mv** command.

The software renames the file in the file system that has the current management focus flash memory by default. However, you do not need to change the focus to rename the file in a file system that does not currently have management focus. In this case, you can specify the */directory/old-file-name /directory/new-file-name* parameter with the **rename** or **mv** command to rename the file in the desired file system.

For example, to rename a file in flash memory, if flash memory has the current management focus, enter a command such as the following.

```
device# rename oldname newname
```

If the command is successful, the CLI displays a new command prompt.

Syntax: **[no] rename mv** | **[/directory/] old-file-name [/directory/] new-file-name**

You can enter either **rename** or **mv** for the command name.

The */directory/* parameter specifies a directory in a file system that does not have current management focus. When moving a file, the path must remain at the same directory level. You cannot rename the directory and the directory can be nested a maximum of five levels.

NOTE

Moving files up to the root directory is not supported.

The *old-file-name* parameter specifies the original filename that you want to change.

The *new-file-name* parameter specifies the new filename that you want to assign to the original file. The new filename must have an equal or greater number of characters than the old filename. The new filename cannot exceed 32 characters.

NOTE

A new filename with fewer characters than the old filename is not supported.

For example, to rename a file on the flash card inserted in slot 2, if flash memory has the current management focus, enter a command similar to the following.

```
device# rename /slot2/oldname /slot2/newname
```

Changing the read-write attribute of a file

You can specify the read-write attribute of a file on a flash card as follows:

- **Read-only** - You can display or copy the file but you cannot replace (copy over) or delete the file.
- **Read-write** - You can replace (copy over) or delete the file. This is the default.

NOTE

All files in flash memory are set to the read-write attribute, which cannot be changed. You cannot change this attribute. Therefore, the **attrib** command does not apply to the flash memory file system.

To determine the current setting of the read-write attribute for a file, use the **dir** command to list the directory information for the file. Files set to read-only are listed with "R" in front of the file name. For information about the **dir** command, refer to [Displaying a directory of the files](#) on page 173.

The software will change the read-write attribute of the file in the file system that has the current management focus (flash memory by default). However, you do not need to change the focus to change this file attribute in a file system that does not currently have management focus. In this case, you can specify the **slot1** or **slot2** keyword with the **attrib** command to change the attribute of the file in the desired file system.

For example, to change the attribute of a file in slot2 to read-only, if flash memory has the management focus, enter a command similar to the following.

```
device# attrib slot2 ro goodcfg.cfg
```

Syntax: [no] attrib [slot1 | slot2] ro | rw file-name

Specify the **slot1** or **slot2** keyword to change the attribute of a file on the flash card in slot 1 or slot 2, respectively. If you do not specify one of these keywords, the command applies to the file system that currently has the management focus.

The **ro** parameter specifies that the attribute of the file is set to read-only. The **rw** parameter specifies that the attribute of the file is set to read-write.

The *file-name* parameter specifies the file for which to change the attribute.

For example, to change the attribute of all files on the flash card in slot 2 to read-only, if flash memory has the current management focus, enter a command similar to the following.

```
device# attrib slot2 ro *.*
```

Deleting a file

You can delete a file from flash memory or a flash card inserted in slot 1 or slot 2 on the management module using the **delete** or **rm** command.

NOTE

The **delete** or **rm** command deletes all files in a file system unless you explicitly specify the files you want to delete.

NOTE

The software does not support an undelete option for the flash memory file system. Be sure you really want to delete the file before you issue this command.

The software will delete the file in the file system that has the current management focus. By default, flash memory has the management focus. However, you do not need to change the focus to delete the file in a file system that does not currently have management focus. In this case, you can specify the */directory/file-name* parameter with the **delete** or **rm** command to delete the file in the desired file system.

For example, to delete a file in flash memory, if flash memory has the current management focus, enter a command similar to the following.

```
device# delete cfg.cfg
```

If the command is successful, the CLI displays a new command prompt.

Syntax: `delete rm [slot1 | slot2] [directory] [file-name]`

You can enter either **delete** or **rm** for the command name.

Specify the **slot1** or **slot2** keywords to delete all files on the flash card in slot 1 or slot 2, respectively.

The *directory* parameter specifies the directory in a file system that does not have the current management focus.

The *file-name* parameter specifies the file that you want to delete.

For example, to delete all files with names that start with "test" from flash memory, if flash memory has the current management focus, enter a command similar to the following.

```
device# delete test*.*
```

For example, to delete all files on the flash card in slot 2, if flash memory has the current management focus, you can enter one of the following commands.

```
device# delete /slot2/
```

or

```
device# delete slot2
```

Recovering ("undeleting") a file

You can recover or undelete a file you have deleted from a flash card file system using the **undelete** command.

NOTE

You can not recover or undelete a file from the flash memory file system. Therefore, the **undelete** command does not apply to the flash memory file system.

The software will recover the file in the file system that has the current management focus (flash memory by default). If you want to recover a file in a file system that does not have the current management focus, you must switch the management focus to the desired file system using the **cd** command. For more information about switching the management focus, refer to [Switching the management focus](#) on page 173.

For example, to undelete a file on the flash card in slot 2, if flash memory has the current management focus, enter a command such as the following.

```
device# cd slot2
device# undelete
Undelete file ?RIMARY ? (enter y or n) :y
Input one character: P
File recovered successfully and named to PRIMARY
```

For each file that can be undeleted from the flash card in slot 2, the CLI displays the remaining name entry in the file directory and prompts you for the first character of the file name. You can enter any valid file name character. You do not need to enter the character that was used before in the deleted file name.

Once you enter a character and the CLI undeletes the file, the CLI continues with the next file that can be undeleted. For each file, specify "y" or "n", and specify a first character for the files that you select to undelete.

NOTE

When you delete a file from a flash card, the CLI leaves the file intact but removes the first letter in the file name from the file directory. However, if you save file changes or new files that use part of the space occupied by the deleted file, you cannot undelete the file. The **undelete** command lists only the files that can be undeleted.

To end the undelete process, enter CTRL + C.

Appending a file to another file

You can append a file in flash memory or on a flash card to the end of another file in one of these file systems.

The software will append one file to another in the file system that has the current management focus (flash memory by default). However, you do not need to change the focus to append one file to another in a file system that does not currently have management focus. In this case, you can specify the */source-dir-path/* or */dest-dir-path/* parameters with the **append** command to append one file to another in the desired file system.

To append one file to another in flash memory, if flash memory has the current management focus, enter a command similar to the following.

```
device# append newac1s.cfg startup-config.cfg
```

Syntax: **[no] append** [**source-file-system** **dest-file-system**] [*/source-dir-path/*] **source-file-name** [*/dest-dir-path/*] **dest-file-name**

Specify the *source-file-system* and *dest-file-system* parameters when you are appending a file on one file system to a file on another file system.

The */source-dir-path/* *source-file-name* parameter specifies the file you are appending to the end of another file. If the file is not located in the current subdirectory (the subdirectory that currently has the management focus), specify the subdirectory path in front of the file name.

The */dest-dir-path/* *dest-file-name* parameter specifies the file to which you are appending the other file. If the file is not located in the current subdirectory, specify the subdirectory path in front of the file name.

For example, to append a file in the root directory of slot 1 to another file in a subdirectory of slot 2, enter a command similar to the following.

```
device# append slot1 slot2 newac1s.cfg /TEST/startup-config.cfg
```

Copying files using the copy command

For information about copying files using the **copy** command while upgrading software images, refer to "Basic Tasks in the Software Upgrade Process" in the appropriate hardware installation guide.

You can perform the following additional copy operations using the **copy** command:

- Copy files from one flash card to the other
- Copy files between a flash card and the flash memory on the management module
- Copy software images between active and standby management modules
- Copy files from a management module to an interface module
- Copy Extreme Multi-Service IronWare management module images from flash memory to a TFTP server
- Copy files between a flash card and a TFTP server
- Copy a startup-config file between a flash card and flash memory on the management module
- Copy a startup-config file between flash memory on the management module and a TFTP server
- Copy the running-config to a flash card or a TFTP server
- Load a running-config from a flash card or TFTP server into the running-config on the device

NOTE

Since the copy options require you to explicitly specify the flash card, you can perform a copy regardless of which flash card is on the currently active management module.

Copying files from one flash card to the other

To copy a file from one flash card to the other, enter the following command.

```
device# copy slot1 slot2 sales.cfg
```

Syntax: `copy from-card to-card [/from-dir-path/] from-name [/to-dir-path/] [to-name]`

For the *from-card* and *to-card* parameters, you can specify **slot1** or **slot2**.

The command shown in the example copies a file from the flash card in slot 1 to the flash card in slot 2. In this case, the software uses the same name for the original file and for the copy. Optionally, you can specify a different file name for the copy.

Copying files between a flash card and flash memory

To copy a file from a flash card to the primary area in flash memory, enter a command similar to the following.

```
device# copy slot1 flash  
nmpr02200.bin primary
```

Syntax: `copy slot1 | slot2 flash [/from-dir-path/] from-name monitor | primary | secondary`

To copy a file from flash memory to a flash card, enter a command similar to the following.

```
device# copy flash slot2  
nmpr02200.bin primary
```

Syntax: `copy flash slot1 | slot2 source-name monitor | primary | secondary | startup-config [dest-name]`

The command in this example copies a Extreme Multi-Service IronWare image file from the primary area in flash memory onto the flash card in slot 2. In this case, the software uses the same name for the source file and for the destination file. Optionally, you can specify a different file name for the destination file.

Copying software images between active and standby management modules

To copy the monitor image from flash memory of the active management module to flash memory of the standby module, enter the following command.

```
device# copy flash flash monitor standby
```

To copy the Extreme Multi-Service IronWare image from the secondary location in flash memory on the active management module to the primary location in flash memory, enter the following command.

```
device # copy flash flash primary
```

Syntax: `copy flash flash primary [standby]`

Specify the optional **standby** keyword to copy the Extreme Multi-Service IronWare image from the secondary location in flash memory on the active management module to the primary location in flash memory on the standby module.

To copy the Extreme Multi-Service IronWare image from the primary location in flash memory on the active management module to the secondary location in flash memory on the active module, enter the following command.

```
device# copy flash flash secondary
```

Syntax: `copy flash flash secondary [standby]`

Specify the optional **standby** keyword to copy the Extreme Multi-Service IronWare image from the primary location in the flash memory on the active management module to the secondary location in the flash memory on the standby module.

Copying Extreme Multi-Service IronWare images from flash memory to a TFTP Server

You can copy Extreme Multi-Service IronWare images from the primary and secondary locations in flash memory on the management module to a TFTP server.

For example, to copy the Extreme Multi-Service IronWare image in the secondary location in flash memory to a TFTP server, enter a command similar to the following.

```
device# copy flash tftp
10.10.10.1 secondary.bak secondary
```

Syntax: `copy flash tftp ip-addr dest-file-name primary | secondary`

Copying files between a flash card and a TFTP server

Use the following methods to copy files between a flash card and a TFTP server.

NOTE

The Extreme system must have network access to the TFTP server.

To copy a file from a flash card to a TFTP server, enter a command similar to the following.

```
device# copy slot1 tftp 192.168.1.17 notes.txt
```

Syntax: `copy slot1 | slot2 tftp ip-addr [/from-dir-path/] source-file [dest-file]`

The command in this example copies a file from slot 1 to a TFTP server. In this case, the software uses the same name for the source file and for the destination file. Optionally, you can specify a different file name for the destination file.

To copy a software image from a TFTP server to a flash card, enter a command similar to the following.

```
device# copy tftp slot1 192.168.1.17
nmp02200.bin primary
```

Syntax: `copy tftp slot1 | slot2 ip-addr [/from-dir-path/] source-file path-name | monitor | primary | secondary`

The command in this example copies the primary Extreme Multi-Service IronWare image from a TFTP server to a flash card in slot 1.

Copying the startup-config file between a flash card and flash memory

Use the following methods to copy a startup-config file between flash memory and a flash card. By default, the Extreme device uses the startup-config in the primary area of flash memory when you boot or reload the device.

NOTE

The Extreme device cannot configure from a startup-config file on a flash card. You cannot boot or reload from a flash card.

To copy a startup-config file from a flash card to flash memory, enter a command similar to the following.

```
device# copy slot1 startup-config test2.cfg
```

Syntax: `copy slot1 | slot2 startup-config [from-dir-path/file-name]`

This command copies a startup configuration named test2.cfg from the flash card in slot 1 into the flash memory on the device. The next time you reboot or reload, the device uses the configuration information in test2.cfg.

To copy the startup-config file on the device from flash memory onto a flash card, enter a command similar to the following.

```
device# copy startup-config slot1 mfgtest.cfg
```

Syntax: `copy startup-config slot1 | slot2 [/to-dir-path/] to-name`

This command copies the startup configuration from the flash memory on the device to a flash card in slot 1 and names the file `mfgtest.cfg`.

Copying the startup-config file between flash memory and a TFTP server

Use the following methods to copy a startup-config between flash memory and a TFTP server to which the Extreme system has access. By default, the device configures from the startup-config in the primary area of flash memory when you boot or reload the device.

