

RUCKUS FastIron Monitoring Configuration Guide, 09.0.10

Supporting FastIron Software Release 09.0.10

Copyright, Trademark and Proprietary Rights Information

© 2021 CommScope, Inc. All rights reserved.

No part of this content may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without written permission from CommScope, Inc. and/or its affiliates ("CommScope"). CommScope reserves the right to revise or change this content from time to time without obligation on the part of CommScope to provide notification of such revision or change.

Export Restrictions

These products and associated technical data (in print or electronic form) may be subject to export control laws of the United States of America. It is your responsibility to determine the applicable regulations and to comply with them. The following notice is applicable for all products or technology subject to export control:

These items are controlled by the U.S. Government and authorized for export only to the country of ultimate destination for use by the ultimate consignee or end-user(s) herein identified. They may not be resold, transferred, or otherwise disposed of, to any other country or to any person other than the authorized ultimate consignee or end-user(s), either in their original form or after being incorporated into other items, without first obtaining approval from the U.S. government or as otherwise authorized by U.S. law and regulations.

Disclaimer

THIS CONTENT AND ASSOCIATED PRODUCTS OR SERVICES ("MATERIALS"), ARE PROVIDED "AS IS" AND WITHOUT WARRANTIES OF ANY KIND, WHETHER EXPRESS OR IMPLIED. TO THE FULLEST EXTENT PERMISSIBLE PURSUANT TO APPLICABLE LAW, COMMSCOPE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, TITLE, NON-INFRINGEMENT, FREEDOM FROM COMPUTER VIRUS, AND WARRANTIES ARISING FROM COURSE OF DEALING OR COURSE OF PERFORMANCE. CommScope does not represent or warrant that the functions described or contained in the Materials will be uninterrupted or error-free, that defects will be corrected, or are free of viruses or other harmful components. CommScope does not make any warranties or representations regarding the use of the Materials in terms of their completeness, correctness, accuracy, adequacy, usefulness, timeliness, reliability or otherwise. As a condition of your use of the Materials, you warrant to CommScope that you will not make use thereof for any purpose that is unlawful or prohibited by their associated terms of use.

Limitation of Liability

IN NO EVENT SHALL COMMSCOPE, COMMSCOPE AFFILIATES, OR THEIR OFFICERS, DIRECTORS, EMPLOYEES, AGENTS, SUPPLIERS, LICENSORS AND THIRD PARTY PARTNERS, BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER, EVEN IF COMMSCOPE HAS BEEN PREVIOUSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, WHETHER IN AN ACTION UNDER CONTRACT, TORT, OR ANY OTHER THEORY ARISING FROM YOUR ACCESS TO, OR USE OF, THE MATERIALS. Because some jurisdictions do not allow limitations on how long an implied warranty lasts, or the exclusion or limitation of liability for consequential or incidental damages, some of the above limitations may not apply to you.

Trademarks

ARRIS, the ARRIS logo, COMMSCOPE, RUCKUS, RUCKUS WIRELESS, the Ruckus logo, the Big Dog design, BEAMFLEX, CHANNELFLY, FASTIRON, ICX, SMARTCELL and UNLEASHED are trademarks of CommScope, Inc. and/or its affiliates. Wi-Fi Alliance, Wi-Fi, the Wi-Fi logo, Wi-Fi Certified, the Wi-Fi CERTIFIED logo, Wi-Fi Protected Access, the Wi-Fi Protected Setup logo, Wi-Fi Protected Setup, Wi-Fi Multimedia and WPA2 and WMM are trademarks or registered trademarks of Wi-Fi Alliance. All other trademarks are the property of their respective owners.

Contents

Preface	7
Contacting RUCKUS Customer Services and Support.....	7
What Support Do I Need?.....	7
Open a Case.....	7
Self-Service Resources.....	8
Document Feedback.....	8
RUCKUS Product Documentation Resources.....	8
Online Training Resources.....	8
Document Conventions.....	9
Notes, Cautions, and Safety Warnings.....	9
Command Syntax Conventions.....	9
About This Document	11
Supported Hardware.....	11
What's new in this document	11
How Command Information is Presented in this Configuration Guide.....	11
Operations, Administration, and Maintenance	13
OAM Overview.....	13
Software Versions Installed and Running on a Device.....	14
Determining the Flash Image Version Running on the Device.....	14
Displaying the Image Versions Installed in Flash Memory.....	15
Flash Image Verification	16
Verifying the Flash Image.....	16
Software Image File Types.....	17
Flash Timeout.....	18
Software Upgrades.....	18
Boot Code Synchronization.....	18
Viewing the Contents of Flash Files.....	18
Using SNMP to Upgrade Software.....	19
Software Reboot.....	19
Displaying the Boot Preference.....	20
System Reload Scheduling.....	22
Diagnostic Error Codes and Remedies for TFTP Transfers.....	23
Network Connectivity Testing.....	24
IEEE 802.3ah EFM-OAM.....	25
Network Deployment Use Case.....	25
EFM-OAM Protocol.....	25
Process Overview.....	26
Remote Failure Indication.....	27
Enabling Battleshort Mode.....	27
Remote Loopback.....	28
EFM-OAM Error Disable Recovery	28
Configuring EFM-OAM.....	28
Displaying OAM Information.....	30
Displaying OAM Statistics.....	31
EFM-OAM Syslog Messages.....	33

Displaying Management Redundancy Information	34
Layer 3 Hitless Route Purge	34
Setting the IPv4 Hitless Purge Timer.....	34
Setting the IPv6 Hitless Purge Timer.....	35
Energy Efficient Ethernet.....	35
Port Support for Energy Efficient Ethernet.....	36
Enabling Energy Efficient Ethernet.....	37
Histogram Information Overview.....	37
Displaying CPU Histogram Information.....	38
External USB Hotplug.....	38
External USB Hotplug Considerations.....	38
Using External USB Hotplug.....	38
Basic System Management.....	39
Viewing System Information.....	39
Viewing Configuration Information.....	41
Enabling the Display of the Elapsed Timestamp for Port Statistics Reset.....	41
Viewing port statistics.....	42
Viewing STP Statistics.....	43
Clearing Statistics.....	44
Viewing Egress Queue Counters.....	44
Clearing the Egress Queue Counters.....	45
Collecting CPU Packet Statistics.....	45
Link Fault Signaling for 10Gbps Ethernet Devices.....	46
Enabling Link Fault Signaling.....	47
Viewing the Status of LFS-enabled Links.....	47
Locating a Device Using Port LEDs.....	48
Blink LED.....	48
Limitations.....	48
Hardware Component Monitoring.....	51
Virtual Cable Testing.....	51
VCT Configuration Notes.....	51
VCT Restrictions.....	51
Crosstalk Between Ports	52
Diagnosing a Cable using Time Domain Reflectometry.....	53
Viewing the Results of the Cable Analysis.....	53
Digital Optical Monitoring.....	54
DOM Show and Configuration Commands.....	55
Enabling DOM.....	56
DOM Configuration Example.....	57
Syslog Messages for Optical Transceivers.....	58
Port Mirroring and Monitoring.....	61
Port Mirroring and Monitoring Overview.....	61
Port Mirroring and Monitoring Configuration.....	61
Configuration Notes for Port Mirroring and Monitoring.....	62
Commands for Port Mirroring and Monitoring.....	62
Mirroring Configuration on a Traditional Stack.....	63
Configuration Notes for Traditional Stack Mirroring.....	64
Configuring Mirroring for Ports on Different Members in a Traditional Stack Example.....	64
Configuring Mirroring for Ports on the Same Stack Member in a Traditional Stack Example.....	64

ACL-based Inbound Mirroring.....	64
Creating an ACL-based Inbound Mirror Clause	65
Destination mirror port	65
MAC ACL Mirroring.....	67
MAC ACL Configuration Notes.....	67
Configuring MAC ACL Mirroring.....	68
VLAN-based Mirroring.....	68
Configuration Notes for VLAN-based Mirroring.....	68
Configuring VLAN-based Mirroring.....	68
Displaying VLAN-based Mirroring Status.....	69
Remote Switched Port Analyzer.....	69
RSPAN Feature Limitations and Considerations.....	70
Configuring RSPAN.....	71
Encapsulated Remote Switched Port Analyzer (ERSPAN)	73
ERSPAN Configuration Steps.....	74
ERSPAN feature limitations.....	74
Configuring an ERSPAN Profile.....	74
Configuring a Monitor Port for ERSPAN.....	79
sFlow.....	81
sFlow Overview.....	81
sFlow Version 5.....	81
sFlow Support for IPv6 Packets.....	82
sFlow Configuration Considerations.....	82
Configuring sFlow.....	85
sFlow Version 5 Feature Configuration.....	88
Specifying Maximum Packet Size Values.....	88
Configuring sFlow Version 5.....	89
Configuring sFlow with Multi-VRF.....	90
Displaying sFlow Information.....	91
Clearing sFlow Statistics.....	91
HMON - Health Monitor Service.....	93
HMON overview.....	93
HMON process registration.....	93
Dynamic start and stop.....	93
Process availability based on stack role.....	93
Clients marked as faulty.....	93
Critical processes.....	94
Determining the administrative and operational state of an HMON client.....	94
Dependent process restart.....	94
Troubleshooting HMON.....	95
Syslog.....	101
Syslog Messages.....	101
Enabling Real-Time Display of Syslog Messages.....	102
Disabling Real-Time Display of Syslog Messages.....	102
Broadcast, Unknown Unicast, and Multicast Suppression Syslog and SNMP notification.....	103
BUM Suppression Restrictions and Limitations.....	103
Enabling BUM Suppression Logging.....	103
Viewing BUM Suppression Syslog Notifications.....	104
Displaying Syslog Messages.....	104

Displaying Real-Time Syslog messages	105
Syslog Service Configuration.....	105
Static and Dynamic Buffers.....	106
Clearing Log Entries.....	106
Timestamps.....	106
Configuring Syslog Service.....	107
Disabling or Re-enabling Syslog.....	109
Displaying the Syslog Configuration.....	109
Generating the Syslog Specific to RFC 5424.....	109
Syslog Messages.....	113
Syslog Message Descriptions.....	113
Syslog messages IPsec and IKEv2.....	143
Syslog messages system.....	145

Preface

• Contacting RUCKUS Customer Services and Support.....	7
• Document Feedback.....	8
• RUCKUS Product Documentation Resources.....	8
• Online Training Resources.....	8
• Document Conventions.....	9
• Command Syntax Conventions.....	9

Contacting RUCKUS Customer Services and Support

The Customer Services and Support (CSS) organization is available to provide assistance to customers with active warranties on their RUCKUS products, and customers and partners with active support contracts.

For product support information and details on contacting the Support Team, go directly to the RUCKUS Support Portal using <https://support.ruckuswireless.com>, or go to <https://www.ruckuswireless.com> and select **Support**.

What Support Do I Need?

Technical issues are usually described in terms of priority (or severity). To determine if you need to call and open a case or access the self-service resources, use the following criteria:

- Priority 1 (P1)—Critical. Network or service is down and business is impacted. No known workaround. Go to the **Open a Case** section.
- Priority 2 (P2)—High. Network or service is impacted, but not down. Business impact may be high. Workaround may be available. Go to the **Open a Case** section.
- Priority 3 (P3)—Medium. Network or service is moderately impacted, but most business remains functional. Go to the **Self-Service Resources** section.
- Priority 4 (P4)—Low. Requests for information, product documentation, or product enhancements. Go to the **Self-Service Resources** section.

Open a Case

When your entire network is down (P1), or severely impacted (P2), call the appropriate telephone number listed below to get help:

- Continental United States: 1-855-782-5871
- Canada: 1-855-782-5871
- Europe, Middle East, Africa, Central and South America, and Asia Pacific, toll-free numbers are available at <https://support.ruckuswireless.com/contact-us> and Live Chat is also available.
- Worldwide toll number for our support organization. Phone charges will apply: +1-650-265-0903

We suggest that you keep a physical note of the appropriate support number in case you have an entire network outage.

Self-Service Resources

The RUCKUS Support Portal at <https://support.ruckuswireless.com> offers a number of tools to help you to research and resolve problems with your RUCKUS products, including:

- Technical Documentation—<https://support.ruckuswireless.com/documents>
- Community Forums—<https://forums.ruckuswireless.com/ruckuswireless/categories>
- Knowledge Base Articles—<https://support.ruckuswireless.com/answers>
- Software Downloads and Release Notes—https://support.ruckuswireless.com/#products_grid
- Security Bulletins—<https://support.ruckuswireless.com/security>

Using these resources will help you to resolve some issues, and will provide TAC with additional data from your troubleshooting analysis if you still require assistance through a support case or RMA. If you still require help, open and manage your case at https://support.ruckuswireless.com/case_management.

Document Feedback

RUCKUS is interested in improving its documentation and welcomes your comments and suggestions.

You can email your comments to RUCKUS at #Ruckus-Docs@commscope.com.

When contacting us, include the following information:

- Document title and release number
- Document part number (on the cover page)
- Page number (if appropriate)

For example:

- RUCKUS SmartZone Upgrade Guide, Release 5.0
- Part number: 800-71850-001 Rev A
- Page 7

RUCKUS Product Documentation Resources

Visit the RUCKUS website to locate related documentation for your product and additional RUCKUS resources.

Release Notes and other user documentation are available at <https://support.ruckuswireless.com/documents>. You can locate the documentation by product or perform a text search. Access to Release Notes requires an active support contract and a RUCKUS Support Portal user account. Other technical documentation content is available without logging in to the RUCKUS Support Portal.

White papers, data sheets, and other product documentation are available at <https://www.ruckuswireless.com>.

Online Training Resources

To access a variety of online RUCKUS training modules, including free introductory courses to wireless networking essentials, site surveys, and products, visit the RUCKUS Training Portal at <https://training.ruckuswireless.com>.

Document Conventions

The following table lists the text conventions that are used throughout this guide.

TABLE 1 Text Conventions

Convention	Description	Example
monospace	Identifies command syntax examples	<code>device(config)# interface ethernet 1/1/6</code>
bold	User interface (UI) components such as screen or page names, keyboard keys, software buttons, and field names	On the Start menu, click All Programs .
<i>italics</i>	Publication titles	Refer to the <i>RUCKUS Small Cell Release Notes</i> for more information.

Notes, Cautions, and Safety 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 some information that you must read before continuing with the current action or task.



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.

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. Default responses to system prompts are enclosed in square brackets.
{x y z}	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <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.

About This Document

- Supported Hardware..... 11
- What's new in this document 11
- How Command Information is Presented in this Configuration Guide..... 11

Supported Hardware

This guide supports the following RUCKUS products:

- RUCKUS ICX 7850 Switch
- RUCKUS ICX 7650 Switch
- RUCKUS ICX 7550 Switch
- RUCKUS ICX 7450 Switch
- RUCKUS ICX 7250 Switch
- RUCKUS ICX 7150 Switch

For information about what models and modules these devices support, refer to the hardware installation guide for the specific product family.

What's new in this document

The following tables describe information added or modified in this guide for FastIron software release 09.0.10.

TABLE 2 Summary of enhancements in FastIron release 09.0.10

Feature	Description	Described in
Syslog threshold	Configure a threshold value to generate a warning message when log entries exceed the specified threshold.	Configuring Syslog Service on page 107
Updates to address defects	Minor updates on content throughout to address defects.	All chapters.
Minor editorial updates	Minor editorial updates were made throughout the Configuration Guide.	All chapters.

How Command Information is Presented in this Configuration Guide

For all new content supported in FastIron release 08.0.20 and later, command information is documented in a standalone command reference guide.

In the *RUCKUS FastIron Command Reference*, the command pages are in alphabetical order and follow a standard format to present syntax, parameters, mode, usage guidelines, examples, and command history.

NOTE

Many commands introduced before FastIron release 08.0.20 are also included in the guide.

Operations, Administration, and Maintenance

• OAM Overview.....	13
• Software Versions Installed and Running on a Device.....	14
• Software Image File Types.....	17
• Flash Timeout.....	18
• Software Upgrades.....	18
• Boot Code Synchronization.....	18
• Viewing the Contents of Flash Files.....	18
• Using SNMP to Upgrade Software.....	19
• Software Reboot.....	19
• Displaying the Boot Preference.....	20
• System Reload Scheduling.....	22
• Diagnostic Error Codes and Remedies for TFTP Transfers.....	23
• Network Connectivity Testing.....	24
• IEEE 802.3ah EFM-OAM.....	25
• Displaying Management Redundancy Information	34
• Layer 3 Hitless Route Purge	34
• Energy Efficient Ethernet.....	35
• Histogram Information Overview.....	37
• External USB Hotplug.....	38
• Basic System Management.....	39
• Collecting CPU Packet Statistics.....	45
• Link Fault Signaling for 10Gbps Ethernet Devices.....	46
• Locating a Device Using Port LEDs.....	48

OAM Overview

For easy software image management, all RUCKUS devices support the download and upload of software images between the flash modules on the devices and a Trivial File Transfer Protocol (TFTP) server on the network.

RUCKUS devices have two flash memory modules:

- Primary flash - The default local storage device for image files and configuration files.
- Secondary flash - A second flash storage device. You can use the secondary flash to store redundant images for additional booting reliability or to preserve one software image while testing another one.

Only one flash device is active at a time. By default, the primary image will become active upon reload.

You can update the software contained on a flash module using TFTP to copy the update image from a TFTP server onto the flash module. In addition, you can copy software images and configuration files from a flash module to a TFTP server.

NOTE

RUCKUS devices are TFTP clients but not TFTP servers. You must perform the TFTP transaction from the RUCKUS device. You cannot "put" a file onto the RUCKUS device using the interface of your TFTP server.

NOTE

If you are attempting to transfer a file using TFTP but have received an error message, refer to [Diagnostic Error Codes and Remedies for TFTP Transfers](#) on page 23.

Software Versions Installed and Running on a Device

Use the following methods to display the software versions running on the device and the versions installed in flash memory.

Determining the Flash Image Version Running on the Device

You can determine the flash image version running on a device.

The following example shows the flash image version running on a Compact device and provides the following version information.

- "SW: Version 08.0.40q017T213" indicates the flash code version number.
- "labeled as SPR08040q017" indicates the flash code image label. The label indicates the image type and version and is especially useful if you change the image file name.
- "Secondary SW: Version 08.0.40q017T213" indicates the flash code image file name that was loaded.

```
device> show version
Copyright (c) 1996-2015 Ruckus Networks. All rights reserved.
UNIT 1: compiled on Aug 31 2015 at 04:56:36 labeled as SPR08040q017
(24061724 bytes) from Secondary
  SW: Version 08.0.40q017T213
Compressed Boot-Monitor Image size = 786944, Version:10.1.05T215 (spz10105b008)
  Compiled on Thu Jul 16 06:27:06 2015

HW: Stackable ICX7450-24
Internal USB: Serial #: 9900614090900038
  Vendor: ATP Electronics, Total size = 1919 MB
=====
UNIT 1: SL 1: ICX7450-24 24-port Management Module
  Serial #:CYT3346K035
  License: ICX7450_L3_SOFT_PACKAGE (LID: eavIIJLmFIK)
  License Compliance: ICX7450-PREM-LIC-SW is Compliant
  P-ASIC 0: type B548, rev 01 Chip BCM56548_A0
=====
UNIT 1: SL 2: ICX7400-4X10GF 4-port 40G Module
  Serial #:CYV3346K07G
=====
UNIT 1: SL 3: ICX7400-1X40GQ 1-port 40G Module
  Serial #:CYX3346K06F
=====
UNIT 1: SL 4: ICX7400-1X40GQ 1-port 40G Module
  Serial #:CYX3346K00A
=====
1000 MHz ARM processor ARMv7 88 MHz bus
8192 KB boot flash memory
2048 MB code flash memory
2048 MB DRAM
STACKID 1 system uptime is 6 day(s) 5 hour(s) 36 minute(s) 29 second(s)
The system : started=cold start
```

The following example shows the flash image version running on a chassis device and provides the following version information.