To copy the startup-config on the device from flash memory to a TFTP server, enter a command similar to the following.

```
device# copy startup-config tftp 10.10.10.1 /backups/startup.cfg
```

Syntax: `copy startup-config tftp ip-addr [/to-dir-path] to-name`

To copy a startup-config file from a TFTP server to flash memory, enter a command similar to the following.

```
device# copy tftp startup-config 10.10.10.1 test.cfg
```

Syntax: `copy tftp startup-config ip-addr [/from-dir-path] from-name`

Copying the running-config to a flash card or a TFTP server

Use the following method to copy the config file on the Extreme device to a flash card or a TFTP server. The running-config contains currently active configuration information for the device. When you copy the running-config to a flash card or TFTP server, you are making a copy of the current configuration, including any configuration changes you have not saved to the startup-config.

To copy the running configuration for the device into a file on a flash card, enter a command similar to the following.

```
device# copy running-config slot1 runip.1
```

Syntax: `copy running-config slot1 | slot2 [/to-dir-path/] to-name`

To copy the running configuration for the device into a file on a TFTP server, enter a command such as the following.

```
device# copy running-config tftp 10.10.10.1 runip.1
```

Loading a running-config from a flash card or a TFTP server

Use the following method to load configuration commands into the active configuration for the Extreme device.

NOTE

A configuration file that you create must follow the same syntax rules as the startup-config the device creates. Refer to "Dynamic Configuration Loading" in the *appropriate hardware installation guide*.

To copy a running-config from a flash card, enter a command such as the following.

```
device# copy slot2 running-config runacl.2
```

Syntax: `copy slot1 | slot2 running-config [/from-dir-path/] from-name`

The command in this example changes the active configuration for the device based on the information in the file.

To copy a running-config from a TFTP server, enter a command similar to the following.

```
device# copy tftp running-config 10.10.10.1 run.cfg overwrite
```

Syntax: `copy tftp running-config ip-addr [/from-dir-path/] from-name [overwrite]`

This command copies a running-config from a TFTP server and overwrites the active configuration for the device.

NOTE

You cannot use the overwrite option from non-console sessions, as it will disconnect the session.

When a configuration file is loaded using the `copy tftp running-config` command, the following commands within the configuration file are supported.

- **isis metric command**
- **set-overload-bit command**
- **admin-group**
- **cspf-group**
- **bypass-lsp**

Copying files using the cp command

Use the **cp** command to do the following:

- Copy files from flash memory to flash memory
- Copy files from flash memory to a flash card or vice versa
- Copy files from one flash card to another flash card

The software will copy a file in a file system to another location in the file system that has the current management focus (flash memory by default). However, you do not need to change the focus to copy a file from one location to another in a file system that does not currently have management focus. In this case, you can specify the */source-dir-path/* or */dest-dir-path/* parameters with the **cp** command to copy a file to or from a file system that does not have current management focus.

For example, to copy a file from flash memory, which has the current management focus, to flash memory, enter a command similar to the following.

```
device# cp primary primary2
```

For example, to copy a file from flash memory, which has the current management focus, to the flash card in slot 2, enter a command similar to the following.

```
device# cp new.cfg /slot2
/cfg/new.cfg
```

Syntax: **cp** [*source-dir-path*] *source-file-name* [*dest-dir-path*] *dest-file-name*

The *source-dir-path* parameter specifies the directory pathname of the source file. Specify this parameter if the source file is in a file system that does not have current management focus. The *source-file-name* specifies the name of the file you want to copy.

The *dest-dir-path* parameter specifies the directory pathname of the destination file. Specify this parameter if you want to copy the source file to a file system that does not have current management focus. The *dest-file-name* specifies the name of the file you copied to a new destination.

For example, to copy a file from a flash card in slot 2 to flash memory, which has current management focus, enter the following command.

```
device# cp /slot2
/cfg/new.cfg new.cfg
```

For example, to copy a file from a flash card in slot 1 to a flash card in slot 2, neither of which has current management focus, enter the following command.

```
device# cp /slot1/cfg/new.cfg /slot2
/cfg/new.cfg
```

Loading the software

By default, the management module loads an Extreme Multi-Service IronWare image from the primary location in flash memory. You can change the Extreme Multi-Service IronWare image source for the system to one of the following sources for a single reboot or for all future reboots:

- The secondary location in flash memory
- A flash card inserted in slot 1 or 2
- A TFTP server
- A BOOTP server

If you specify a source other than the primary location in flash memory and for some reason the source or the Extreme Multi-Service IronWare image is unavailable, the system uses the primary location in flash memory as a default backup source.

Rebooting from the system

To use a source besides the Multi-Service IronWare image in the primary location in flash memory for a single reboot, enter a command similar to the following at the Privileged EXEC level of the CLI.

```
device# boot system slot1 /slot1/xmr03000.bin
```

The command in this example reboots the system using the image xmr03000.bin located on the flash card in slot 1. This example assumes that the flash card in slot 1 is not the management focus.

Syntax: `boot system slot1 | slot2 [/dir-path/] file-name`

The **slot1** and **slot2** keywords specify the flash card slot.

The *file-name* parameter specifies the file name. If the file is in a subdirectory, specify the subdirectory path in front of the file name. If the file name you specify is not a full path name, the CLI assumes that the name (and path, if applicable) you enter are relative to the subdirectory that currently has the management focus.

NOTE

This command also is supported at the boot PROM.

For example, to reboot the system using the image xmr03000.bin on a TFTP server, enter a command similar to the following.

```
device# boot system tftp 10.10.10.1 xmr03000.bin
```

Syntax: `boot system tftp ip-address file-name`

The *ip-address* parameter specifies the address of the TFTP server on which the desired image resides.

The *file-name* parameter specifies the name of the Extreme Multi-Service IronWare image on the TFTP server.

For example, to reboot the system using the secondary location in flash memory, enter the following command.

```
device# boot system flash secondary
device# Are you sure? (enter 'y' or 'n'): y
```

Syntax: `boot system flash secondary`

To reboot the system from a BOOTP server, enter the following command.

```
device# boot system bootp
```

Syntax: `boot system bootp`

Configuring the boot source for future reboots

To change the Extreme Multi-Service IronWare image source from the primary location in flash memory to another source for future reboots, enter a command similar to the following at the global CONFIG level of the CLI.

```
device(config)# boot system slot1 xmr03000.bin
```

The command in this example sets Auxiliary flash slot 1 as the primary boot source for the Extreme device. When you reload the software or power cycle the device, the device will look for the Extreme Multi-Service IronWare image on the flash card in slot 1.

Syntax: `boot system slot1 file-name | slot2 file-name | flash secondary | tftp ip-address file-name | bootp`

NOTE

The command syntax is the same for immediately reloading and for changing the primary source, except the *file-name* must be the full path name. You cannot specify a relative path name. If the first character in the path name is not a slash (/), the CLI treats the name you specify as relative to the root directory. How the device responds to the command depends on whether you enter the command at the Privileged EXEC level or the global CONFIG level.

If you enter multiple **boot system** commands at the global CONFIG level, the software places them in the running-config in the order you enter them, and saves them to the startup-config in the same order when you save the configuration. When you reload or power cycle the device, the device tries the boot sources in the order they appear in the startup-config and running-config.

Saving configuration changes

You can configure the Extreme system to save configuration changes to a startup-config in flash memory or on a flash card in slot 1 or 2.

Displaying the current location for saving configuration changes

Enter the following command at the Privileged EXEC level of the CLI to display the current save location for the startup-config.

```
device# locate startup-config
Startup-config data location is flash memory
```

Specifying the location for saving configuration changes

By default, when you save configuration changes, the changes are saved to the startup-config in flash memory. To change the save location to a flash card in slot 1 or 2, enter a command similar to the following.

```
device# locate startup-config slot1 router1.cfg
device# write memory
```

The first command in this example sets the device to save configuration changes to the file named "switch1.cfg" in the flash card in slot 1. The second command saves the running-config to the router1.cfg file on the flash card in slot 1.

NOTE

In this example, after you save the configuration changes using the **write memory** command, the router1.cfg file will include the command that designates slot 1 as the save location for configuration changes.

Syntax: `locate startup-config [slot1 | slot2 | flash-memory] [/dir-path-name/] file-name`

The **locate** command is used only for saving the startup-config file to a different location. But once after reload, the system always picks up the startup-config file from the flash memory.

The **slot1**, **slot2**, and **flash-memory** keywords specify the flash card in slot 1 or slot 2 or flash memory as the save location for configuration changes.

Specify the *dir-path-name* parameter if you want to save the configuration changes to a directory other than the root directory of a flash card file system.

The *file-name* parameter indicates the name of the saved configuration file.

To change the save location back to flash memory, enter a command similar to the following.

```
device# locate startup-config flash-memory router1.cfg
device# write memory
```

File management messages

The following table lists the messages the CLI can display in response to file management commands.

TABLE 45 Flash card file management messages

This message...	Means...
File not found	You specified a file name that the software could not find. Verify the command you entered to make sure it matches the source and destination you intended for the file operation.
Current directory is: <i>dir-path</i>	You have successfully changed the management focus to the slot and subdirectory indicated by the message.
Path not found	You specified an invalid path.
There is not enough space on the card	The flash card does not have enough space to hold the file you are trying to copy to it.
Access is denied	You tried to copy or delete a file that has the read-only attribute.
A duplicate file name exists	You tried to rename a file using a name that is already in use by another file.
Fatal error, can not read or write media	A hardware error has occurred. One possible cause of this message is removing the flash card while a file operation involving the card was in progress.
There is sharing conflict between format command and other read/write operations	The flash card is currently undergoing formatting. This message also appears if you enter a command to format the card while the card is being accessed for another file operation.
Invalid DOS file name	A filename you entered contains an invalid character (for example, ":" or "\").
File recovered successfully and named <i>file-name</i>	A file you tried to recover was successfully recovered under the name indicated in the message

LLDP


- Link layer discovery protocol overview..... 189
- General operating principles..... 190
- Configuration considerations..... 194
- Using LLDP..... 195
- Resetting LLDP statistics..... 209

Link layer discovery protocol overview

Link layer discovery protocol (LLDP) enables a station attached to an IEEE 802 LAN or MAN to advertise its capabilities and to discover other stations in the same 802 LAN segments. The advertisements describe the network's physical topology and associated systems within that topology. For example, a station can advertise its management address, the address of the entities that manage the device, and the ID of the port to which the station is connected.

The information distributed through LLDP (the advertisement) is stored by the receiving device in a standard Management Information Base (MIB), accessible using a management protocol such as the Simple Network Management Protocol (SNMP). The information also can be viewed through the CLI, using **show LLDP** commands.

FIGURE 5 LLDP connectivity



The LLDP has been enhanced to allow configuring the sub-type for LLDP port-id as one of the below which can be advertised along with the corresponding port-id string to the receiving device which displays this information as port-id string and subtype info of its neighbor.

- Interface-name
- Interface-alias
- Mac-address

For example, see the below configuration:

```
device(config)#lldp advertise
link-aggregation      Advertise 802.3 link aggregation information
mac-phy-config-status Advertise 802.3 MAC/PHY configuration/status
management-address    Advertise a management address
max-frame-size         Advertise 802.3 maximum frame size
port-description       Advertise port description
port-id-subtype        Advertise 802.1 port-id-subtype
port-protocol-vlan-id  Advertise no 802.1 port-and-protocol VLAN support
port-vlan-id           Advertise 802.1 port untagged VLAN id (PVID)
system-capabilities    Advertise system capabilities
system-description     Advertise system description
system-name            Advertise system name
vlan-name              Advertise 802.1 VLAN name

device(config)#lldp advertise port-id-subtype ?
interface-alias  Advertise 802.1 interface alias
interface-name   Advertise 802.1 interface name
mac-address      Advertise 802.1 interface mac-address
```

`interface-alias` sub-option when configured as above enables advertising the location of a port as Port ID subtype achieved by configuring unique names indicating its location to the interface as port ID subtype per RFC2863, using the global config command **port-name**.

`interface-name` sub-option when configured as above enables advertising the device local interface name. Such as the interface name used at the device local console as Port ID subtype as per RFC2863.

`mac-address` sub-option when configured as above enables advertising the mac-address of the interface as Port ID subtype.

NOTE

By default, when a user does not specify which Port ID subtype has to be advertised using the relevant configuration command, mac-address of the port is advertised to the remote device in the port ID TLV in LLDPDU. Hence, the default configuration of port ID subtype as 'mac-address' is not shown in **show running config** output. Extreme devices use port ID subtype 3, the default MAC address associated with the port (the same port ID sub-type is used to advertise other port ID sub-types to the remote devices such as interface-alias or interface-name or mac-address as per the user configuration).

This information can be retrieved on remote device as advertised from local device using either LLDP show commands or SNMP GET operations.

General operating principles

LLDP use the services of the Data Link sub layers, Logical Link Control and Media Access Control, to transmit and receive information to and from other LLDP Agents (protocol entities that implement LLDP).

LLDP is a one-way protocol. An LLDP agent can transmit and receive information to and from another LLDP agent located on an adjacent device, but it cannot solicit information from another LLDP agent, nor can it acknowledge information received from another LLDP agent.