- "03.1.00aT3e3" indicates the flash code version number. The "T3e3" is used by RUCKUS for record keeping.
- "labeled as SXR03100a" indicates the flash code image label. The label indicates the image type and version and is especially useful if you change the image file name.
- "Primary SXR03100a.bin" indicates the flash code image file name that was loaded.

```
device> show version
=====
Active Management CPU [Slot-9]:
  SW: Version 07.4.00T3e3 Copyright (c) 1996-2012 Ruckus Networks. All rights reserved.
      Compiled on Mar 02 2012 at 11:54:29 labeled as SXR07400
      (4585331 bytes) Primary /GA/SXR07400.bin
  BootROM: Version 07.2.00T3e5 (FEv2)
  Chassis Serial #: Bxxxxxxxxx
      License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE (LID: yGFJGOiFLd)
  HW: Chassis FastIron SX 800-PREM6 (PROM-TYPE SX-FIL3U-6-IPV6)
=====
Standby Management CPU [Slot-10]:
  SW: Version 07.4.00T3e3 Copyright (c) 1996-2012 Ruckus Networks. All rights reserved.
      Compiled on Mar 02 2012 at 11:54:29 labeled as SXR07400
      BootROM: Version 07.2.00T3e5 (FEv2)
  HW: Chassis FastIron SX 800-PREM6 (PROM-TYPE SX-FIL3U-6-IPV6)
=====
SL 1: SX-FI-8XG 8-port 10G Fiber
      Serial #: BQKxxxxxxxxx
      P-ASIC 0: type C341, rev 00 subrev 00
=====
SL 2: SX-FI-24GPP 24-port Gig Copper + PoE+
      Serial #: BTUxxxxxxxxx
      P-ASIC 2: type C300, rev 00 subrev 00
=====
SL 8: SX-FI-48GPP 48-port Gig Copper + PoE+
      Serial #: BFVxxxxxxxxx
      P-ASIC 14: type C300, rev 00 subrev 00
=====
SL 9: SX-FIZMR6 0-port Management
      Serial #: Wxxxxxxxxx
      License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE (LID: yGFJGOiFLd)
=====
SL 10: SX-FIZMR6 0-port Management
      Serial #: Wxxxxxxxxx
      License: SX_V6_HW_ROUTER_IPv6_SOFT_PACKAGE (LID: yyyyyyyy)
=====
Active Management Module:
  660 MHz Power PC processor 8541 (version 0020/0020) 66 MHz bus
  512 KB boot flash memory
  16384 KB code flash memory
  512 MB DRAM
Standby Management Module:
  660 MHz Power PC processor 8541 (version 0020/0020) 66 MHz bus
  512 KB boot flash memory
  16384 KB code flash memory
  512 MB DRAM
The system uptime is 1 minutes 2 seconds
The system : started=warm start   reloaded=by "reload"
```

Displaying the Image Versions Installed in Flash Memory

You can view the boot and flash images installed on the device, as well as the boot image running on the device, as outlined in the following example.

```
device> show flash
Active Management Module (Slot 9):
Compressed Pri Code size = 3613675, Version 03.1.00aT3e3 (sxr03100a.bin)
Compressed Sec Code size = 2250218, Version 03.1.00aT3e1 (sxs03100a.bin)
```

Operations, Administration, and Maintenance

Software Versions Installed and Running on a Device

```
Compressed BootROM Code size = 524288, Version 03.0.01T3e5
Code Flash Free Space = 9699328
Standby Management Module (Slot 10):
Compressed Pri Code size = 3613675, Version 03.1.00aT3e3 (sxr03100a.bin)
Compressed Sec Code size = 2250218, Version 03.1.00aT3e1 (sxs03100a.bin)
Compressed BootROM Code size = 524288, Version 03.0.01T3e5
Code Flash Free Space = 524288
```

The following information is displayed in the output:

- The "Compressed Pri Code size" line lists the flash code version installed in the primary flash area.
- The "Compressed Sec Code size" line lists the flash code version installed in the secondary flash area.
- The "Boot Monitor Image size" line lists the boot code version installed in flash memory. The device does not have separate primary and secondary flash areas for the boot image. The flash memory module contains only one boot image.

NOTE

To minimize the boot-monitor image size on FastIron devices, the **ping** and **tftp** operations performed in the boot-monitor mode are restricted to copper ports on the FastIron Chassis management modules and to the out-of-band management port on the FastIron stackable switches. The other copper or fiber ports on these devices do not have the ability to ping or tftp from the boot-monitor mode.

Refer to the *RUCKUS FastIron Command Reference* for more information.

Flash Image Verification

The Flash Image Verification feature allows you to verify boot images based on hash codes, and to generate hash codes where needed. This feature lets you select from three data integrity verification algorithms:

- MD5 - Message Digest algorithm (RFC 1321)
- **SHA1** - US Secure Hash Algorithm (RFC 3174)
- CRC - Cyclic Redundancy Checksum algorithm

Verifying the Flash Image

You can use various **verify** commands for the verification of boot images based on hash codes and the generation of hash codes where needed.

Use one of the following commands to verify the boot image. The commands do not need to be entered in the specified order. Using these commands is optional and they can be entered in any order. For more information on these commands, refer to the *RUCKUS FastIron Command Reference*.

1. Enter the **verify** command with the **md5** and **secondary** keywords to generate an MD5 hash value for the secondary imaged.

```
device# verify md5 secondary
device#.....Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653862
```

2. Enter the **verify** command with the **sha** and **secondary** keywords to generate a SHA-1 hash value for the secondary image.

```
device# verify sha secondary
device#.....Done
Size = 2044830, SHA1 49d12d26552072337f7f5fcaef4cf4b742a9f525
```

3. Enter the **verify** command with the **crc32** and **secondary** keywords to generate a CRC32 hash value for the secondary image.

```
device# verify crc32 secondary
device#.....Done
Size = 2044830, CRC32 b31fcbc0
```


4. Enter the **verify** command with the **md5** and **secondary** keywords, and specify a value, to verify the hash value of a secondary image with a known value.

```
device# verify md5 secondary 01c410d6d153189a4a5d36c955653861
device#.....Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653862
Verification FAILED.
```

In the example above, the codes did not match and verification failed

5. Enter the **verify** command with the **md5** and **secondary** keywords, and specify a value, to verify the hash value of a secondary image with a known value.

```
device# verify md5 secondary 01c410d6d153189a4a5d36c955653861
device#.....Done
Size = 2044830, MD5 01c410d6d153189a4a5d36c955653861
Verification SUCCEEDED.
```

In the example above, the codes matched and verification succeeded.

6. Enter the **verify** command with the **sha** and **secondary** keywords, specifying a value, to generate a SHA-1 hash value for the secondary image with a known.

```
device# verify sha secondary 49d12d26552072337f7f5fcaef4cf4b742a9f525
device#.....Done
Size = 2044830, sha 49d12d26552072337f7f5fcaef4cf4b742a9f525
Verification SUCCEEDED.
```

7. Enter the **verify** command with the **crc32** and **secondary** keywords, specifying a value, to generate a CRC32 hash value for the secondary image with a known value.

```
device# verify crc32 secondary b31fcbc0
device#.....Done
Size = 2044830, CRC32 b31fcbc0
Verification SUCCEEDED.
```

Software Image File Types

This section lists the boot and flash image file types supported and how to install them on the RUCKUS ICX family of switches. For information about a specific version of code, refer to the release notes.

NOTE

The boot images are applicable to the listed devices only and are not interchangeable.

TABLE 3 Software Image Files

Product	Boot Image	Flash Image
ICX 7250	spzxxxx.bin	SPSxxxx.bin (Layer 2) or
ICX 7450		SPRxxxx.bin (Layer 3)

Flash Timeout

The operations that require access to the flash device are expected to be completed within the default flash timeout value of 12 minutes.

If the operations exceed the timeout value of 120 minutes, the flash device is locked and further flash operations cannot be processed. To facilitate prolonged flash operations without the device being locked, you can manually configure the flash timeout for a longer duration using the **flash-timeout** command. You can configure the flash timeout to a value from 12 through 120 minutes. The new timeout value is applicable for all flash operations and will be effective from the next flash operation. Refer to the *RUCKUS FastIron Command Reference* for more information.

Software Upgrades

For detailed instructions about upgrading the software, refer to the *RUCKUS FastIron Software Upgrade Guide*.

Boot Code Synchronization

RUCKUS devices support automatic synchronization of the boot image in the active and redundant management modules. When the new boot image is copied into the active module, it is automatically synchronized with the redundant management module.

NOTE

There is currently no option for manual synchronization of the boot image.

The following example activates the boot synchronization process.

```
device# copy tftp flash 10.20.65.194 /GA/SXZ07200.bin bootrom
Load to buffer (8192 bytes per dot)
.....Write to boot flash.....
TFTP to Flash Done.
Synchronizing with standby module...
Boot image synchronization done.
```

Viewing the Contents of Flash Files

You can view the list of files stored in flash memory. The following example displays a list of files stored in flash memory.

```
device# show files
Type      Size   Name
-----
F         24018046 primary
F         24018046 secondary
F           520 startup-config.backup
F           610 startup-config.txt

48037222 bytes 4 File(s) in FI root

1768706048 bytes free in FI root
1768706048 bytes free in /
```

Refer to the *RUCKUS FastIron Command Reference* for more information.

Using SNMP to Upgrade Software

You can use a third-party SNMP management application such as HP OpenView to upgrade software on a RUCKUS device.

NOTE

The syntax shown in this section assumes that you have installed HP OpenView in the "/usr" directory.

NOTE

It is recommended that you make a backup copy of the startup-config file before you upgrade the software. If you need to run an older release, you will need to use the backup copy of the startup-config file.

1. (Optional) Enter the **snmp-server community** command, with the **rw** keyword, to configure a read-write community string on the device, if one is not already configured.

```
device(config)# snmp-server community community_string rw
```

2. Enter the **no snmp-server pw-check** to disable password checking for SNMP set requests. If a third-party SNMP management application does not add a password to the password field when it sends SNMP set requests to a RUCKUS device, by default the RUCKUS device rejects the request.

```
no snmp-server pw-check
```

3. From the command prompt in the UNIX shell, enter the following command.

```
/usr/OV/bin/snmpset -c rw-community-string brcd-ip-addr 1.3.6.1.4.1.1991.1.1.2.1.5.0 ipaddress  
tftp-ip-addr 1.3.6.1.4.1.1991.1.1.2.1.6.0 octetstringascii file-name 1.3.6.1.4.1.1991.1.1.2.1.7.0  
integer command-integer
```

where

rw-community-string is a read-write community string configured on the RUCKUS device.

brcd-ip-addr is the IP address of the RUCKUS device.

tftp-ip-addr is the TFTP server IP address.

file-name is the image file name.

command-integer is one of the following.

20 - Download the flash code into the primary flash area.

22 - Download the flash code into the secondary flash area.

Software Reboot

You can use boot commands to immediately initiate software boots from a software image stored in primary or secondary flash on a RUCKUS device, or from a BootP or TFTP server. You can test new versions of code on a RUCKUS device or choose the preferred boot source from the console boot prompt without requiring a system reset.

NOTE

You must verify a successful TFTP transfer of the boot code before you reset the system. If the boot code is not transferred successfully but you try to reset the system, the system will not have the boot code with which to successfully boot.

By default, the RUCKUS device first attempts to boot from the image stored in its primary flash, then its secondary flash, and then from a TFTP server. You can modify this booting sequence from the global configuration mode of the CLI using the **boot system** command.

To initiate an immediate boot from the CLI, enter one of the **boot system** commands. Refer to the *RUCKUS FastIron Command Reference* for more information.

NOTE

When using the **boot system tftp** command, the IP address of the device and the TFTP server should be in the same subnet.

NOTE

If you are booting the device from a TFTP server through a fiber connection, use the following command: **boot system tftp ip-address filename fiber-port**.

NOTE

The **boot system tftp** command is not supported in a stacking environment.

Displaying the Boot Preference

You can view information about the boot sequence in the startup configuration and running configuration files.

Use the following commands to display the boot sequence in the startup configuration and running configuration files. The commands do not need to be entered in the specified order. Using these commands is optional and they can be entered in any order. For more information on these commands, refer to the *RUCKUS FastIron Command Reference*.

1. Enter the **show boot-preference** command to display the boot sequence in the startup configuration and running configuration files.

```
device> show boot-preference

Boot system preference(Configured):
  Boot system flash secondary
Boot system preference(Default):
  Boot system flash primary
  Boot system flash secondary
```

Information about the default boot sequence preference is displayed.

2. Enter the **show run** command to view detailed information about the configuration.

```
device# show run
Current configuration:
!
ver 08.0.40q042T213
!
stack unit 1
  module 1 icx7450-24-port-management-module
  module 2 icx7400-xgf-4port-40g-module
  module 3 icx7400-qsfp-lport-40g-module
  module 4 icx7400-qsfp-lport-40g-module
!
monitor-profile 1 type erspan
destination-ip 2.2.2.2
source-ip 1.1.1.1
!
monitor-profile 2 type erspan
destination-ip 2.2.2.2
source-ip 1.1.1.1
!
monitor-profile 3 type erspan
!
monitor-profile 4 type erspan
!
vlan 1 name DEFAULT-VLAN by port
!
vlan 10 by port
  tagged ethe 1/4/1
  interface ve 10
  multicast6 passive
  multicast6 pimsm-snooping
!
vlan 100 by port
  tagged ethe 1/1/1
  interface ve 100
!
system-max gre-tunnels 24
!
vrf vrf1
  rd 1:11
exit-vrf
!
vrf blue
  rd 1:1
exit-vrf
!
vrf vrf0
exit-vrf
!
vrf 0
exit-vrf
!
vrf v1
  rd 1:5
exit-vrf
!
buffer-sharing-full
!
priority-flow-control enable
optical-monitor 4000
boot sys fl sec
enable telnet authentication
ip dns domain-list englab.ruckus.com
ip dns server-address 10.x.x.x
ip show-service-number-in-log
ip route 0.0.0.0/0 10.xx.xx.xx distance 254
ip multicast passive
!
ipv6 multicast passive
telnet server enable vlan 10
```

```
!
batch buffer 1 c
hello-interval
host-max-num
c
!
dot1x-mka-enable
!
ip multicast-debug-mode
ip multicast-routing
ip multicast-routing rpf-check mac-movement
!
router pim
!
!
ipv6 router pim
!
interface management 1
 ip address 10.xxx.xxx.xxx 255.255.255.0 dynamic
!
interface ethernet 1/1/1
 port-name ERSPAN
 mon profile 1 both
 unknown-unicast limit 3 kbps
 port security
  age 2 absolute
!
interface ethernet 1/1/2
 ip address 1.1.1.1 255.255.255.0
 ip address 2.2.2.2 255.255.255.0
!
interface ve 10
!
interface ve 100
!
router msdp
 sa-filter originate route-map w2
!
end
```

The following example shows a user-configured boot sequence preference.

```
device# show boot-preference

Boot system preference(Configured):
  Boot system flash primary
Boot system preference(Default):
  Boot system flash primary
  Boot system flash secondary
```

System Reload Scheduling

In addition to reloading the system manually, you can configure the RUCKUS device to reload itself at a specific time or after a specific amount of time has passed.

NOTE

The scheduled reload feature requires the system clock. Refer to the NTP version 4 documentation at <http://doc.ntp.org/4.2.2/release.html>.

The following example schedules a system reload from the primary flash module for 6:00:00 AM, December 1, 2020.

```
device# reload at 06:00:00 12-01-15
```

The following example schedules a system reload from the secondary flash one day and 12 hours later.

```
device# reload after 01:12:00 secondary
```

The following example cancels a scheduled system reload.

```
device# reload cancel
```

Refer to the *RUCKUS FastIron Command Reference* for more information.

Diagnostic Error Codes and Remedies for TFTP Transfers

This section describes the error messages associated with TFTP transfer of configuration files, software images or flash images to or from a RUCKUS device.

Error code	Message	Explanation and action
1	Flash read preparation failed.	A flash error occurred during the download. Retry the download. If it fails again, contact customer support.
2	Flash read failed.	
3	Flash write preparation failed.	
4	Flash write failed.	
5	TFTP session timeout.	TFTP failed because of a time out. Check IP connectivity and make sure the TFTP server is running.
6	TFTP out of buffer space.	The file is larger than the amount of room on the device or TFTP server. If you are copying an image file to flash, first copy the other image to your TFTP server, then delete it from flash. (Use the erase flash ... CLI command from the privileged exec model to erase the image in the flash.) If you are copying a configuration file to flash, edit the file to remove unnecessary information, then try again.
7	TFTP busy, only one TFTP session can be active.	Another TFTP transfer is active on another CLI session, or Web management session, or network management system. Wait, then retry the transfer.
8	File type check failed.	You accidentally attempted to copy the incorrect image code into the system. For example, you might have tried to copy a Chassis image into a Compact device. Retry the transfer using the correct image.

Error code	Message	Explanation and action
16	TFTP remote - general error.	The TFTP configuration has an error. The specific error message describes the error. Correct the error, then retry the transfer.
17	TFTP remote - no such file.	
18	TFTP remote - access violation.	
19	TFTP remote - disk full.	
20	TFTP remote - illegal operation.	
21	TFTP remote - unknown transfer ID.	
22	TFTP remote - file already exists.	
23	TFTP remote - no such user.	

This section describes the error messages associated with the TFTP transfer of PoE firmware file to a RUCKUS device.

Message	Explanation and action
Firmware TFTP timeout.	TFTP failed because of a time out. Check IP connectivity and make sure the TFTP server is running.
Firmware is not valid for this platform.	Each Power over Ethernet (PoE) firmware file delivered by RUCKUS is meant to be used on the specific platform only. If the file is used on a platform for which it is not meant, then this error message will display. Download the correct file, then retry the transfer.
Firmware is not valid for the IEEE 802.3at (PoE-Plus) controller type.	Each PoE firmware file delivered by RUCKUS is meant to be used on the specific platform only. If the file is used on a platform for which it is not meant, then this error message will display. Download the correct file, then retry the transfer.
Firmware is not valid for the IEEE 802.3af PoE controller type.	
Firmware type cannot be detected from the firmware content.	Each PoE firmware file delivered by RUCKUS is meant to be used on the specific platform and the specific PoE controller on the specified module. If the file is used for a platform for which it is meant, but the PoE controller is not same then this error message will display. Download the correct file, then retry the transfer.
TFTP File not Valid for PoE Controller Type.	
Firmware TFTP remote file access failed.	The TFTP server needs read access on the PoE firmware file. Check the permissions on the file, then try again.

Network Connectivity Testing

After you install the network cables, you can test network connectivity to other devices by pinging those devices. You also can observe the LEDs related to network connection and perform trace routes.

For more information about observing LEDs, refer to the appropriate Hardware Installation Guide.

The following example verifies that a RUCKUS device can reach another device through the network.

```
device# ping 10.33.4.7
```

NOTE

If the device is a RUCKUS Layer 2 Switch or Layer 3 Switch, you can use the host name only if you have already enabled the Domain Name Server (DNS) resolver feature on the device from which you are sending the ping.

The following example issues an IPv4 traceroute.

```
device# traceroute 10.33.4.7
```

IEEE 802.3ah EFM-OAM

The IEEE 802.3ah Ethernet in the First Mile (EFM) standard specifies the protocols and Ethernet interfaces for using Ethernet over access links as a first-mile technology.

Using the Ethernet in the First Mile solution, you will gain broadcast Internet access, in addition to services, such as Layer 2 transparent LAN services, voice services over Ethernet Access networks, and video and multicast applications, reinforced by security and Quality of Service control in order to build a scalable network.