Operating modes

When LLDP is enabled on a global basis, by default, each port on the Extreme device will be capable of transmitting and receiving LLDP packets. LLDP supports the following operating modes on physical interfaces:

- Transmit and Receive LLDP information. (System default)
- Transmit LLDP information only
- Receive LLDP information only

Transmit mode

An LLDP agent sends LLDP packets to adjacent LLDP-enabled devices. The LLDP packets contain information about the transmitting device and port.

An LLDP agent initiates the transmission of LLDP packets whenever the transmit countdown timing counter expires, or whenever LLDP information has changed. When a transmit cycle is initiated, the LLDP manager extracts the MIB objects and formats this information into TLVs. The TLVs are inserted into an LLDPDU, addressing parameters are prepended to the LLDP packet, and the information is sent out LLDP-enabled ports to adjacent LLDP-enabled devices.

Receive mode

An LLDP agent receives LLDP packets from adjacent LLDP-enabled devices. The LLDP packets contain information about the transmitting device and port.

When an LLDP agent receives LLDP packets, it checks to ensure that the LLDP packets contain the correct sequence of mandatory TLVs, then validates optional TLVs. If the LLDP agent detects any errors in the LLDPDUs and TLVs, it drops them in software. TLVs that are not recognized but do not contain basic formatting errors, are assumed to be valid and are assigned a temporary identification index and stored for future possible alter retrieval by network management. All validated TLVs are stored in the neighbor database.

LLDP packets

LLDP agents transmit information about a sending device or port in packets called LLDP Data Units (LLDPDUs). All the LLDP information to be communicated by a device is contained within a single 1500 byte packet. LLDP information exceeding 1500 bytes will be truncated. A device receiving LLDP packets is not permitted to combine information from multiple packets.

Each LLDPDU consists of an untagged Ethernet header and a sequence of short, variable length information elements known as TLVs.

TLVs have Type, Length, and Value fields, where:

- **Type** identifies the kind of information being sent
- **Length** indicates the length (in octets) of the information string
- **Value** is the actual information being sent (for example, a binary bit map or an alpha-numeric string containing one or more fields).

TLV support

This section lists and describes LLDP TLV support.

LLDP TLVs

There are two types of LLDP TLVs, as specified in the IEEE 802.3AB standard:

- **Basic Management TLVs** consist of both optional general system information TLVs as well as mandatory TLVs.

Mandatory TLVs cannot be manually configured. They are always the first three TLVs in the LLDPDU, and are part of the packet header.

General system information TLVs are optional in LLDP implementations and are defined by the Network Administrator.

Extreme devices support the following Basic Management TLVs:

- Chassis ID (mandatory)
- Port ID (mandatory)
- Time to Live (mandatory)
- Port description
- System name
- System description
- System capabilities
- Management address
- End of LLDPDU

- **Organizationally-specific TLVs** are optional in LLDP implementations and are defined and encoded by individual organizations or vendors. These TLVs include support for, but are not limited to, the IEEE 802.1 and 802.3 standards and the TIA-1057 standard.

Extreme devices support the following Organizationally-specific TLVs:

- 802.1 organizationally-specific TLVs
 - > Port VLAN ID
 - > VLAN name TLV
- 802.3 organizationally-specific TLVs
 - > MAC/PHY configuration/status
 - > Link aggregation
 - > Maximum frame size

Mandatory TLVs

LLDP uses the service provided by the LLDP/LSAP and LLC to transmit and receive LLDPDUs. The LLDPDU shall contain an ordered sequence of three mandatory TLVs followed by one or more optional TLVs plus an end Of LLDPDU TLV.

- Three mandatory TLVs at the beginning of each LLDPDU are:
 - Chasis ID
 - Port ID
 - Time to Live
- Optional TLVs as selected by network management can be inserted in any order.
- The end Of LLDPDU TLV shall be the last TLV in the LLDPDU.

The mandatory and optional TLVs are encoded in the LLDPDU and sent to the remote device. On the receiving device, these TLVs are decoded and information is stored in the database per port pertaining to each neighbor. This information can be retrieved either using **show commands** or **SNMP Get** operations on the MIB.

Chassis ID

The Chassis ID identifies the device that sent the LLDP packets.

There are several ways in which a device may be identified. A Chassis ID subtype, included in the TLV and shown in the following table, indicates how the device is being referenced in the Chassis ID field.

TABLE 46 Chassis ID subtypes

ID Subtype	Description
0	Reserved
1	Chassis component
2	Interface alias
3	Port component
4	MAC address
5	Network address
6	Interface name
7	Locally assigned
8 - 255	Reserved

Extreme devices use Chassis ID subtype 4, the base MAC address of the device. Other third party devices may use a Chassis ID subtype other than 4. The Chassis ID will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
Chassis ID (MAC address): 0012.f233.e2c0
```

The Chassis ID TLV is always the first TLV in the LLDPDU.

Port ID

The Port ID identifies the port from which LLDP packets were sent.

There are several ways in which a port may be identified, as shown in the following table. A port ID subtype, included in the TLV, indicates how the port is being referenced in the Port ID field.

TABLE 47 Port ID subtypes

ID Subtype	ID basis	References
0	Reserved	-
1	Interface alias	ifAlias (IETF RFC 2863)
2	Port component	entPhysicalAlias when entPhysicalClass has a value port (10) or backplane (4) (IETF RFC 2737)
3	MAC address	MAC address (IEEE Std 802-2001)
4	Network address	network address
5	Interface name	ifName (IETF RFC 2863)
6	Agent circuit ID	agent circuit ID (IETF RFC 3046)
7	Locally assigned	local
8 - 255	Reserved	-

Other third party devices may use a port ID subtype other than default type. The port ID appears similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
Port ID (MAC address): 0012.f233.e2d3
```

To advertise interface-alias string to the remote device when Port-ID Subtype is configured using **lldp advertise port-id-subtype interface-alias** command:

```
device(config)#lldp advertise port-id-subtype interface-alias ports ethernet 1/1
device (config) #show lldp local-info ports ethernet 1/1
Local port: 1/1
+ Chassis ID (MAC address): 0024.3863.8cc0
+ Port ID (interface alias): "ExtrBld1Floor3Cube3300"
+ Time to live: 120 seconds
+ System name       : "CER-1"
+ Port description  : "GigabitEthernet1/1"
+ System capabilities : bridge, router
+ Enabled capabilities: bridge, router
+ 802.3 MAC/PHY      : auto-negotiation enabled
+ Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
                          100BaseTX-FD, 1000BaseT-HD, 1000BaseT-FD
+ Operational MAU type : 1000BaseT-FD
+ Link aggregation: not capable
+ Maximum frame size: 1548 octets
+ Port VLAN ID: 100
+ Management address (IPv4): 10.37.71.1
```

NOTE

The LLDP Port ID Subtype supports Enhanced 9-1-1 (E-911) by advertising information about the physical location of a port when configured using **port-name** global config command.

TTL value

The Time to Live (TTL) Value is the length of time the receiving device should maintain the information acquired through LLDP in its MIB.

The TTL value is automatically computed based on the LLDP configuration settings. The TTL value will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**):

```
Time to live: 40 seconds
```

- If the TTL field has a value other than zero, the receiving LLDP agent is notified to completely replace all information associated with the LLDP agent or port with the information in the received LLDPDU.
- If the TTL field value is zero, the receiving LLDP agent is notified that all system information associated with the LLDP agent or port is to be deleted. This TLV may be used, for example, to signal that the sending port has initiated a port shutdown procedure.

Configuration considerations

- LLDP is supported on Ethernet interfaces only.
- If a port is 802.1X-enabled, the transmission and reception of LLDP packets will only take place while the port is authorized.
- Cisco Discovery Protocol (CDP) and Foundry Discovery Protocol (FDP) run independently of LLDP; therefore, these discovery protocols can run simultaneously on the same device.
- LLDP is supported on VPLS/VLL end-points and the behavior is the same as other interfaces.
- LLDP packets have the standard Multicast Destination MAC address and are sent with highest priority (7).

- By default, the Extreme device limits the number of neighbors per port to four (valid range is 1- 64), and staggers the transmission of LLDP packets on different ports, in order to minimize any high-usage spikes to the CPU.
- If the advertisements by the neighbor exceed the maximum value of the neighbor per port or if it exceeds the maximum neighbors configured at the global level then the new advertisements will be dropped.
- LLDP advertisements are limited to a single 1500 byte packet.

Using LLDP

LLDP is disabled by default on individual ports. To run LLDP, it must be enabled on a global basis (on the entire device).

Enabling LLDP

To enable LLDP globally, enter the **lldp run** command at the global configuration level.

```
device(config)# lldp run
device(config)#lldp enable ports ethernet 1/1 to 1/2
device(config)#lldp advertise port-id-subtype interface-alias ports ethernet 1/1
device(config)#lldp advertise port-id-subtype interface-name ports ethernet 1/2

device(config)# show run | inc lldp
device(config)#lldp advertise port-id-subtype interface-alias ports ethe 1/1
device(config)#lldp advertise port-id-subtype interface-name ports ethe 1/2
device(config)#lldp enable ports ethe 1/1 to 1/24
device(config)#lldp run
```

Syntax: [no] **lldp run**

NOTE

The LLDP Port ID Subtype configured as mac-address is not shown in **show running config** output as it is default configuration.

Changing the operating mode of a port

When LLDP is enabled on a global basis, by default, each port on the Extreme device will be capable of transmitting and receiving LLDP packets. Each port can be configured for a different operating mode on the Extreme device.

Configuring transmit and receive mode

To enable receipt and transmission of LLDP packets on individual ports, enter the **lldp enable ports ethernet** command at the Global CONFIG level of the CLI. The enabled ports are placed into transmit and receive mode by default.

```
device(config)# lldp enable ports ethernet 2/1
```

Syntax: [no] **lldp enable ports ethernet** *portlist* | **all**

For *port list*, specify the ports in the format [*slotnum*/]*portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Use the **no** form of the command to disable the receipt and transmission of LLDP packets on a port.

Configuring transmit mode

To change the LLDP operating mode from receive and transmit mode to transmit only mode, disable the transmit and receive mode, and enter the **lldp enable transmit ports ethernet** command.

```
device(config)# no lldp enable ports ethernet 2/4 2/5 2/6
device(config)# lldp enable transmit ports ethernet 2/4 2/5 2/6
The above command changes the LLDP operating mode on ports 2/4, 2/5, and 2/6 from transmit and receive mode to transmit only mode.
```

Syntax: **[no] lldp enable transmit ports ethernet** *portlist* | **all**

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Configuring receive mode

To change the LLDP operating mode from receive and transmit mode to receive only mode, disable the transmit and receive mode, and enter the **lldp enable receive ports ethernet** command at the Global CONFIG level of the CLI.

```
device(config)# no lldp enable ports ethernet 2/4
device(config)# lldp enable receive ports ethernet 2/4
The above command changes the LLDP operating mode on port 2/4 from transmit and receive mode to receive only mode.
```

Syntax: **[no] lldp enable receive ports ethernet** *portlist* | **all**

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Specifying the maximum number of LLDP neighbors

You can change the limit of the number of LLDP neighbors for which LLDP data will be retained, per device as well as per port.

Per device

To change the maximum number of neighbors for which LLDP data is retained for the entire system, use the **lldp max-total-neighbors** command. The default number of LLDP neighbors per device is 392.

```
device(config)# lldp max-total-neighbors 392
```

Syntax: **[no] lldp max-total-neighbors** *value*

The *value* variable specifies the total number of LLDP neighbors per device with a range of 16 to 8192.

Per port

To change the maximum number of LLDP neighbors for which LLDP data is retained for each port, use the **lldp max-neighbors-per-port** command. The default is number of LLDP neighbors per port is 4.

```
device(config)# lldp max-neighbors-per-port 4
```

Syntax: **[no] lldp max-neighbors-per-port** *value*

The *value* variable specifies the number of LLDP neighbors per port with a range of 1 to 64.

Enabling bridging of LLDP BPDUs when LLDP not enabled

An interface which does not have LLDP enabled can be configured to bridge LLDP packets instead of dropping them. This action has to be specified explicitly by using the **forward-lldp** command.

NOTE

When LLDP is enabled this command will not have any effect on the behavior of LLDP. In other words, BPDUs will not be bridged.

The forward-lldp command must be issued on the physical port configuration, not in LAG configuration.

The LLDP BPDU forward command can be used at the interface level to allow bridging of LLDP BPDUs (LLDP BPDUs are normally dropped if LLDP is not configured on that interface).

```
device(config)# int e 2/1
device(config-if-e1000-1/2)#forward-lldp
```

Syntax: forward-lldp

Enabling LLDP SNMP notifications and Syslog messages

SNMP notifications and Syslog messages for LLDP provide data updates and general status.

When LLDP SNMP notifications are enabled, corresponding Syslog messages are enabled as well. When LLDP SNMP notifications are enabled, the device sends traps and corresponding Syslog messages whenever there are changes to the LLDP data received from neighboring devices.

LLDP SNMP notifications and corresponding Syslog messages are disabled by default. To enable SNMP notifications and Syslog messages on all interfaces, enter the **lldp enable snmp notifications ports all** command at the Global CONFIG level of the CLI.

```
device(config)# lldp enable snmp notifications ports all
```

Syntax: [no] lldp enable snmp notifications ports all

To enable or disable SNMP notifications and Syslog messages on a specific interface, enter the **lldp enable snmp notifications ports ethernet** command at the config level of the CLI.

```
device(config)# lldp enable snmp notifications ports ethernet 4/1
```

Syntax: [no] lldp enable snmp notifications ports ethernet *slot/port*

Specifying the minimum time between SNMP traps and Syslog messages

When SNMP notifications and Syslog messages for LLDP are enabled, the device sends no more than one SNMP notification and Syslog message within a 5 second period. You can adjust the amount of time between transmission of SNMP traps (lldpRemTablesChange) and Syslog messages from five seconds up 3600 seconds.

Use the **lldp snmp-notification-interval** command to change the amount of time between SNMP notifications.

```
device(config)# lldp snmp-notification-interval 5
```

Syntax: [no] lldp snmp-notification-interval *seconds*

The *seconds* variable specifies the notification interval with a range of 5 to 3600 seconds.

Changing the minimum time between LLDP transmissions

The LLDP transmit delay timer limits the number of LLDP frames an LLDP agent can send within a specified time frame. The LLDP transmit delay timer prevents an LLDP agent from transmitting a series of successive LLDP frames during a short time period, when rapid changes occur in LLDP. It also increases the probability that multiple changes, rather than single changes, will be reported in each LLDP frame.

When LLDP is enabled, the system automatically sets the LLDP transmit delay timer to the default of 2 seconds. To change the LLDP transmit delay timer setting, use the **lldp transmit-delay** command.

```
device(config)# lldp transmit-delay 2
```

Syntax: **[no] lldp transmit-delay** *seconds*

The *seconds* variable specifies the notification interval with a range of 1 to 8192 seconds.

NOTE

The LLDP transmit delay timer must not be greater than one quarter of the LLDP transmission interval (CLI command **lldp transmit-interval**).

Changing the interval between regular LLDP transmissions

The LLDP transmit interval specifies the number of seconds between regular LLDP packet transmissions. When LLDP is enabled, by default, the device waits 30 seconds between regular LLDP packet transmissions. To change the LLDP transmission interval, enter the **lldp transmit-interval** command.

```
device(config)# lldp transmit-interval 5
```

Syntax: **[no] lldp transmit-interval** *seconds*

The *seconds* variable specifies the notification interval with a range of 5 to 32768 seconds.

NOTE

Setting the transmit interval or transmit holdtime multiplier to inappropriate values can cause the LLDP agent to transmit LLDPDUs with TTL values that are excessively high. This in turn can affect how long a receiving device will retain the information if it is not refreshed.

Changing the holdtime multiplier for transmit TTL

The holdtime multiplier for transmit TTL is used to compute the actual time-to-live (TTL) value used in an LLDP frame. The TTL value is the length of time the receiving device should maintain the information. The default setting of holdtime multiplier for TTL to 4. This is the age out time for that particular advertisement. To compute the TTL value, the system multiplies the LLDP transmit interval by the holdtime multiplier. For example, if the LLDP transmit interval is 30 and the holdtime multiplier for TTL is 4, then the value 120 is encoded in the TTL field in the LLDP header.

To change the holdtime multiplier for TTL from the default value, use the **lldp transmit-hold** command.

```
device(config)# lldp transmit-hold 4
```

Syntax: **lldp transmit-hold** *value*

The *value* variable specifies holdtime multiplier for transmit TTL with a range of 4 to 10.

NOTE

Setting the transmit interval or transmit holdtime multiplier to inappropriate values can cause the LLDP agent to transmit LLDPDUs with TTL values that are excessively high.

Changing the minimum time between port reinitializations

The LLDP re-initialization delay timer specifies the minimum amount of time the device waits from when LLDP is disabled on a port, until it honors a request to re-enable LLDP on that port. When LLDP is enabled, the default is set to 2 seconds. The LLDP re-initialization delay timer ensures that there is a defined minimum amount of time between successive LLDP frame transmission, thereby preventing a large number of LLDP frames to be sent at one time.

To change the LLDP re-initialization delay timer, enter the **lldp reinit-delay** command.

```
device(config)# lldp reinit-delay 2
```

Syntax: **lldp reinit-delay** *seconds*

The *seconds* variable specifies the LLDP re-initialization delay timer with a range of 1 to 10 seconds.

LLDP TLVs advertised by the Extreme device

When LLDP is enabled on a global basis, the Extreme device automatically advertises the following information, except as specified.

General system information

- Management address
- Port description
- System capabilities
- System description (not automatically advertised)
- System name

Management address

The management address is an IPv4 address that can be used to manage the device. If no management address is explicitly configured to be advertised, the Extreme device will use the first available IPv4 address configured on the following types of interfaces, in the following order of preference:

- Physical port on which LLDP will be transmitting the packet
- Loopback interface
- Virtual routing interface (VE)
- Router interface on a VLAN of which the port is a member
- Other physical interface

If no IP address is configured, the port's current MAC address will be advertised.

To advertise the IPv4 management address, enter the **lldp advertise management-address ipv4** command.

```
device(config)#lldp advertise management-address ipv4 10.157.2.1 ports e 1/4
```

The management address will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
Management address (IPv4): 10.157.2.1
```

Syntax: **[no] lldp advertise management-address ipv4 *ipv4address* ports ethernet *portlist* | all**

ipv4 address is the address that may be used to reach higher layer entities to assist discovery by network management. In addition to the management address, the advertisement will include the system interface number and OID associated with the management address, if either or both are known.

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Port description

The port description TLV identifies the port from which the LLDP agent transmitted the advertisement. The port description is taken from the ifDescr MIB object from MIB-II.

By default, the port description is automatically advertised when LLDP is enabled on a global basis. To disable advertisement of the port description, enter the **lldp advertise port-description ports ethernet** command.

```
device(config)#no lldp advertise port-description ports e 2/4 to 2/12
```

The port description will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
Port description: "GigabitEthernet20"
```

Syntax: **[no] lldp advertise port-description ports ethernet *portlist* | all**

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements.

System capabilities

The system capabilities TLV identifies the primary functions of the device and indicates whether these primary functions are enabled. The primary functions can be one or more of the following:

- Repeater
- Bridge
- WLAN access point
- Router
- Telephone
- DOCSIS cable device
- Station only (devices that implement end station capability)
- Other

System capabilities for Extreme devices are based on the type of software image in use.

By default, the system capabilities are automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter the **lldp advertise system-capabilities ports ethernet** command.

```
device(config)#no lldp advertise system-capabilities ports e 2/4 to 2/12
```

The system capabilities will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
System capabilities :   bridge
Enabled capabilities:   bridge
```

Syntax: **[no] lldp advertise system-capabilities ports ethernet *portlist* | all**

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements.

System description

The system description is the network entity. The information corresponds to the sysDescr MIB object. To advertise the system description, enter the **lldp advertise system-description ports ethernet** command.

```
device(config)#lldp advertise system-description ports e 2/4 to 2/12
```

The system description will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**):

```
device# show lldp local-info
Local port: 8/13
+ Chassis ID (MAC address): 0024.3891.1100
+ Port ID (MAC address): 0024.3891.125c
+ Time to live: 40 seconds
+ System name       : "IxANVL-1"
+ Port description  : "GigabitEthernet8/13"
+ System description: "Extreme MLXe (System Mode: MLX), IronWare Version V\
                    5.3.0T163 Compiled on Jan  3 2012 at 18:01:00 label\
                    ed as V5.3.00b460"
+ System capabilities : bridge, router
  Enabled capabilities: bridge, router
+ 802.3 MAC/PHY       : auto-negotiation enabled
  Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
                        100BaseTX-FD, 1000BaseT-HD, 1000BaseT-FD
  Operational MAU type  : 1000BaseT-FD
+ Link aggregation: not capable
+ Maximum frame size: 9216 octets
+ Port VLAN ID: 813
+ Management address (IPv4): 10.1.1.190
+ Management address (IPv4): 10.20.103.190
+ Management address (IPv6): 2001:DB8
+ Port-Protocol VLAN ID: not supported
```

Syntax: **[no] lldp advertise system-description ports ethernet *portlist* | all**

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements.

System name

The system name is taken from the sysName MIB object. The sysName MIB object corresponds to the name defined with the CLI command **hostname** . By default, the system name is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter the **lldp advertise system-name ports ethernet** command.

```
device(config)#no lldp advertise system-name ports e 2/4 to 2/12
```

The system name will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
System name: "NI"
```

Syntax: [no] **lldp advertise system-name ports ethernet slotnum/portnum | all**

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements.

802.1 capabilities

Except for the VLAN name, the Extreme device will advertise the following 802.1 attributes when LLDP is enabled on a global basis:

- VLAN name (not automatically advertised)
- Untagged VLAN ID

VLAN name

The VLAN name TLV contains the name and VLAN ID of a VLAN configured on a port. An LLDPDU may include multiple instances of this TLV, each for a different VLAN.

To advertise the VLAN name, enter the **lldp advertise vlan-name vlan** command.

```
device(config)#lldp advertise vlan-name vlan 99 ports e 2/4 to 2/12
```

The VLAN name will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
VLAN name (VLAN 99): "Voice-VLAN-99"
```

Syntax: [no] **lldp advertise vlan-name vlan vlanID ports ethernet portlist | all**

For *vlan ID*, enter the VLAN ID to advertise.

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements.

Port and Protocol VLAN ID

The port and protocol VLAN TLV indicates if a port is capable of supporting port and protocol VLANs and whether it is enabled on the port. If port and protocol VLANs are enabled on the port, the advertisement also contains the port and protocol VLAN ID (PPVID). If the port is not capable of supporting port and protocol VLANs, or if the port is not enabled with any port and protocol VLAN, the PPVID number will be zero.

Use the **lldp advertise port-protocol-vlan-id ports ethernet** command to enable or disable advertising the port and protocol VLAN ID.

```
device(config)#lldp advertise port-protocol-vlan-id ports e 2/4 to 2/12
```

The port and protocol VLAN ID advertisement will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**):

```
Port-Protocol VLAN ID: not supported
```

Syntax: **[no] lldp advertise port-protocol-vlan-id ports ethernet** *portlist* | **all**

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Untagged VLAN ID

The port VLAN ID TLV advertises the Port VLAN Identifier (PVID) that will be associated with untagged or priority-tagged frames. If the port is not an untagged member of any VLAN (that is, the port is strictly a tagged port), the value zero will indicate that.

By default, the port VLAN ID is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise port-vlan-id ports e 2/4 to 2/12
```

The untagged VLAN ID will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**):

```
Port VLAN ID: 99
```

Syntax: **[no] lldp advertise port-vlan-id ports ethernet** *portlist* | **all**

For *port list*, specify the ports in the format *[slotnum/]portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements.

802.3 capabilities

Except for Power-via-MDI information, the Extreme device will advertise the following 802.3 attributes when LLDP is enabled on a global basis:

- Link aggregation information
- MAC/PHY configuration and status
- Maximum frame size

Link aggregation

Extreme devices advertise link aggregation information about standard link aggregation (LACP) as well as static trunk configuration. By default, link-aggregation information is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter the **lldp advertise link-aggregation ports ethernet** command.

```
device(config)#no lldp advertise link-aggregation ports e 2/12
```

Syntax: [no] **lldp advertise link-aggregation ports ethernet** *portlist* | **all**

The link aggregation advertisement will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
Link aggregation: not capable
```

For *port list*, specify the ports in the format [*slotnum*/]*portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements.

MAC/PHY configuration status

The MAC/PHY configuration and status TLV includes the following information:

- Auto-negotiation capability and status
- Speed and duplex mode
- Flow control capabilities for auto-negotiation
- Port speed down-shift and maximum port speed advertisement
- If applicable, indicates if the above settings are the result of auto-negotiation during link initiation or of a manual set override action

By default, the MAC/PHY configuration and status information are automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter the **lldp advertise mac-phy-config-status ports ethernet** command.

```
device(config)#no lldp advertise mac-phy-config-status ports e 2/4 to 2/12
```

The MAC/PHY configuration advertisement will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
+ 802.3 MAC/PHY      : auto-negotiation enabled
  Advertised capabilities: 10baseT-HD, 10baseT-FD, 100baseTX-HD, 100baseTX-FD,
  fdxSPause, fdxBPause, 1000baseT-HD, 1000baseT-FD
  Operational MAU type: 100BaseTX-FD
```

Syntax: [no] **lldp advertise mac-phy-config-status ports ethernet** *portlist* | **all**

For *port list*, specify the ports in the format [*slotnum*/]*portnum*, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements. Maximum frame size

Maximum frame size TLV

The maximum frame size TLV provides the maximum 802.3 frame size capability of the port. This value is expressed in octets and includes the four-octet Frame Check Sequence (FCS). The default maximum frame size is 1522. The advertised value may change depending on whether the **aggregated-vlan** or **jumbo** commands are in effect.