The in-band management specified by IEEE 802.3ah EFM standard defines the operations, administration and maintenance (OAM) mechanism needed for the advanced monitoring and maintenance of Ethernet links in the first mile. The OAM capabilities facilitate network operation and troubleshooting. Basic 802.3 frames convey OAM data between two ends of the physical link. EFM-OAM is optional and can be disabled on each physical port.

When OAM is present, two connected OAM sub-layers exchange protocol data units (OAMPDUs). OAMPDUs are standard-size frames that can be sent at a maximum rate of 10 frames per second. This limitation is necessary for reducing the impact on the usable bandwidth. It is possible to send each frame several times in order to increase the probability of reception. A combination of the destination MAC address, the Ethernet type/length field and subtype allow distinguishing OAMPDU frames from other frames.

OAM functionality is designed to provide reliable service assurance mechanisms for both provider and customer networks.

Network Deployment Use Case

The data-link layer OAM is targeted at last-mile applications, and service providers can use it for demarcation point OAM services.

Ethernet last-mile applications require robust infrastructure that is both passive and active. 802.3ah OAM aims to solve validation and testing problems in such an infrastructure.

Using the Ethernet demarcation, service providers can additionally manage the remote device without utilizing an IP layer. This can be done by using link-layer SNMP counters, request and reply, loopback testing, and other techniques.

EFM-OAM Protocol

The functionality of the EFM-OAM can be summarized under the following categories:

- **Discovery:** Discovery is the mechanism to detect the presence of an OAM sub-layer on the remote device. During the discovery process, information about OAM entities, capabilities, and configurations are exchanged.
- **Remote fault detection:** Provides a mechanism for an OAM entity to convey error conditions to its peer by way of a flag in the OAMPDUs.
- **Remote loopback:** This mechanism is used to troubleshoot networks and to isolate problem segments in a large network by sending test segments.

Discovery

Discovery is the first phase of EFM-OAM. At this phase, EFM-OAM identifies network devices along with their OAM capabilities. The Discovery process relies on the Information OAMPDUs. During discovery, the following information is advertised through the TLVs within periodic information OAMPDUs:

- OAM capabilities: Advertises the capabilities of the local OAM entity. Using this information, a peer can determine what functions are supported and accessible (for example, loopback capability).
- OAM mode: The OAM mode is conveyed to the remote OAM entity. The mode can be either active or passive, and can also be used to determine a device's functionality.
- OAMPDU configuration: This configuration includes the maximum OAMPDU size to delivery. In combination with the limited rate of 10 frames per second, this information can be used to limit the bandwidth allocated to OAM traffic.

Timers

Two configurable timers control the protocol, one determining the rate at which OAMPDUs are to be sent, and the second controlling the rate at which OAMPDUs are to be received to maintain the Discovery procedure from resetting.

- The timer should generate PDUs in the range of 1 - 10 PDUs per second. The default value is 1 PDU per second.
- The Hold timer assumes the peer is dead if no packet is received for a period of 1 - 10 seconds. The default value is 5 seconds.

Flags

Included in every OAMPDU is a flags field, which contains, besides other information, the status of the discovery process. There are three possible values for the status:

- Discovering: Discovery is in progress.
- Stable: Discovery is completed. Once aware of this, the remote OAM entity can start sending any type of OAMPDU.
- Unsatisfied: When there are mismatches in the OAM configuration that prevent OAM from completing the discovery, the discovery process is considered unsatisfied and cannot continue.

Process Overview

The discovery process allows local Data Terminating Entity (DTE) to detect OAM on a remote DTE. Once OAM support is detected, both ends of the link exchange state and configuration information (such as mode, PDU size, loopback support, and so on). If both DTEs are satisfied with the settings, OAM is enabled on the link. However, the loss of a link or a failure to receive OAMPDUs for five seconds may cause the discovery process the start over again.

DTEs may be in either active or passive mode. Active mode DTEs instigate OAM communications and can issue queries and commands to a remote device. Passive mode DTEs generally wait for the peer device to instigate OAM communications and respond to, but do not instigate, commands and queries. Rules of what DTEs in active or passive mode can do are discussed in the following sections.

Rules for Active Mode

A DTE in active mode:

- Initiates the OAM Discovery process
- Sends information PDUs
- May send event notification PDUs
- May send variable request or response PDUs

- May send loopback control PDUs

Exceptions

- A DTE in active mode does not respond to variable request PDUs from DTEs in passive mode
- A DTE in active mode does not react to loopback control PDUs from DTEs in passive mode

Rules for Passive Mode

A DTE in passive mode:

- Waits for the remote device to initiate the Discovery process
- Sends information PDUs
- May send event notification PDUs
- May respond to variable request PDUs
- May react to received loopback control PDUs
- Is not permitted to send variable request or loopback control OAMPDUs

Remote Failure Indication

Faults in Ethernet that are caused by slowly deteriorating quality are more difficult to detect than completely disconnected links. A flag in the OAMPDU allows an OAM entity to send failure conditions to its peer. The failure conditions are defined as follows:

- Dying gasp: This condition is detected when the receiver goes down. The dying gasp condition is considered as unrecoverable. The conditions for a dying gasp condition include:
 - Reload command (Warm reboot)
 - Boot system flash pri/sec command (Warm reboot)
 - Failure on the box (Cold reboot)
- Critical event: On any critical event, the DTE will set the critical event bit in the information OAMPDU. The device will generate critical event in the following cases:
 - When the temperature of the box breaches the warning/shutdown threshold
 - Fan failure

The battleshort mode allows you to prevent the shutdown of ICX 7450 when the temperature of the box breaches the warning or shutdown threshold. This is intended to be used in emergency conditions to allow the switches to function in a hostile environment as long as possible.

Enabling Battleshort Mode

The following task enables battleshort mode on a standalone device or globally on all stack units.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **ignore-temp-shutdown** command to enable battleshort mode on a standalone device or globally on all stack units.

```
device(config)# ignore-temp-shutdown  
Ignore temperature shutdown threshold has been enabled
```

If the **ignore-temp-shutdown** command is used at global level, it applies to all the units which are part of the stack.

The following example enables battleshort mode on a standalone device or globally on all stack units.

```
device# configure terminal
device(config)# ignore-temp-shutdown
Ignore temperature shutdown threshold has been enabled
```

The following example enables battleshort mode on an individual stack member.

```
device# configure terminal
device(config)# stack unit 2
device(config-unit-2)# ignore-temp-shutdown
Ignore temperature shutdown threshold has been enabled in Stack unit 2
```

Remote Loopback

An OAM entity can put its remote entity into loopback mode using a loopback control OAMPDU. This helps you ensure quality of links during installation or when troubleshooting. In loopback mode, each frame received is transmitted back on that same port except for OAMPDUs and pause frames. The periodic exchange of OAMPDUs must continue while in the loopback state to maintain the OAM session. The loopback command is acknowledged by responding with an information OAMPDU with the loopback state indicated in the state field.

NOTE

RUCKUS recommends to ensure that any higher layer protocol running over the local and remote loopback ports does not block the interfaces in the VLAN on which loopback traffic testing is being performed.

NOTE

Ethernet loopback and EFM-OAM remote loopback cannot be configured on the same interface.

NOTE

If EEE is enabled globally, the port ceases to be in the remote loopback mode.

EFM-OAM Error Disable Recovery

The error disable recovery feature enables the device to recover the EFM-OAM interface from the error-disabled state caused by reception of a critical event from the remote device.

The following example configures the device to recover the EFM-OAM interface from the error-disabled state caused by reception of a critical event from the remote device.

```
device# configure terminal
device(config)# errdisable recovery cause loam-critical-event
```

The ports recover automatically from the error-disabled state upon the expiry of the error disable recovery timeout value.

Configuring EFM-OAM

The EFM-OAM configuration includes the following procedural steps to enable EFM-OAM on an interface or multiple interfaces for advanced monitoring and maintenance of Ethernet network.

1. Enter the **configure terminal** command to enter global configuration mode.

```
device# configure terminal
```

2. Enter the **link-oam** command to enable the EFM-OAM protocol and enter EFM-OAM protocol configuration mode.

```
device(config)# link-oam
```

3. Enter the **timeout** command to configure the time in seconds for which the local Data Terminal Equipment (DTE) waits to receive OAM Protocol Data Units (OAM-PDUs) from the remote entity.

```
device(config-link-oam)# timeout 5
```

4. Enter the **pdu-rate** command to configure the number of PDUs to be transmitted per second by the DTE.

```
device(config-link-oam)# pdu-rate 2
```

5. Enter the **ethernet** command to enable EFM-OAM on an interface.

EFM-OAM can be enabled on more than one interface. You can also specify a range of interfaces to enable EFM-OAM on multiple interfaces.

You can set the operational mode of EFM-OAM as Active or Passive.