By default, the maximum frame size is automatically advertised when LLDP is enabled on a global basis. To disable this advertisement, enter a command such as the following.

```
device(config)#no lldp advertise max-frame-size ports e 2/4 to 2/12
```

The maximum frame size advertisement will appear similar to the following on the remote device, and in the CLI display output on the Extreme device (**show lldp local-info**).

```
Maximum frame size: 1522 octets
```

Syntax: `[no] lldp advertise max-frame-size ports ethernet portlist | all`

For *port list*, specify the ports in the format `[slotnum/]portnum`, where *slotnum* is required on chassis devices only. You can list all of the ports individually, use the keyword **to** to specify a range of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Note that using the keyword **all** may cause undesirable effects on some ports. The configuration will be applied to all ports, however, the ports that are not members of any VLAN will not send VLAN name advertisements. Maximum frame size

Displaying LLDP statistics and configuration settings

You can use the following CLI **show** commands to display information about LLDP settings and statistics:

- **show lldp** - Displays a summary of the LLDP configuration settings.
- **show lldp statistics** - Displays LLDP global and per-port statistics.
- **show lldp neighbors** - Displays a list of the current LLDP neighbors.
- **show lldp neighbors detail** - Displays the details of the latest advertisements received from LLDP neighbors.
- **show lldp local-info** - Displays the details of the LLDP advertisements that will be transmitted on each port.

LLDP configuration summary

To display a summary of the LLDP configuration settings on the device, enter the **show lldp** command at any level of the CLI.

The following shows an example report.

```
device#show lldp
LLDP transmit interval           : 10 seconds
LLDP transmit hold multiplier    : 4   (transmit TTL: 40 seconds)
LLDP transmit delay              : 1 seconds
LLDP SNMP notification interval  : 5 seconds
LLDP reinitialize delay          : 1 seconds
LLDP maximum neighbors          : 392
LLDP maximum neighbors per port  : 4
```

Syntax: `show lldp`

[Table 48](#) describes the information displayed by the **show lldp statistics** command.

TABLE 48 Output descriptions for the **show lldp statistics** command

Field	Description
LLDP transmit interval	The number of seconds between regular LLDP packet transmissions.

TABLE 48 Output descriptions for the **show lldp statistics** command (continued)

Field	Description
LLDP transmit hold multiplier	The multiplier used to compute the actual time-to-live (TTL) value of an LLDP advertisement. The TTL value is the transmit interval multiplied by the transmit hold multiplier.
LLDP transmit delay	The number of seconds the LLDP agent will wait after transmitting an LLDP frame and before transmitting another LLDP frame.
LLDP reinitialize delay	The minimum number of seconds the device will wait from when LLDP is disabled on a port, until a request to re-enable LLDP on that port will be honored.
LLDP maximum neighbors	The maximum number of LLDP neighbors for which LLDP data will be retained, per device.
LLDP maximum neighbors per port	The maximum number of LLDP neighbors for which LLDP data will be retained, per port.

LLDP statistics

The **show lldp statistics** command displays an overview of LLDP neighbor detection on the device, as well as packet counters and protocol statistics. The statistics are displayed on a global and per-port basis.

The following shows an example report.

```

device#show lldp statistics
Last neighbor change time: 23 hours 50 minutes 40 seconds ago
Neighbor entries added      : 14
Neighbor entries deleted    : 5
Neighbor entries aged out   : 4
Neighbor advertisements dropped : 0
Port      Tx Pkts  Rx Pkts  Rx Pkts  Rx Pkts  Rx TLVs  Rx TLVs  Neighbors
          Total    Total    w/Errors Discarded Unrecognz Discarded Aged Out
1         60963    75179      0         0         0         0         4
2          0         0         0         0         0         0         0
3         60963    60963      0         0         0         0         0
4         60963    121925     0         0         0         0         0
5          0         0         0         0         0         0         0
6          0         0         0         0         0         0         0
7          0         0         0         0         0         0         0
8          0         0         0         0         0         0         0
9          0         0         0         0         0         0         0
10         60974     0         0         0         0         0         0
11          0         0         0         0         0         0         0
12          0         0         0         0         0         0         0
13          0         0         0         0         0         0         0
14          0         0         0         0         0         0         0

```

Syntax: show lldp statistics

NOTE

You can reset LLDP statistics using the CLI command **clear LLDP statistics**. Refer to [Resetting LLDP statistics](#) on page 209.

NOTE

LLDP statistics are not preserved in the event of a module switchover.

[Table 49](#) describes the information displayed by the **show lldp statistics** command.

TABLE 49 Neighbor detection output description for the **show lldp statistics** command

Field	Description
Last neighbor change time	The elapsed time (in hours, minutes, and seconds) since a neighbor last advertised information. For example, the elapsed time since a neighbor was last added, deleted, or its advertised information changed.
Neighbor entries added	The number of new LLDP neighbors detected since the last reboot or since the last time the clear lldp statistics all command was issued. This number includes the number entries added after timing out or aging out.

TABLE 49 Neighbor detection output description for the **show lldp statistics** command (continued)

Field	Description
Neighbor entries deleted	The number of LLDP neighbors deleted since the last reboot or since the last time the clear lldp statistics all command was issued. This number includes the number of entries deleted after timing out or aging out.
Neighbor entries aged out	The number of LLDP neighbors dropped on all ports after the time-to-live expired. Note that LLDP entries age out naturally when a port's cable or module is disconnected or when a port becomes disabled. However, if a disabled port is re-enabled, the system will delete the old LLDP entries.
Neighbor advertisements dropped	The number of valid LLDP neighbors the device detected, but could not add. This can occur, for example, when a new neighbor is detected and the device is already supporting the maximum number of neighbors possible. This can also occur when an LLDPDU is missing a mandatory TLV or is not formatted correctly.
Port	The local port number.
Tx Pkts Total	The number of LLDP packets the port transmitted.
Rx Pkts Total	The number of LLDP packets the port received.
Rx Pkts w/Errors	The number of LLDP packets the port received that have one or more detectable errors.
Rx Pkts Discarded	The number of LLDP packets the port received then discarded.
Rx TLVs Unrecognz	The number of TLVs the port received that were not recognized by the LLDP local agent. Unrecognized TLVs are retained by the system and can be viewed in the output of the show LLDP neighbors detail command or retrieved through SNMP.
Rx TLVs Discarded	The number of TLVs the port received then discarded.
Neighbors Aged Out	The number of times a neighbor's information was deleted because its TTL timer expired.

LLDP neighbors

The **show lldp neighbors** command displays a list of the current LLDP neighbors per port.

The following shows an example report.

```
device#show lldp neighbors
Lcl Port Chassis ID      Port ID      Port Description      System Name
1         0004.1234.0fc0 0004.1234.0fc0 GigabitEthernet9/1    BigIron RX 32~
1         00e0.5201.4000 00e0.5201.4000 GigabitEthernet0/1/1  BigIron RX 4~
3         00e0.5211.0200 00e0.5211.0203 GigabitEthernet4      BigIron RX 4~
4         00e0.5211.0200 00e0.5211.0202 GigabitEthernet3      BigIron RX 16~
4         00e0.5211.0200 00e0.5211.0210 GigabitEthernet17     BigIron RX 4~
15        00e0.5211.0200 00e0.5211.020f GigabitEthernet16     BigIron RX 8~
16        00e0.5211.0200 00e0.5211.020e GigabitEthernet15     BigIron RX 16~
17        00e0.5211.0200 00e0.5211.0211 GigabitEthernet18     BigIron RX 4~
18        00e0.5211.0200 00e0.5211.0210 GigabitEthernet17     BigIron RX 4~
```

Syntax: show lldp neighbors

The following table describes the information displayed by the **show lldp neighbors** command.

TABLE 50 Output descriptions of the show lldp neighbors command

Field	Description
Lcl Port	The local LLDP port number.
Chassis ID	The identifier for the device. Extreme devices use the base MAC address of the device as the Chassis ID.
Port ID	The identifier for the port. Extreme devices use the permanent MAC address associated with the port as the port ID.

TABLE 50 Output descriptions of the show lldp neighbors command (continued)

Field	Description
Port Description	The description for the port. Extreme devices use the ifDescr MIB object from MIB-II as the port description.
System Name	The administratively-assigned name for the system. Extreme devices use the sysName MIB object from MIB-II, which corresponds to the CLI hostname command setting. NOTE: A tilde (~) at the end of a line indicates that the value in the field is too long to display in full and is truncated.

LLDP neighbors detail

The **show lldp neighbors detail** command displays the LLDP advertisements received from LLDP neighbors.

The following shows an example **show lldp neighbors detail** report.

NOTE

The **show lldp neighbors detail** output will vary depending on the data received. Also, values that are not recognized or do not have a recognizable format, may be displayed in hexadecimal binary form.

```
device#show lldp neighbors detail ports e 8/13
Local port: 8/13
+ Chassis ID (MAC address): 0024.3891.1100
+ Port ID (MAC address): 0024.3891.125c
+ Time to live: 40 seconds
+ System name       : "IxANVL-1"
+ Port description  : "GigabitEthernet8/13"
+ System description : "Extreme MLXe (System Mode: MLX), IronWare Version V\
5.3.0T163 Compiled on Jan  3 2012 at 18:01:00 label\
ed as V5.3.00b460"
+ System capabilities : bridge, router
  Enabled capabilities: bridge, router
+ 802.3 MAC/PHY       : auto-negotiation enabled
  Advertised capabilities: 10BaseT-HD, 10BaseT-FD, 100BaseTX-HD,
100BaseTX-FD, 1000BaseT-HD, 1000BaseT-FD
  Operational MAU type  : 1000BaseT-FD
+ Link aggregation: not capable
+ Maximum frame size: 9216 octets
+ Port VLAN ID: 813
+ Management address (IPv4): 10.1.1.190
+ Management address (IPv4): 10.20.103.190
+ Management address (IPv6): 2001:DB8
+ Port-Protocol VLAN ID: not supported
Local port: 8/23
+ Chassis ID (MAC address): 0024.3891.1100
+ Port ID (MAC address): 0024.3891.1266
+ Time to live: 40 seconds
+ System name       : "IxANVL-1"
+ System description : "Extreme MLXe (System Mode: MLX), IronWare Version V\
5.3.0T163 Compiled on Jan  3 2012 at 18:01:00 label\
ed as V5.3.00b460"
+ Port VLAN ID: 1
+ Management address (IPv4): 10.1.1.190
+ Management address (IPv4): 10.20.103.190
+ Management address (IPv6): 2001:DB8
```

A backslash (\) at the end of a line indicates that the text continues on the next line.

Except for the Neighbor field, the fields in the previous output are described in the individual TLV advertisement sections in this chapter. The Neighbor field displays the source MAC address from which the packet was received, and the remaining TTL for the neighbor entry.

Syntax: **show lldp neighbors detail** [**ports ethernet slotnum/portnum** | **all**]

If you do not specify any ports or use the keyword **all**, by default, the report will show the LLDP neighbor details for all ports.

You can list all of the ports individually, use the keyword **to** to specify ranges of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

LLDP configuration details

The **show lldp local-info** command displays the local information advertisements (TLVs) that will be transmitted by the LLDP agent.

NOTE

The **show lldp local-info** output will vary based on LLDP configuration settings.

The following shows an example report.

```
device#show lldp local-info ports ethernet 1/40
Local port: 1/40
+ Chassis ID (MAC address): 001b.edb3.f180
+ Port ID (MAC address): 001b.edb3.f1a8
+ Time to live: 40 seconds
+ System name      : "CES-151"
+ Port description : "GigabitEthernet1/40"
+ System description : "Extreme NetIron CES, IronWare Version V5.3.0T183 Co\
                        mpiled on Jan 03 2012 at 18:18:17 labeled as V5.3.0\
                        0b460"
+ System capabilities : bridge, router
  Enabled capabilities: bridge, router
+ 802.3 MAC/PHY      : auto-negotiation enabled
  Advertised capabilities: 1000BaseX-FD
  Operational MAU type : 1000BaseT-FD
+ Link aggregation: aggregated (aggregated port ifIndex: 3)
+ Maximum frame size: 9216 octets
+ Port VLAN ID: 1
+ Management address (IPv4): 10.1.1.151
+ Management address (IPv4): 10.20.103.151
+ Management address (IPv6): 2001:DB8
+ Port-Protocol VLAN ID: not supported
```

A backslash (\) at the end of a line indicates that the text continues on the next line.

The fields in the above output are described in the individual TLV advertisement sections in this chapter.

Syntax: **show lldp local-info** [**ports ethernet slot num/portnum** | **all**]

If you do not specify any ports or use the keyword **all**, by default, the report will show the local information advertisements for all ports.

You can list all of the ports individually, use the keyword **to** to specify ranges of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

Resetting LLDP statistics

To reset LLDP statistics, enter the **clear lldp statistics** command at the Global CONFIG level of the CLI. The Extreme device will clear the global and per-port LLDP neighbor statistics on the device (refer to [LLDP statistics](#) on page 206).

```
device#clear lldp statistics
```

Syntax: **clear lldp statistics** [**ports ethernet slot num/portnum** | **all**]

If you do not specify any ports or use the keyword **all**, by default, the system will clear lldp statistics on all ports.

You can list all of the ports individually, use the keyword **to** to specify ranges of ports, or a combination of both. To apply the configuration to all ports on the device, use the keyword **all** instead of listing the ports individually.