- Enter the **ethernetstackid/slot/portactive** command to set the EFM-OAM operational mode as active on an interface.

```
device(config-link-oam)# ethernet 1/1/3 active
device(config-link-oam)# ethernet 1/1/4 active
```

- Enter the **ethernetstackid/slot/porttostackid/slot/portactive** command to set the EFM-OAM operational mode as active on a range of interfaces.

```
device(config-link-oam)# ethernet 1/1/5 to 1/1/8 active
```

- Enter the **ethernetstackid/slot/portpassive** command to set the EFM-OAM operational mode as passive on an interface.

```
device(config-link-oam)# ethernet 2/1/1 passive
```

- Enter the **ethernetstackid/slot/porttostackid/slot/portpassive** command to set the EFM-OAM operational mode as passive on a range of interfaces.

```
device(config-link-oam)# ethernet 2/1/1 to 2/1/8 passive
```

6. (Optional) Enter the **ethernetstackid/slot/portallow-loopback** command to enable the interface to respond to a loopback request from the remote device.

```
device(config-link-oam)# ethernet 1/1/3 allow-loopback
```

7. (Optional) Enter the **ethernetstackid/slot/portremote-failure** command to set the device for the remote-failure action to be taken upon the reception of critical event information on the interface.

```
device(config-link-oam)# ethernet 1/1/3 remote-failure critical-event action block-interface
```

8. (Optional) Enter the **remote-loopbackethernetstackid/slot/port** command to start or stop the remote loopback procedure on a remote device.

```
device(config-link-oam)# remote-loopback ethernet 2/1/1 start
device(config-link-oam)# remote-loopback ethernet 2/1/1 stop
```

The following shows an example of EFM-OAM configuration.

```
device# configure terminal
device(config)# link-oam
device(config-link-oam)# timeout 5
device(config-link-oam)# pdu-rate 2
device(config-link-oam)# ethernet 1/1/3 active
device(config-link-oam)# ethernet 1/1/3 allow-loopback
device(config-link-oam)# remote-loopback ethernet 2/1/1 start
device(config-link-oam)# ethernet 1/1/3 remote-failure critical-event action block-interface
```

Displaying OAM Information

The following sample output of the **show link-oam info** command displays the OAM information on all OAM-enabled ports.

```
device# show link-oam info
Ethernet Link Status   OAM Status   Mode   Local Stable   Remote Stable
1/1/1      up           up        active   satisfied     satisfied
1/1/2      up           up        passive  satisfied     satisfied
1/1/3      up           up        active   satisfied     satisfied
1/1/4      up           init      passive  unsatisfied   unsatisfied
1/1/5      down        down      passive  unsatisfied   unsatisfied
1/1/6      down        down      passive  unsatisfied   unsatisfied
1/1/7      down        down      passive  unsatisfied   unsatisfied
```

The following sample output of the **show link-oam info detail** command displays detailed OAM information on all OAM-enabled ports.

```
device# show link-oam info detail
OAM information for Ethernet port: 10/1/1
+link-oam mode:      passive
+link status:        down
+oam status:         down
Local information
  multiplexer action: forward
  parse action:      forward
  stable:            unsatisfied
  state:             linkFault
  loopback state:    disabled
  dying-gasp:        false
  critical-event:    false
  link-fault:        true
Remote information
  multiplexer action: forward
  parse action:      forward
  stable:            unsatisfied
  loopback support:  disabled
  dying-gasp:        false
  critical-event:    true
  link-fault:        false

OAM information for Ethernet port: 10/1/3
+link-oam mode:      active
+link status:        up
+oam status:         down
Local information
  multiplexer action: forward
  parse action:      forward
  stable:            unsatisfied
  state:             activeSend
  loopback state:    disabled
  dying-gasp:        false
  critical-event:    false
  link-fault:        false
Remote information
  multiplexer action: forward
  parse action:      forward
  stable:            unsatisfied
  loopback support:  disabled
  dying-gasp:        false
  critical-event:    false
  link-fault:        false

OAM information for Ethernet port: 10/1/4
+link-oam mode:      active
+link status:        up
+oam status:         up
Local information
  multiplexer action: forward
  parse action:      forward
  stable:            satisfied
  state:             up
```

```

        loopback state:      disabled
        dying-gasp:          false
        critical-event:      false
        link-fault:          false
Remote information
        multiplexer action:  forward
        parse action:        forward
        stable:               satisfied
        loopback support:    disabled
        dying-gasp:          false
        critical-event:      true
        link-fault:          false

```

The following sample output of the **show link-oam info detail ethernet** command displays detailed OAM information on a specific Ethernet port.

```

device# show link-oam info detail ethernet 1/1/3
OAM information for Ethernet port: 1/1/3
+link-oam mode:      active
+link status:        up
+oam status:         up
Local information
        multiplexer action:  forward
        parse action:        forward
        stable:               satisfied
        state:                up
        loopback state:      disabled
        dying-gasp:          false
        critical-event:      false
        link-fault:          false
Remote information
        multiplexer action:  forward
        parse action:        forward
        stable:               satisfied
        loopback support:    disabled
        dying-gasp:          false
        critical-event:      false
        link-fault:          false

```

Displaying OAM Statistics

The following sample output of the **show link-oam statistics** command displays the OAM statistics on all OAM-enabled ports.

```

device# show link-oam statistics
Ethernet Tx Pdus      Rx Pdus
10/1/1  377908        377967
10/1/3   400           44
10/1/4   400           385
10/1/5   400           385
10/1/6   400           385

```

The following sample output of the **show link-oam statistics detail** command displays detailed OAM statistics on all OAM-enabled ports.

```

device# show link-oam statistics detail
OAM statistics for Ethernet port: 10/1/1
Tx statistics
        information OAMPDUs:      377908
        loopback control OAMPDUs: 0
        variable request OAMPDUs: 0
        variable response OAMPDUs: 0
        unique event notification OAMPDUs: 0
        duplicate event notification OAMPDUs: 0
        organization specific OAMPDUs: 0
        link-fault records:        0
        critical-event records:    0
        dying-gasp records:        0
Rx statistics
        information OAMPDUs:      377967
        loopback control OAMPDUs: 0
        loopback control OAMPDUs dropped: 0

```

Operations, Administration, and Maintenance

IEEE 802.3ah EFM-OAM

```
variable request OAMPDUs:      0
variable response OAMPDUs:     0
unique event notification OAMPDUs: 0
duplicate event notification OAMPDUs: 0
organization specific OAMPDUs: 0
unsupported OAMPDUs:           0
link-fault records:           0
critical-event records:       377395
dying-gasp records:           0
discarded TLVs:               0
unrecognized TLVs:            0
```

OAM statistics for Ethernet port: 10/1/3

```
Tx statistics
  information OAMPDUs:          427
  loopback control OAMPDUs:    0
  variable request OAMPDUs:    0
  variable response OAMPDUs:   0
  unique event notification OAMPDUs: 0
  duplicate event notification OAMPDUs: 0
  organization specific OAMPDUs: 0
  link-fault records:          0
  critical-event records:      0
  dying-gasp records:          0
Rx statistics
  information OAMPDUs:          44
  loopback control OAMPDUs:    0
  loopback control OAMPDUs dropped: 0
  variable request OAMPDUs:    0
  variable response OAMPDUs:   0
  unique event notification OAMPDUs: 0
  duplicate event notification OAMPDUs: 0
  organization specific OAMPDUs: 0
  unsupported OAMPDUs:         0
  link-fault records:          0
  critical-event records:      0
  dying-gasp records:          0
  discarded TLVs:              0
  unrecognized TLVs:           0
```

OAM statistics for Ethernet port: 10/1/4

```
Tx statistics
  information OAMPDUs:          428
  loopback control OAMPDUs:    0
  variable request OAMPDUs:    0
  variable response OAMPDUs:   0
  unique event notification OAMPDUs: 0
  duplicate event notification OAMPDUs: 0
  organization specific OAMPDUs: 0
  link-fault records:          0
  critical-event records:      0
  dying-gasp records:          0
Rx statistics
  information OAMPDUs:          413
  loopback control OAMPDUs:    0
  loopback control OAMPDUs dropped: 0
  variable request OAMPDUs:    0
  variable response OAMPDUs:   0
  unique event notification OAMPDUs: 0
  duplicate event notification OAMPDUs: 0
  organization specific OAMPDUs: 0
  unsupported OAMPDUs:         0
  link-fault records:          0
  critical-event records:      350
  dying-gasp records:          0
  discarded TLVs:              0
  unrecognized TLVs:           0
```


The following sample output of the **show link-oam statistics detailethernet** command displays detailed OAM statistics on a specific Ethernet port.

```
device# show link-oam statistics detail ethernet 1/1/3
OAM statistics for Ethernet port: 1/1/3
  Tx statistics
    information OAMPDUs:          122474
    loopback control OAMPDUs:    0
    variable request OAMPDUs:    0
    variable response OAMPDUs:   0
    unique event notification OAMPDUs: 0
    duplicate event notification OAMPDUs: 0
    organization specific OAMPDUs: 0
    link-fault records:          0
    critical-event records:      0
    dying-gasp records:          0
  Rx statistics
    information OAMPDUs:          94691
    loopback control OAMPDUs:    0
    loopback control OAMPDUs dropped: 0
    variable request OAMPDUs:    0
    variable response OAMPDUs:   0
    unique event notification OAMPDUs: 0
    duplicate event notification OAMPDUs: 0
    organization specific OAMPDUs: 0
    unsupported OAMPDUs:         0
    link-fault records:          0
    critical-event records:      0
    dying-gasp records:          0
    discarded TLVs:              0
    unrecognized TLVs:           0
```

EFM-OAM Syslog Messages

When EFM-OAM is enabled on an interface, the syslog messages in the following table are generated when the link goes up or down, or when loopback mode is entered or cleared on an interface.

TABLE 4 EFM-OAM Syslog Messages

Event	Syslog output
Port 1 is LOAM logically Up	Link-OAM: Logical link on interface Ethernet 1/1/1 is up.
Port 1 is LOAM logically Down	Link-OAM: Logical link on interface Ethernet 1/1/1 is down.
Port 1 entered remote Loopback mode	Link-OAM: Link entered remote loopback on ethernet 1/1/1
Port 1 cleared remote Loopback mode	Link-OAM: Link cleared remtote loopback on ethernet 1/1/1
Port 1 entered local Loopback mode	Link-OAM: Link entered local loopback on ethernet 1/1/1
Port 1 cleared local Loopback mode	Link-OAM: Link cleared local loopback on ethernet 1/1/1
Dying gasp event on port 1	Link-OAM: Link received dying-gasp event on ethernet 1/1/1
Critical event on port 1	Link-OAM: Link received critical event on ethernet 1/1/1

SNMP trap support is enabled for EFM-OAM from 08.0.70 release onwards.

Displaying Management Redundancy Information

You can view the redundancy parameter settings and statistics.

The following example displays the redundancy parameter settings and statistics.

```
device(config)# show redundancy
=== MP Redundancy Settings ===
Configured Active Slot = 9
Running-Config Sync Period = (upon "write mem")
=== MP Redundancy Statistics ===
Current Active Session:
Active mgmt slot = 9, Standby mgmt slot = 10 (Absent)
Switchover cause = No Switchover
Start Time       = Jan  1 00:00:09
Sxr Sys Hitless Enable Status = 0
Total number of Switchover/Failovers = 0
L3 slib baseline sync status: 0 [complete]
```

NOTE

Management redundancy is not available for the ICX 7450-24 platform.

Layer 3 Hitless Route Purge

Layer 3 traffic is forwarded seamlessly during a failover, switchover, or OS upgrade when hitless management is enabled.

Some protocols support non-stop routing. On enabling non-stop routing, after switchover the management module quickly re-converge the protocol database. Whereas, some protocols support graceful restart, in which the protocol state is re-established with the help of neighboring devices. Once all the protocols converge the routes which were removed from the network during the convergence period, the routes are deleted from the devices. You can set the route purge timer per VRF instance. Configure the timer to set the duration for which the routes should be preserved after switchover. Once this period elapses, the route purging starts, if by then all other protocols have finished non-stop routing or graceful restart.

When switchover occurs, the route purge timer starts. If non-stop routing or graceful restart is also configured, the route validation and purging starts only when they are complete and the purge timer has elapsed. If for some reason more delay is expected in learning the routes, you can configure a larger period for the purge timer.

Setting the IPv4 Hitless Purge Timer

You can configure the timer to set the duration for which the routes should be preserved after switchover. The following task configures the IPv4 hitless purge timer for the default VRF.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **ip hitless-route-purge-timer** command, specifying a value, to set the IPv4 hitless purge timer for the default VRF.

```
device(config)# ip hitless-route-purge-timer 60
```

The following example sets the IPv4 hitless purge timer on the default VRF to 60 seconds.

```
device(config)# ip hitless-route-purge-timer 60
```

The following example sets the IPv4 purge timer for a non-default VRF instance to 120 seconds.

```
device(config)# vrf blue
device(config-vrf-blue)# rd 10:10
device(config-vrf-blue)# address-family ipv4
device(config-vrf-blue-ipv4)# ip hitless-route-purge-timer 120
```

Setting the IPv6 Hitless Purge Timer

You can configure the timer to set the duration for which the routes should be preserved after switchover. The following task configures the IPv6 hitless purge timer for the default VRF.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **ipv6 hitless-route-purge-timer** command, and specify a value, to set the IPv6 hitless purge timer for the default VRF.

```
device(config)# ipv6 hitless-route-purge-timer 60
```

The following example sets the IPv6 hitless purge timer on the default VRF to 60 seconds.

```
device(config)# ipv6 hitless-route-purge-timer 60
```

The following example sets the IPv6 purge timer for a non-default VRF instance to 180 seconds.

```
device(config)# vrf blue
device(config-vrf-blue)# rd 10:10
device(config-vrf-blue)# address-family ipv6
device(config-vrf-blue-ipv4)# ipv6 hitless-route-purge-timer 60
```

Energy Efficient Ethernet

Energy Efficient Ethernet (EEE) regulates and saves power consumed by the active hardware components in the switch and conserves power during idle time.

EEE allows RUCKUS devices to conform to green computing standards. This functionality is achieved by moving the data ports to a low-power state when their function is not necessary or when they are in a passive, no traffic condition. EEE switching platforms reduces overall energy consumption, cooling, noise, and operating costs for energy and cooling. Lower power consumption also means lower heat dissipation, increased system stability, and less energy usage, thereby reducing costs and impact on the environment.

EEE is a set of enhancements to the Ethernet specification to address power consumption during periods of low data activity. EEE is specified in IEEE Std 802.3az-2010, which is an amendment to the IEEE Std 802.3-2008 specification. The optional EEE capability combines the IEEE 802.3 Media Access Control (MAC) sublayer with a family of physical layers defined to support operation in the Low Power Idle (LPI) mode. When the LPI mode is enabled, systems on both sides of the link can save power during periods of low link utilization. LPI signaling allows the LPI client to indicate to the PHY, and to the link partner, that a break in the data stream is expected. The LPI client can then use this information to enter power-saving modes that require additional time to resume normal operation. LPI signaling also informs the LPI client when the link partner sends such an indication. The client device connected to the EEE-enabled switch port must also support the LPI functionality in order to take advantage of EEE on the switch.

NOTE

Port flap may occur on the port when EEE is enabled or disabled.

Port Support for Energy Efficient Ethernet

- On ICX 7150 models, EEE is supported only on data ports.
- On ICX 7150 models, EEE is supported in standalone and homogeneous stacking topologies.
- EEE is supported on the following ports of ICX 7150 models:
 - ICX 7150-24 and ICX 7150-24P: Ports 1/1/1 through 1/1/8
 - ICX 7150-48, ICX 7150-48P, and ICX 7150-48PF: Ports 1/1/1 through 1/1/8 and 1/1/25 through 1/1/32
 - ICX 7150-C10ZP: Ports 1/1/1 through 1/1/8 (front panel data ports) in 2.5-Gbps and 1-Gbps speed
 - ICX 7150-48ZP: Ports 1/1/1 through 1/1/16 (front panel data ports) in 2.5-Gbps and 1-Gbps speed
- EEE is not supported on ICX 7150 MultiGig ports at 100-Mbps speed.
- EEE is not supported on ICX 7150 stacking ports, and uplink ports.
- On ICX 7250 models, EEE is supported on 1-Gbps copper ports.
- On ICX 7450 models, EEE is supported on 1-Gbps copper ports and 10-Gbps copper module ports.
- EEE is not supported on 1-Gbps fiber ports (ICX7450-48F), 4x10F module ports, and 1x40Q module ports.
- On the ICX 7650-48ZP, EEE is supported on ports 1/1/25 through 1/1/48 in 10-, 5-, 2.5-, and 1-Gbps speed.

NOTE

EEE statistics counters are not supported on multigigabit ports.

TABLE 5 Ports Supporting EEE on ICX 7550 Models

Models	Module Number	Type	Media	Purpose	Speed	Details	Port Number
ICX 7550-24	1	24x1-Gbps	Copper	Downlink	1000BASE-T	Base Module	1/1/1 to 1/1/24
ICX 7550-48	1	48x1-Gbps	Copper	Downlink	1000BASE-T	Base Module	1/1/1 to 1/1/48
ICX 7550-24P (PoE)	1	24x1-Gbps (PoE)	Copper	Downlink	1000BASE-T	Base Module, 802.3bt Type-2 supported, 2-pair, 802.3bt Type-2, Class 5, up to 45W per port PoE	1/1/1 to 1/1/24
ICX 7550-48P (PoE)	1	48x1-Gbps (PoE)	Copper	Downlink	1000BASE-T	Base Module, 802.3bt Type-2 supported, 2-pair, 802.3bt Type-2, Class 5, up to 45W per port PoE	1/1/1 to 1/1/48
ICX 7550-24ZP	1	12x2.5-Gbps (Multigigabit-lite)	Copper	Downlink	2.5GBASE-T and 1000BASE-T	Base Module, 4-pair, 802.3bt Type-4, up to 90W per port PoE	1/1/1 to 1/1/12
		12x10-Gbps (Multigigabit)	Copper	Downlink	10GBASE-T, 5GBASE-T, 2.5GBASE-T, and 1000BASE-T	Base Module, 4-pair, 802.3bt Type-4, up to 90W per port PoE	1/1/13 to 1/1/24

TABLE 5 Ports Supporting EEE on ICX 7550 Models (continued)

Models	Module Number	Type	Media	Purpose	Speed	Details	Port Number
ICX 7550-48ZP	1	36x2.5-Gbps (Multigigabit-lite)	Copper	Downlink	2.5GBASE-T and 1000BASE-T	Base Module, 4-pair, 802.3bt Type-4, up to 90W per port PoE	1/1/1 to 1/1/36
		12x10-Gbps (Multigigabit)	Copper	Downlink	10GBASE-T, 5GBASE-T, 2.5GBASE-T, and 1000BASE-T	Base Module, 4-pair, 802.3bt Type-4, up to 90W per port PoE	1/1/37 to 1/1/48

NOTE

Enable auto-negotiation for Energy Efficient Ethernet (EEE) support on ICX 7550 models.

Enabling Energy Efficient Ethernet

Energy Efficient Ethernet (EEE) is supported on select ICX devices and can be enabled globally or per port. Complete the following steps to enable EEE globally or per port.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **eee** command to enable EEE globally.

```
device(config)# eee
EEE Feature Enabled
```

3. Enter the **interface ethernet** command, and specify an interface, to enter interface configuration mode for the specified interface.

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

4. Enter the **eee** command to enable EEE for the interface.

```
device(config-if-e1000-1/1/1)# eee
EEE Feature Enabled on port 1/1/1
```

Histogram Information Overview

The histogram feature monitors and records system resource usage information. The main objective of the histogram is to record resource allocation failures and task CPU usage information. The histogram feature keeps track of task execution information, context switch history of tasks, buffer allocation failure, and memory allocation failure.

The histogram information is collected and maintained internally, in a cyclical buffer. It can be reviewed to determine if resource allocation failures or task CPU usage may have contributed to an application failure.

NOTE

Histogram information is not maintained across reboot.

Displaying CPU Histogram Information

The CPU histogram provides information about task CPU usage. The CPU histogram is viewed in the form of buckets (task usage is divided into different interval levels called *buckets*). For example, the task run time is divided into buckets: bucket 1 (0-50 ms), bucket 2 (50-100 ms), bucket 3 (100-150 ms), and so on. The CPU histogram collects the task CPU usage in each bucket. This includes how many times a task run time or hold time falls in each bucket, and the maximum run time and total run time for each bucket. CPU histogram information is measured for the hold-time and wait-time of the task.

- Hold time - The time that the task is holding the CPU without yield.
- Wait time - The time that the task is waiting for execution.

External USB Hotplug

External USB Hotplug support allows you to copy images, cores, logs, and configurations between the external USB and the internal eUSB.

RUCKUS device images are stored in the raw partition. Cores, logs, and configurations are stored in the ext4 filesystem partition. The introduction of the External USB Hotplug gives you the option to easily copy device images, cores, logs, and configurations between the external USB and the internal flash.