SNMP

• SNMP overview.....	211
• Adding an SNMP community string.....	211
• Displaying the SNMP community strings.....	212
• Using the User-Based Security model.....	213
• Defining SNMP views.....	218
• SNMP v3 configuration examples.....	218
• Configuring SNMP traps.....	219
• Configuring SNMP management of VRFs	222
• Configuring SNMP ifIndex	224
• SNMP scalability optimization.....	225
• Configuring SNMP to revert ifType to legacy values	226
• Configuring snAgentConfigModuleType to return original values.....	226
• Preserving interface statistics in SNMP.....	227

SNMP overview

Simple Network Management Protocol (SNMP) is a set of protocols for managing complex networks. SNMP sends messages, called protocol data units (PDUs), to different parts of a network. An SNMP-compliant device, called an agent, stores data about itself in Management Information Bases (MIBs) and SNMP requesters or managers.

SNMP versions 1 and 2c use community strings to restrict SNMP access. The default passwords for SNMP access are the SNMP community strings configured on the device:

- The default read-only community string is "public"
- Use this community string for any SNMP Get, GetNext, or GetBulk request

By default, you cannot perform any SNMP Set operations since a read-write community string is not configured.

You can configure as many additional read-only and read-write community strings as you need. The number of strings you can configure depends on the memory on the device. There is no practical limit. If you delete all read-only community strings, the device automatically re-adds the default "public" read-only community string the next time you load the software, or you disable and re-enable the SNMP feature.

Encryption of SNMP community strings

Encryption is enabled by default. The software automatically encrypts SNMP community strings. Users with read-only access or who do not have access to management functions in the CLI cannot display the strings. For users with read-write access, the strings are encrypted in the CLI but are shown in the clear in the Web Management Interface.

To display the community strings in the CLI, first use the **enable password-display** command and then use the **show snmp server** command. This will display both the read-only and read-write community strings in the clear.

Adding an SNMP community string

By default, the string is encrypted. To add a community string, enter commands such as the following.

```
device(config)# snmp-server community private rw
```

The command adds the read-write SNMP community string "private".

Syntax: `[no] snmp-server community string ro | rw [view viewstring] [standard-acl-name | standard-acl-id | ipv6 ipv6-acl-name]`

The *string* parameter specifies the community string name. The string can be up to 32 characters long.

The system modifies the configuration to `session 10.1.1.1 key 2 $XkBTb24tb0RuXA==`

For example, the following portion of the code has the encrypted code "2".

```
snmp-server community 2
$D?@d=8 rw
```

The prefix can be one of the following:

- 1 = the community string uses proprietary simple cryptographic 2-way algorithm (only for NetIron CES and NetIron CER)
- 2 = the community string uses proprietary base64 cryptographic 2-way algorithm (only for NetIron XMR and NetIron MLX)

The **ro** parameter specifies the string is read-only.

The **rw** parameter specifies the string is read-write.

The **view viewstring** parameter is optional. It allows you to associate a view to the members of this community string. Enter up to 32 alphanumeric characters. If no view is specified, access to the full MIB is granted. The view that you want must exist before you can associate it to a community string. Here is an example of how to use the view parameter in the community string command.

```
device(config)# snmp-s community myread ro view sysview
```

The command in this example associates the view "sysview" to the community string named "myread". The community string has read-only access to "sysview". For information on how create views, refer to the section "Defining SNMP views".

The *standard-acl-name | standard-acl-id | ipv6ipv6-acl-name* parameter is optional. It allows you to specify which ACL is used to filter the incoming SNMP packets. You can enter either the ACL name or its ID for an IPv4 ACL; for an IPv6 ACL, you must enter the keyword **ipv6** followed by the name of the IPv6 ACL. Here are examples.

```
device(config) # snmp-s community myread ro view sysview 2
device(config) # snmp-s community myread ro view sysview myacl
```

The command in the first example specifies that ACL group 2 filters incoming SNMP packets, whereas the command in the second example uses the IPv4 ACL group called "myacl" to filter incoming packets.

Displaying the SNMP community strings

To display the community strings in the CLI, first use the **enable password-display** command and then use the **show snmp server** command. This will display both the read-only and read-write community strings in the clear.

To display the configured community strings, enter the following command at any CLI level.

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

Syntax: `show snmp server`

NOTE

If display of the strings is encrypted, the strings are not displayed. Encryption is enabled by default.

Using the User-Based Security model

SNMP version 3 (RFC 2570 through 2575) introduces a User-Based Security model (RFC 2574) for authentication and privacy services.

SNMP version 1 and version 2 use community strings to authenticate SNMP access to management modules. This method can still be used for authentication. In SNMP version 3, the User-Based Security model of SNMP can be used to secure against the following threats:

- Modification of information
- Masquerading the identity of an authorized entity
- Message stream modification
- Disclosure of information

Furthermore, SNMP version 3 supports View-Based Access Control Mechanism (RFC 2575) to control access at the PDU level. It defines mechanisms for determining whether or not access to a managed object in a local MIB by a remote principal should be allowed. (Refer to the section "Defining SNMP views" .)

Configuring your NMS

To be able to use the SNMP version 3 features.

1. Make sure that your Network Manager System (NMS) supports SNMP version 3.
2. Configure your NMS agent with the necessary users.
3. Configure the SNMP version 3 features in the device.

Configuring SNMP version 3 on the device

To configure SNMP version 3 on the device, perform the tasks listed below.

1. Enter an engine ID for the management module using the **snmp-server engineid** command if you will not use the default engine ID. Refer to "Defining the engine ID".
2. Create views that will be assigned to SNMP user groups using the **snmp-server view** command. Refer to the "Defining SNMP views" for details.
3. (Optional) Create access lists (ACLs) to filter incoming SNMP packets according to rules in the ACL. The following ACL types are supported for SNMP:
 - Standard, named and numbered IPv4 ACLs.
 - IPv6 ACLs.
4. Create user groups using the **snmp-server group** command. You can optionally assign an ACL to a user group. Refer to "Defining an SNMP group".
5. Create user accounts and associate these accounts to user groups using the **snmp-server user** command. Refer to "Defining an SNMP user account".

If SNMP version 3 is not configured, then community strings by default are used to authenticate access.

Even if SNMP version 3 users are configured on the device, the system will still accept SNMP version 1, 2c and 3 PDUs from the remote manager.

Defining the engine ID

A default engine ID is generated during system start up. The format of the default engine ID is derived from RFC 2571 (Architecture for SNMP frameworks) within the MIB description for object `SnmpEngineID`.

To determine what the default engine ID of the device is, enter the **show snmp engineid** command and find the following line.

```
Local SNMP Engine ID: 800007c70300e05290ab60
```

Refer to the Displaying the engine ID section for details.

The default engine ID guarantees the uniqueness of the engine ID for SNMP version 3. If you want to change the default engine ID, enter a command such as the following.

```
device(config)# snmp-server engineid local 800007c70300e05290ab60
```

Syntax: `[no] snmp-server engineid local hex-string`

The **local** parameter indicates that engine ID to be entered is the ID of this device, representing an SNMP management entity.

NOTE

Since the current implementation of SNMP version 3 does not support Notification, remote engine IDs cannot be configured at this time.

The *hex-string* variable consists of 11 octets, entered as hexadecimal values. Each octet has two hexadecimal characters. The engine ID should contain an even number of hexadecimal characters.

The default engine ID has a maximum of 11 octets:

- Octets 1 through 4 represent the agent's SNMP management private enterprise number as assigned by the Internet Assigned Numbers Authority (IANA). The most significant bit of Octet 1 is "1". For example, "000007c7" is the ID for Extreme Networks in hexadecimal. With Octet 1 always equal to "1", the first four octets in the default engine ID is always "800007c7" (which is 1991 in decimal).
- Octet 5 is always 03 in hexadecimal and indicates that the next set of values represent a MAC address.
- Octets 6 through 11 form the MAC address of the lowest port in the management module.

NOTE

Engine ID must be a unique number among the various SNMP engines in the management domain. Using the default engine ID ensures the uniqueness of the numbers.

Defining an SNMP group

SNMP groups map SNMP users to SNMP views. For each SNMP group, you can configure a notify view, a read view, a write view, or combinations of the above. Users who are mapped to a group will use its views for access control.

NOTE

This topic is for SNMP v3, but not for SNMP v1 or SNMP v2c. In those versions, groups and group views are created internally using community strings. (Refer to "Establishing SNMP community strings".) When a community string is created, two groups are created, based on the community string name. One group is for SNMP version 1 packets, while the other is for SNMP version 2 packets.

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

```
device# configure terminal
```

2. Enter the **snmp-server group** command.

```
device(config)# snmp-server group admin v3 auth ipv6 acl_1 read all write all notify all
```

Defining an SNMP user account

The **snmp-server user** command does the following:

- Creates an SNMP user.
- Defines the group to which the user will be associated.
- Defines the type of authentication to be used for SNMP access by this user.

Here is an example of how to create the account.

```
device(config)# snmp-s user bob admin v3 access 2 auth md5 bobmd5 priv des bobdes
```

The CLI for creating SNMP version 3 users has been updated as follows.

Syntax: **[no] snmp-server user** *name groupname v3* **[[access** *standard-acl-id* **] [[encrypted] auth** *md5 md5-password* **| sha** *sha-password* **[priv [encrypted] des** *des-password-key* **| aes** *aes-password-key* **]]**

The *name* parameter defines the SNMP user name or security name used to access the management module.

The *groupname* parameter identifies the SNMP group to which this user is associated or mapped. All users must be mapped to an SNMP group. Groups are defined using the **snmp-server group** command.

NOTE

The SNMP group to which the user account will be mapped should be configured before creating the user accounts; otherwise, the group will be created without any views. Also, ACL groups must be configured before configuring user accounts.

The **v3** parameter is required.

The **access** *standard-acl-id* **parameter is optional. It indicates that incoming SNMP packets are filtered based on the ACL attached to the user account.**

NOTE

The ACL specified in a user account overrides the ACL assigned to the group to which the user is mapped. If no ACL is entered for the user account, the ACL configured for the group is used to filter packets.

The **encrypted** parameter means that the MD5 or SHA password will be a digest value. MD5 has 16 octets in the digest. SHA has 20. The digest string has to be entered as a hexadecimal string. In this case, the agent need not generate any explicit digest. If the **encrypted** parameter is not used, the user is expected to enter the authentication password string for MD5 or SHA. The agent converts the password string to a digest, as described in RFC 3414.

The optional **auth** *md5* **|** *sha* **parameter defines the type of encryption the user must have to be authenticated. The choices are MD5 and SHA encryption (the two authentication protocols used in SNMP version 3).**

The *md5-password* and *sha-password* define the password the user must use to be authenticated. These password must have a minimum of 8 characters. If the **encrypted** parameter is used, then the digest has 16 octets for MD5 or 20 octets for SHA.

NOTE

Once a password string is entered, the generated configuration displays the digest (for security reasons), not the actual password.

The **priv** [encrypted] parameter is optional after you enter the md5 or sha password. The **priv** parameter specifies the encryption that is used to encrypt the privacy password. If the **encrypted** keyword is used, do the following:

- If DES is the privacy protocol to be used, enter **desdes-password-key** and enter a 16-octet DES key in hexadecimal format for the *des-password-key*. If you include the **encrypted** keyword, enter a password string of at least 8 characters.
- If AES is the privacy protocol to be used, enter **aes** and an *aes-password-key*. Enter either 12 (for a small key) or 16 (for a big key) characters for the *aes-password-key*. If you include the **encrypted** keyword, enter a password string containing 32 hexadecimal characters.

Displaying the engine ID

To display the engine ID of a management module, enter a command such as the following.

```
device(config)# show snmp engineid
Local SNMP Engine ID: 800007c70300e05290ab60
Engine Boots: 3
Engine time: 5
```

Syntax: show snmp engineid

The engine ID identifies the source or destination of the packet.

The engine boots represents the number of times that the SNMP engine reinitialized itself with the same engine ID. If the engineID is modified, the boot count is reset to 0.

The engine time represents the current time with the SNMP agent.

Displaying SNMP groups

To display the definition of an SNMP group, enter a command such as the following.

```
device# show snmp group
groupname = snmp_group_ipv6
security model = v3
security level = none
ACL id = 0
IPv6 ACL name: snmp_acl
readview = <none>
writeview = <none>
notifyview = <none>

groupname = snmp_group
security model = v3
security level = none
ACL id = 1
IPv6 ACL name: <none>
readview = <none>
writeview = <none>
notifyview = <none>
```

Syntax: show snmp group

The value for security level can be one of the following.

TABLE 51 Security Level and Authentication

Security level	Authentication
<i>none</i>	If the security model shows v1 or v2, then security level is blank. User names are not used to authenticate users; community strings are used instead.
noauthNoPriv	Displays if the security model shows v3 and user authentication is by user name only.

TABLE 51 Security Level and Authentication (continued)

Security level	Authentication
authNoPriv	Displays if the security model shows v3 and user authentication is by user name and the MD5 or SHA algorithm.
authPriv	Authentication uses MD5 or SHA. Encryption uses DES and AES protocol.

Displaying user information

To display the definition of an SNMP user account, enter a command such as the following.

```
device(config)# show snmp user
username = bob
acl id = 0
group = bobgroup
security model = v3
group acl id = 0
authtype = md5
authkey = ad172674ebc09cd9448c8276db0d12f8
privtype = aes
privkey = 3c154b47996534b22b22758e23f9a71a
engine ID= 800007c703000cdbf48a00
```

Syntax: show snmp user

Interpreting varbinds in report packets

If an SNMP version 3 request packet is to be rejected by an SNMP agent, the agent sends a report packet that contains one or more varbinds. The varbinds contain additional information, showing the cause of failures. An SNMP manager application decodes the description from the varbind. The following table presents a list of varbinds supported by the SNMP agent.

TABLE 52 Varbinds supported by the SNMP agent

Varbind object identifier	Description
1. 3. 6. 1. 6. 3. 11. 2. 1. 3. 0	Unknown packet data unit.
1. 3. 6. 1. 6. 3. 12. 1. 5. 0	The value of the varbind shows the engine ID that needs to be used in the snmp-server engineid command
1. 3. 6. 1. 6. 3. 15. 1. 1. 1. 0	Unsupported security level.
1. 3. 6. 1. 6. 3. 15. 1. 1. 2. 0	Not in time packet.
1. 3. 6. 1. 6. 3. 15. 1. 1. 3. 0	Unknown user name. This varbind can also be generated if either the: <ul style="list-style-type: none"> Configured ACL for the user filters out the packet. Group associated with the user is unknown.
1. 3. 6. 1. 6. 3. 15. 1. 1. 4. 0	Unknown engine ID. The value of this varbind would be the correct authoritative engineID that should be used.
1. 3. 6. 1. 6. 3. 15. 1. 1. 5. 0	Wrong digest.
1. 3. 6. 1. 6. 3. 15. 1. 1. 6. 0	Decryption error.

Defining SNMP views

SNMP views are named groups of MIB objects that can be associated with user accounts to allow limited access for viewing and modification of SNMP statistics and system configuration. SNMP views can also be used with other commands that take SNMP views as an argument. SNMP views reference MIB objects using object names, numbers, wildcards, or a combination of the three. The numbers represent the hierarchical location of the object in the MIB tree. You can reference individual objects in the MIB tree or a subset of objects from the MIB tree.

You can create up to 10 views on the device. This number cannot be changed.

To create an SNMP view, enter one of the following commands.

```
device(config)# snmp-server view Maynes system included
device(config)# snmp-server view Maynes system.2 excluded
device(config)# snmp-server view Maynes 2.3.*.6 included
device(config)# write mem
```

NOTE

The **snmp-server view** command supports the MIB objects as defined in RFC 1445.

Syntax: `[no] snmp-server view name mib_tree included | excluded`

The *name* parameter can be any alphanumeric name you choose to identify the view. The names cannot contain spaces.

The *mib_tree* parameter is the name of the MIB object or family. MIB objects and MIB sub-trees can be identified by a name or by the numbers called Object Identifiers (OIDs) that represent the position of the object or sub-tree in the MIB hierarchy. You can use a wildcard (*) in the numbers to specify a sub-tree family.

The **included | excluded** parameter specifies whether the MIB objects identified by the *mib_family* parameter are included in the view or excluded from the view.

NOTE

All MIB objects are automatically excluded from any view unless they are explicitly included; therefore, when creating views using the **snmp-server view** command, indicate which portion of the MIB you want users to access. For example, you may want to assign the view called "admin" a community string or user group. The "admin" view will allow access to the Unified IP MIB objects that begin with the 10.3.6.1.4.1.1991 object identifier. Enter the following command.

```
device(config)# snmp-server view admin 10.3.6.1.4.1.1991 included
```

You can exclude portions of the MIB within an inclusion scope. For example, if you want to exclude the snAgentSys objects, which begin with 10.3.6.1.4.1.1991.1.1.2 object identifier from the admin view, enter a second command such as the following.

```
device(config)# snmp-server view admin 10.3.6.1.4.1.1991.1.1.2 excluded
```

Note that the exclusion is within the scope of the inclusion.

To delete a view, use the **no** parameter before the command.

SNMP v3 configuration examples

The following examples show how to configure SNMP v3.

Simple SNMP v3 configuration

```
device(config)#snmp-s group admingrp v3 priv read all write all notify all
device(config)#snmp-s user adminuser admingrp v3 auth md5 auth password priv privacy password
device(config)#snmp-s host dest-ip adminuser
```

More detailed SNMP v3 configuration

```
device(config)#snmp-server view internet internet included
device(config)#snmp-server view system system included
device(config)#snmp-server community ..... ro
device(config)#snmp-server community ..... rw
device(config)#snmp-server contact isc-operations
device(config)#snmp-server location sdh-pillbox
device(config)#snmp-server host 10.91.255.32 .....
device(config)#snmp-server group ops v3 priv read internet write system
device(config)#snmp-server group admin v3 priv read internet write internet
device(config)#snmp-server group restricted v3 priv read internet
device(config)#snmp-server user ops ops v3 encrypted auth md5 ab8e9cd6d46e7a270b8c9549d92a069 priv
encrypted des 0e1b153303b6188089411447dbc32de
device(config)#snmp-server user admin admin v3 encrypted auth md5 0d8a2123f91bfbd8695fef16a6f4207b priv
encrypted des 18e0cf359fce4fcd60df19c2b6515448
device(config)#snmp-server user restricted restricted v3 encrypted auth md5
261fd8f56a3ad51c8bcecle4609f54dc priv encrypted des d32e66152f89de9b2e0cb17a65595f43
```

Configuring SNMP traps

This section explains how to do the following:

- Specify an SNMP trap receiver.
- Specify a source address and community string for all traps that the Extreme device sends.
- Change the holddown time for SNMP traps.
- Disable individual SNMP traps. (All traps are enabled by default.)
- Disable traps for CLI access that is authenticated by a local user account, a RADIUS server, or a TACACS or TACACS+ server.

Specifying an SNMP trap receiver

You can specify a trap receiver to ensure that all SNMP traps sent by the Extreme device go to the same SNMP trap receiver or set of receivers, typically one or more host devices on the network. When you specify the host, you also specify a community string. The Extreme device sends all the SNMP traps to the specified hosts and includes the specified community string. Administrators can therefore filter for traps from a Extreme device based on IP address or community string. The number of SNMP Trap receivers that can be configured is limited to 32.

If the string is in the clear format, the system will internally encrypt it. When you display or save the configuration, the encrypted string is used.

To specify an SNMP trap receiver, enter a command such as the following.

```
device(config)# snmp-server host 10.2.2.2 version v2c mypublic port 200
```

The command adds trap receiver 10.2.2.2 and designates the UDP port that will be used to receive traps.

Syntax: `[no] snmp-server host ip-addr version [v1 | v2c | v3] string [port value]`

The *ip-addr* parameter specifies the IP address of the trap receiver.

The v1, v2c, or v3 parameter indicates which version of SNMP is used.

The *string* parameter specifies an SNMP community string configured on the Extreme device. It is not used to authenticate access to the trap host, but it is a useful method for filtering traps on the host. For example, if you configure each of your Extreme devices that use the trap host to send a different community string, you can easily distinguish among the traps from the devices based on the community strings.

By default, *string* is encrypted. If you want *string* to be in clear text, insert a *O* preceding *string*.

```
device(config)# snmp-server host 10.2.2.2 version v2c O mypublic port 200
```

The software adds a prefix to the string in the configuration. For example, the following portion of the code has the encrypted code "2".

```
snmp-server host 10.2.2.2 version v2c 12
$Si2^=d
```

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 simple encryption (only for CES 2000 Series and CER 2000 Series)
- 2 = the key string uses base64 encryption format (only for XMR Series and MLX Series)

The **port** *value* parameter specifies the UDP port that will be used to receive traps. This parameter allows you to configure several trap receivers in a system. With this parameter, Extreme Network Advisor and another network management application can co-exist in the same system. The Extreme devices can be configured to send copies of traps to more than one network management application.

Specifying a single trap source

You can specify a single trap source to ensure that all SNMP traps sent by the Extreme device use the same source IP address. When you configure the SNMP source address, you can specify the Ethernet port, loopback interface, virtual routing interface, or management interface as the source for the traps. The Extreme device uses the lowest-numbered IP address configured on the port or interface as the source IP address in the SNMP traps it sends.

Identifying a single source IP address for SNMP traps provides the following benefits:

- If your trap receiver is configured to accept traps only from specific links or IP addresses, you can simplify configuration of the trap receiver by configuring the Extreme device to always send the traps from the same link or source address.
- If you specify a loopback interface as the single source for SNMP traps, SNMP trap receivers can receive traps regardless of the states of individual links. Thus, if a link to the trap receiver becomes unavailable but the receiver can be reached through another link, the receiver still receives the trap, and the trap still has the source IP address of the loopback interface.

To configure the Extreme device to send all SNMP traps from the first configured IP address on port 4/11, enter the following commands.

```
device(config)# snmp-server trap-source ethernet 4/11
device(config)# write memory
```

Syntax: [no] snmp-server trap-source { ethernet *slot/port* | loopback *num* | management *num* | ve *num* }

If you do not configure this command, the device will use the device router ID as the source IP address of the notification packet. The router ID of the device can be obtained from the "show ip" command output.

In the case when the SNMP trap comes from the IPv4 or IPv6 management interface, the management IP is used as SNMP trap source. By default, this occurs when the trap source does not come from the router ID. You do not need to configure the management interface or the port as a trap source using the **snmp-server trap-source** command.

From NetIron 06.1.00 release onwards, you can explicitly configure management interface as the SNMP trap source. If management interface is specified as the SNMP trap source, the outgoing SNMP trap notifications sent by the Extreme device use the management

IP address as the source IP address. If multiple IP addresses are configured as management IP address, the first IP address displayed in the output of the **show ip interfaces** command, which is the primary IP address, is used as the trap source.

NOTE

IPv6 address is not supported as SNMP trap source.

Following is a configuration example to specify a management interface as the device's SNMP trap source:

```
device(config)# snmp-server trap-source management 1
device(config)# write memory
```

To specify a loopback interface as the device's SNMP trap source, enter following commands.

```
device(config)# int loopback 1
device(config-lbif-1)# ip address 10.0.0.1/24
device(config-lbif-1)# exit
device(config)# snmp-server trap-source loopback 1
```

The commands configure loopback interface 1, gives it IP address 10.0.0.1/24, then designate it as the SNMP trap source for the Extreme device. Regardless of the port the Extreme uses to send traps to the receiver, the traps always arrive from the same source IP address.

Setting the SNMP trap holddown time

When a Extreme device starts up, the software waits for Layer 2 convergence (STP) and Layer 3 convergence (OSPF) before beginning to send SNMP traps to external SNMP servers. Until convergence occurs, the Extreme device might not be able to reach the servers, in which case the messages are lost.

By default, the Extreme device uses a one-minute holddown time to wait for the convergence to occur before starting to send SNMP traps. After the holddown time expires, the Extreme device sends the traps, including traps such as "cold start" or "warm start" that occur before the holddown time expires.

You can change the holddown time to a value from one second to ten minutes.

To change the holddown time for SNMP traps, enter a command such as the following at the global CONFIG level of the CLI.

```
device(config)# snmp-server enable traps holddown-time 30
```

The command changes the holddown time for SNMP traps to 30 seconds. The Extreme device waits 30 seconds to allow convergence in STP and OSPF before sending traps to the SNMP trap receiver.

Syntax: [no] snmp-server enable traps holddown-time secs

The secs parameter specifies the number of seconds (1 - 600). The default is 60.

Disabling SNMP traps

The Extreme device comes with SNMP trap generation enabled by default for all traps.

NOTE

By default, all SNMP traps are enabled at system startup.

You can selectively disable one or more of the following traps:

- SNMP authentication key
- Temperature
- Power supply failure

- Fan failure
- Cold start
- Link up
- Link down
- Bridge new root
- Bridge topology change
- Module insert
- Module remove
- Redundant module
- Metro-ring
- MPLS
- BGP4
- OSPF
- VRRP
- VSRP

To stop link down occurrences from being reported, enter the following command.

```
device(config)# no snmp-server enable traps link-down
```

Syntax: `[no] snmp-server enable traps trap-type`

A list of traps is available in the *Unified IP MIB Reference* .

Configuring SNMP management of VRFs

The SNMP agent can now support SNMP management of VRFs for multiple instances of routing protocol MIBs in addition to the default VRF. This section explains how to configure SNMP management for multiple instances of routing protocol MIBs.

SNMPv3 polling

For SNMPv3 polling, you can use the dedicated field used for identifying contexts to distinguish among multiple routing instances. SNMPv3 polling supports contexts using the `contextName` field in the SNMPv3 PDU. Use the following command to create an SNMP context and associate it with a routing instance (VRF).

```
device (Config)# snmp-server context context-name vrf vrf-name
```

SNMPv3 traps

The 'contextName' field in the SNMPv3 trap PDU contains the context name associated with the VRF for all SNMP traps originating from the routing instance. The SNMP manager application uses 'contextName' to distinguish between various VRF instances from which the trap originates. If there are no contexts configured for the VRF, the traps sent to the trap host will have null 'contextName'.

SNMPv1/v2c polling

For SNMPv1/v2c polling, you must map the community name to the context name. SNMP-COMMUNITY-MIB (RFC3584) is supported as part of this feature to help with the mapping and then enable SNMP-COMMUNITY-MIB support using the commands shown below.

```
device (Config)# snmp-server mib community-map community-name context context-name
device (Config)# snmp-server enable mib snmp-community-mib
```

SNMPv1/v2c traps

You can now configure 1 default trap community (community not mapped to any context) which is sent in 'communityName' field in the trap PDU for all traps generated from default VRFs and VRFs which do not have any context name configured and 1 trap community per context (community mapped to context name) which will be sent in 'communityName' field in the trap PDU for all traps generated from the VRF mapped to the context.

This allows the SNMP manager application to distinguish among various VRF instances of the trap even for SNMPv1/v2c traps.

The following example shows a typical configuration sequence for setting up multi-VRF support for SNMPv1/v2c.

```
//The following creates contexts for VRFs
device (Config)# snmp-server context ctxtA vrf VRFA
device (Config)# snmp-server context ctxtB vrf VRFB
//The following creates communities
device (Config)# snmp-server community comA ro
device (Config)# snmp-server community comB ro
device (Config)# snmp-server community comRest ro
//The following maps communities to contexts
device (Config)# snmp-server mib community-map comA context ctxtA
device (Config)# snmp-server mib community-map comB context ctxtB
//The following configures trap host with community names
device (Config)# snmp-server host 10.10.10.10 version v2c comA
device (Config)# snmp-server host 10.10.10.10 version v2c comB
device (Config)# snmp-server host 10.10.10.10 version v2c comRest
device (Config)# snmp-server host 20.20.20.20 version v3 noauth adminuser
```

The following command is enhanced to support multiple community names per host.

```
device (Config)# snmp-server host 20.20.20.20 version v1 community-map
```

Getting VRF information

You can now execute the following commands as shown below to get the OSPF area IDs associated with the VRFs and the information on the device.

```
device $snmp-server -v2c -c comA 10.37.73.178 ospfAreaId
OSPF-MIB::ospfAreaId.0.0.0.1 = IPAddress: 0.0.0.1
OSPF-MIB::ospfAreaId.12.12.12.12 = IPAddress: 12.12.12.12
```

```
device $snmp-server -v2c -c comB 10.37.73.178 ospfAreaId
OSPF-MIB::ospfAreaId.5.5.5.5 = IPAddress: 5.5.5.5
OSPF-MIB::ospfAreaId.6.6.6.6 = IPAddress: 6.6.6.6
```

```
device # ip ospf vrf VRFA area
Number of Areas is 2
Indx Area          Type    Cost      SPFR      ABR      ASBR     LSA      Chksum      Translator
1      1              normal  0          16        0        0        1          0x0000a76e  --
2      12.12.12.12    normal  0          16        0        0        1          0x00009d77  --
```

```
MLX#sh ip ospf vrf VRFB area
Number of Areas is 2
Indx Area          Type    Cost      SPFR      ABR      ASBR     LSA      Chksum      Translator
1      5.5.5.5          normal  0          4         0        0        0          0x00000000  --
2      6.6.6.6          normal  0          3         0        0        0          0x00000000  --
```

For a complete list of MIBs supported for SNMP VRFs, refer to the *Unified IP MIB Reference*.

Configuring SNMP ifIndex

This section explains how ifIndex values are assigned on Extreme devices.

On Extreme NetIron CES and CER only

On the CES 2000 Series and CER 2000 Series, the system automatically assign 64 indexes to each module on the device. This value is not configurable.

On Extreme NetIron XMR and MLX Series only

On XMR Series and MLX Series devices, SNMP Management Information Base (MIB) uses Interface Index (ifIndex) to assign a unique value to each port on a module or slot. The number of indexes that can be assigned per module is 20, 40, or 64, depending on the number of ports on the module.

Enter the following to change the number of indexes per module.

```
device(config)# snmp-server max-ifindex-per-module 40
```

Syntax: [no] snmp-server max-ifindex-per-module [20 | 40 | 64]

20 is the default.

You cannot change the maximum ifIndex per module to a number less than the number of ports.

After this command is issued the following are generated:

- "System: Ifindex assignment was changed" is logged in the Syslog.
- The snTrapIfIndexAssignmentChanged trap is generated.

Configuration notes for Extreme NetIron XMR and MLX Series

Note the following if you are upgrading the software on the XMR Series and MLX Series:

- If you are running an earlier version of the software and you will not be installing the NI-MLX-1Gx48-T module, you do not need to change your ifIndex allocation scheme. The current definition is maintained. The maximum ifIndex per module can remain at 20 or 40.
- If you are running an earlier version of the software and you will be installing the NI-MLX-1Gx48-T module on your MLX Series, you must configure the maximum ifIndex per module to 64. **You must change the ifIndex allocation before installing the NI-MLX-1Gx48-T module** ; otherwise, the module status remains in the Offline state.
- If you have a new MLX Series (no previous software installed), but will not be installing an NI-MLX-1Gx48-T module, it is recommended that you configure the maximum ifIndex per module to 64 to avoid future ifIndex problems in case an NI-MLX-1Gx48-T module is installed in the future.
- If you have a new MLX Series (no previous software installed), and you will be installing an NI-MLX-1Gx48-T module, you **must** configure the maximum ifIndex per module to 64; otherwise, the module remains in the Offline state.

SNMP scalability optimization

To ensure that SNMP requests are responded to promptly and that SNMP loads do not impact other device activities, the Extreme device speeds SNMP tasks and limit their effects on the CPU by a combination of throughput optimization and load throttling.

Configuring SNMP throughput optimization

SNMP throughput is optimized on the Extreme device through a combination of SNMP value caching, conditional yielding by the SNMP agent, and acceptance of incoming packets during queue processing.

SNMP agent yielding behavior

When an SNMP agent yields CPU control unconditionally between processing of queued packets, it can result in low throughput for packets which are processed quickly. To increase throughput for these packets, the SNMP agent in the Extreme device yields CPU control between packets only when the agent has controlled the CPU for more than 10 milliseconds.

SNMP queue processing

To ensure that SNMP packets are not dropped, the SNMP task on the Extreme device continues to accept newly received SNMP packets from the IP stack while processing the SNMP queue.

Configuring SNMP load throttling

To ensure that high SNMP loads do not interfere with the performance of the device, the Extreme device limits the percentage of CPU time that can be occupied by SNMP processing. This limit is not imposed when the CPU is idle.

NOTE

This command tries to fix the maximum percentage of time SNMP task can run in a non-idle system environment. This implies that SNMP task can't run for more than the specified percentage of time if the system is having zero idle time. But this constraint is checked only between processing of 2 SNMP PDU's. If the processing of a single SNMP PDU takes longer time then we may overrun the maximum limit. This command also tries to fix the minimum percentage of time SNMP task can run in a non-idle system environment. But if there is another task which is continuously hogging the CPU and SNMP is not getting time to run then we may under run the specified limit.

To configure the maximum percentage of CPU time that can be used by SNMP processing, use the following command at the configuration level of the CLI.

```
device(config)#snmp-server cpu max-non-idle-utilization 25
```

Syntax: `[no] snmp-server cpu max-non-idle-utilization percent`

- The preceding example raises the maximum percentage of non-idle CPU time to be used by SNMP processing to 25%.
- The *percent* parameter is the maximum percentage of non-idle CPU time to be used by SNMP processing. The range for this parameter is from 1 through 25.
- Use the *no* form of this command to return the SNMP non-idle CPU time maximum to the default value of 10%.

Configuring SNMP to revert ifType to legacy values

The ifType for all Ethernet interfaces (10/100/1G/10G) returns the value ethernetCsmacd(6) as mandated by RFC 2665. If you want ifType to return gigabitEthernet (117) or fastEther(62) for Ethernet interfaces, enter the following command.

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

Syntax: `[no] snmp-server legacy iftype`

When this command is configured, the values gigabitEthernet (117) or fastEther(62) are returned for ifType. If you issue a **no snmp-server legacy iftype**, ifType returns ethernetCsmacd(6) for Ethernet interfaces.

Configuring snAgentConfigModuleType to return original values

Enumeration values for snAgentConfigModuleType object in the SNMP MIB have been changed in Release 04.0.00 for the XMR Series and MLX Series to resolve enumeration conflicts with other hardware modules in the Unified IP MIB. For example, an SNMP get of the snAgentConfigModuleType of the 10x1GC module returned xmr20PortGigCopperSPModule(84). Beginning with Release 04.0.00, snAgentConfigModuleType returns fdryXmr20PortGigCopperSPModule(1084) for the 10x1GC module.

If you want snAgentConfigModuleType to return the enumeration values used before Release 04.0.00, configure the following command.

```
device(config)# snmp-server legacy module-type
```

Syntax: `[no] snmp-server legacy module-type`

Refer to the *Unified IP MIB Reference* for details on snAgentConfigModuleType.

Preserving interface statistics in SNMP

By default, statistics for an interface is cleared from both the CLI and SNMP when the following commands are entered on the CLI:

- **clear statistics ethernet** *slot-number/port-number*
- **clear statistics** *slot-number/port-number*
- **clear rmon statistics**
- **clear statistics log** *slot-number/port-number*

If you want to preserve interface statistics in SNMP when these commands are entered, configure the following command at the Global level of the CLI.

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

Syntax: [no] **snmp-server preserve-statistics**

For details on which interface statistics are preserved in SNMP, refer to the "Preserved interface statistics for SNMP" section of the "Supported Standard MIBs" chapter in the *Unified IP MIB Reference* .

NOTE

Statistics for an interface will be different between the CLI and SNMP if **snmp-server preserve-statistics** is configured and the clear commands listed above are executed.

NIAP-CCEVS

- NIAP-CCEVS-certified Extreme equipment and IronWare releases..... 229
- Web management access to NIAP-CCEVS-certified Extreme equipment.....230
- Warning: local user password changes.....230

NIAP-CCEVS-certified Extreme equipment and IronWare releases

Some devices have passed the Common Criteria (CC) certification testing. This testing is sponsored by the National Information Assurance Partnership (NIAP) - Common Criteria Evaluation and Validation Scheme (CCEVS). For more information regarding the NIAP-CCEVS certification process refer to the following link: <http://www.nap-ccevs.org/>

In an effort to maintain a proper level of security as it relates to access to network infrastructure resources, Extreme recommends that all Extreme hardware be installed within a secure location that is accessible by approved personnel only.

The following devices have been NIAP-CCEVS certified. The following IronWare software release must be used to remain compliant with this certification:

TABLE 53 NIAP-CCEVS certified equipment and IronWare software releases

Extreme product	Extreme IronWare software version	Discussed in
Extreme XMR Family	3.8.00a	<i>Extreme NetIron Administration Guide</i>
Extreme MLX Family	3.8.00a	<i>Extreme NetIron Administration Guide</i>
BigIron RX Family	2.5.00b	<i>BigIron RX Series Configuration Guide</i>
FastIron SuperX/SX Family	4.1.00	<i>FastIron and Turbolron Configuration Guide</i>
FastIron Edge X Family	4.1.00	<i>FastIron and Turbolron Configuration Guide</i>
FastIron GS/LS Family	4.2.00a	<i>FastIron and Turbolron Configuration Guide</i>
FastIron Edge Switch Family	4.0.00a	<i>FastIron Security Guide</i>
ServerIron JetCore Family	11.0.00a	<i>ServerIron TrafficWorks Graphical User Interface</i> <i>ServerIron TrafficWorks Server Load Balancing Guide</i> <i>ServerIron TrafficWorks Advanced Server Load Balancing Guide</i> <i>ServerIron TrafficWorks Global Server Load Balancing Guide</i> <i>ServerIron TrafficWorks Security Guide</i> <i>ServerIron TrafficWorks Administration Guide</i> <i>ServerIron TrafficWorks Switching and Routing Guide</i> <i>ServerIron Firewall Load Balancing Guide</i>
ServerIron ADX Family	12.0.00	<i>ServerIron ADX TrafficWorks Graphical User Interface</i> <i>ServerIron ADX TrafficWorks Server Load Balancing Guide</i> <i>ServerIron ADX TrafficWorks Advanced Server Load Balancing Guide</i> <i>ServerIron ADX TrafficWorks Global Server Load Balancing Guide</i> <i>ServerIron ADX TrafficWorks Security Guide</i> <i>ServerIron ADX TrafficWorks Administration Guide</i>

TABLE 53 NIAP-CCEVS certified equipment and IronWare software releases (continued)

Extreme product	Extreme IronWare software version	Discussed in
		ServerIron ADX TrafficWorks Switching and Routing Guide ServerIron ADX Firewall Load Balancing Guide

Web management access to NIAP-CCEVS-certified Extreme equipment

All devices that are to remain in compliancy with the NIAP-CCEVS certification must disable all remote access through the integrated Web management graphical user interface (GUI). In accordance with NIAP-CCEVS this functionality is considered a security risk and must be disabled.

Refer to the Extreme Configuration Guides associated with each product in [NIAP-CCEVS-certified Extreme equipment and IronWare releases](#) on page 229 for detailed instructions on how to disable the Web Management Interface feature.

Warning: local user password changes

Please note that if existing usernames and passwords have been configured on a device with specific privilege levels (super-user, read-only, port-config) and if you attempt to change a user's password by executing the following command.

```
device(config)# user fdryreadonly password <value>
```

The privilege level of this particular user will be changed from its current value to "super-user". The "super-user" level username and password combination provides full access to the Extreme command line interface (CLI). To prevent this from occurring, use the following command.

```
device(config)# user fdryreadonly privilege <value> password <value>
```