External USB Hotplug Considerations

- Only USB drives of up to 128 GB of any vendor type are supported.
- USB 3.0 is not supported.
- You can copy files of less than 2 GB only.
- Make sure the external USB is formatted as a "FAT" filesystem before attempting to use it. Formatting can be done on a PC or on the RUCKUS device using the **format disk0** command.
- You should not insert a USB-based disk drive or a USB hub to connect multiple USB disks.
- Copying TFTP/SCP to disk0 and disk0 to TFTP/SCP is not supported.
- Only an administrator can execute operations on an external USB, similar to TFTP.
- You cannot access the active unit's local external USB from a member unit and vice versa.
- Booting from an external USB is not supported.
- You must run the **unmount disk0** command before unplugging the external USB. The external USB can be mounted using the **mount disk0** command.
- The USB drive is only functional on the active member in a stacked environment.

Using External USB Hotplug

Plug in the External USB to begin using the External USB Hotplug commands.

You can use the commands in the following table as part of the External USB Hotplug functionality.

TABLE 6 External USB Hotplug Commands

Command	Description
show files disk0	Verifies that the external USB is mounted and ready to use, and displays the files in the external USB drive.
format disk0	Formats the external USB.
mount disk0	Mounts the file system in the external USB drive.

TABLE 6 External USB Hotplug Commands (continued)

Command	Description
unmount disk0	Unmounts the file system of the external USB drive. This command is required to safely plug out the USB, so that files are not lost or corrupted.
copy flash disk0 <code>{{primary secondary}}filename</code>	Copies the image binary stored in the primary or secondary partition of the flash to a destination file in the external USB.
copy flash disk0 file	Copies any file from a source file in the system flash to an external USB destination file.
copy disk0 license	Copies the license file present in the external USB drive to the system.
copy disk0 running-config	Copies the configuration file present on the external USB drive to the system's running configuration.
copy disk0 startup-config	Copies the configuration file present on the external USB drive to the system's startup configuration file.
copy disk0 system-manifest <code>{filename{primary secondary}}</code>	Copies the system-manifest file present on the external USB drive to the primary or secondary flash image on the device.

Refer to the *RUCKUS FastIron Command Reference* for details and examples on using the External USB Hotplug commands.

Basic System Management

The following sections contain procedures and examples for basic system management tasks.

Viewing System Information

You can access software and hardware specifics for a RUCKUS Layer 2 switch or Layer 3 switch. For software specifics, refer to the section [Software Versions Installed and Running on a Device](#) on page 14.

The following shows example displays software and hardware details for the system.

```
device# show version
Copyright (c) 1996-2015 Ruckus Networks. All rights reserved.
UNIT 1: compiled on Oct 1 2015 at 11:29:56 labeled as SPR08040q042
(24018046 bytes) from Secondary SPR08040q042.bin
SW: Version 08.0.40q042T213
Compressed Boot-Monitor Image size = 786944, Version:10.1.05T215 (spz10105b008)
Compiled on Thu Jul 16 06:27:06 2015

HW: Stackable ICX7450-24
Internal USB: Serial #: 9900614090900038
Vendor: ATP Electronics, Total size = 1919 MB
=====
UNIT 1: SL 1: ICX7450-24 24-port Management Module
Serial #:CYT3346K035
License: ICX7450_L3_SOFT_PACKAGE (LID: eavIIJLmFIK)
License Compliance: ICX7450-PREM-LIC-SW is Compliant for next 45 days
P-ASIC 0: type B548, rev 01 Chip BCM56548_A0
=====
UNIT 1: SL 2: ICX7400-4X10GF 4-port 40G Module
Serial #:CYV3346K07G
=====
UNIT 1: SL 3: ICX7400-1X40GQ 1-port 40G Module
Serial #:CYX3346K06F
=====
UNIT 1: SL 4: ICX7400-1X40GQ 1-port 40G Module
Serial #:CYX3346K00A
=====
1000 MHz ARM processor ARMv7 88 MHz bus
8192 KB boot flash memory
2048 MB code flash memory
2048 MB DRAM
```

Operations, Administration, and Maintenance

Basic System Management

```
STACKID 1 system uptime is 31 day(s) 1 hour(s) 1 minute(s) 5 second(s)
The system : started=cold start
```

The following hardware details are listed in the output of the **show version** command:

- Chassis type
- PROM type (if applicable)
- Chassis serial number
- Management and interface module serial numbers and ASIC types

Refer to the *RUCKUS FastIron Command Reference* for more information.

For a description of the software details in the output of the **show version** command, refer to the section [Software Versions Installed and Running on a Device](#) on page 14.

You can also view the serial number pluggable modules. If there are no pluggable modules on the device, the serial number of the fixed modules on the device is displayed. The following is an example of the **show version** output on an ICX 7850/ICX 7850.

```
device(config)#show version
Copyright (c) Ruckus Networks, Inc. All rights reserved.
UNIT 1: compiled on May 27 2019 at 05:17:20 labeled as TNR08091
(61269864 bytes) from Secondary TNR08091.bin (UFI)
SW: Version 08.0.91T233
Compressed Secondary Boot Code size = 1573376, Version:10.1.16T235 (tnu10116)
Compiled on Sat May 25 15:39:48 2019
UNIT 2: compiled on May 27 2019 at 05:17:20 labeled as TNR08091
(61269864 bytes) from Secondary TNR08091.bin (UFI)
SW: Version 08.0.91T233
Compressed Secondary Boot Code size = 1573376, Version:10.1.16T235 (tnu10116)

HW: Stackable ICX7850-48FS
=====
UNIT 1: SL 1: ICX7850-48FS-L3-BASE 48-port Management Module
Serial #:FLV3334P010
Software Package: ICX7850_L3_SOFT_PACKAGE
Current License: l3-prem
P-ASIC 0: type B873, rev 01 Chip BCM56873_A0
=====
UNIT 1: SL 2: ICX7800-8X100G 8-port 800G Module
=====
UNIT 2: SL 1: ICX7850-48F-L3-BASE 48-port Management Module
Serial #:FLW3332P01E
Software Package: ICX7850_L3_SOFT_PACKAGE
Current License: l3-prem
=====
UNIT 2: SL 2: ICX7800-8X100G 8-port 800G Module
=====
2000 MHz ARM processor ARMv8 88 MHz bus
8192 KB boot flash memory
2048 MB code flash memory
3910 MB DRAM
STACKID 1 system uptime is 7 day(s) 2 hour(s) 7 minute(s) 23 second(s)
STACKID 2 system uptime is 7 day(s) 3 hour(s) 17 minute(s) 55 second(s)
The system started at 14:15:36 GMT+05:30 Thu May 30 2019

The system : started=warm start reloaded=by "stack"
My stack unit ID = 1, bootup role = active
```

Starting with FastIron 8.0.40, there is a **show version** command option that specifies the version running on a single unit. The following is an example of the **show versionunit1** command output on an ICX 7850.

```
device(config)#show version unit 1
Copyright (c) Ruckus Networks, Inc. All rights reserved.
UNIT 1: compiled on May 27 2019 at 05:17:20 labeled as TNR08091
(61269864 bytes) from Secondary TNR08091.bin (UFI)
SW: Version 08.0.91T233
Compressed Secondary Boot Code size = 1573376, Version:10.1.16T235 (tnu10116)
```


Compiled on Sat May 25 15:39:48 2019

```
HW: Stackable ICX7850-48FS
=====
UNIT 1: SL 1: ICX7850-48FS-L3-BASE 48-port Management Module
      Serial #:FLV3334P010
      Software Package: ICX7850_L3_SOFT_PACKAGE
      Current License: l3-prem
      P-ASIC 0: type B873, rev 01 Chip BCM56873_A0
=====
UNIT 1: SL 2: ICX7800-8X100G 8-port 800G Module
=====
2000 MHz ARM processor ARMv8 88 MHz bus
8192 KB boot flash memory
2048 MB code flash memory
3910 MB DRAM
STACKID 1 system uptime is 7 day(s) 2 hour(s) 7 minute(s) 29 second(s)
The system started at 14:15:37 GMT+05:30 Thu May 30 2019

The system : started=warm start reloaded=by "stack"
My stack unit ID = 1, bootup role = active
```

Viewing Configuration Information

You can view a variety of configuration details and statistics with the **show** option. The **show** command provides a convenient way to check configuration changes before saving them to flash.

The available show commands vary for Layer 2 and Layer 3 switches and by configuration mode.

To determine the available show commands for the system or a specific mode of the CLI, enter the following command.

```
device# show ?
```

You can also enter **show** at the command prompt, and then press the TAB key.

Enabling the Display of the Elapsed Timestamp for Port Statistics Reset

Whenever the port statistics of a device are cleared globally or on an interface, the counter values of the received and transmitted packets on the device are reset for all the ports or for an interface, respectively.

The elapsed time after the most recent reset of the port statistics counters can be displayed in the output of the **show statistics** command by configuring the **port-statistics-reset-timestamp enable** command. By default, the display of the elapsed timestamp information is disabled.

The elapsed time is calculated as the time between the most recent reset of the port statistics counters and the time when the **show statistics** command is executed.

The following list provides details of the conditions under which the port statistics counters are reset and also explains the elapsed time calculation considerations.

- When the port statistics are cleared individually using the **clear statistics ethernet** command. The elapsed time is calculated and displayed only for that particular interface.
- When the port statistics are cleared globally using the **clear statistics** command. The port statistics counters for all the ports, including management ports, are cleared and the elapsed time is calculated and displayed for each of the interfaces.
- When the management interface is cleared using the **clear statistics management** command. The port statistics counters specific to management ports are cleared. The elapsed time is calculated and displayed for the management interface.
- If the system is reloaded (hard reboot or soft reboot), the port statistics on the device are cleared automatically. In this case, the time when the ports are cleared during the reload is considered as the most recent reset time.

- In a stacking device, the Elapsed Timestamp information is applicable for other unit's ports. In case of a switchover, all the port statistics are cleared and the elapsed time is calculated and displayed for all ports.
- If hitless failover is enabled and if any unit is reloaded, the statistics of the reloading device's interfaces are cleared. In this case, the time when the ports are cleared during the reload is considered as the most recent reset time.
- The elapsed time is not impacted when the Network Time Protocol (NTP) syncs up with a different time other than the recorded time.

For more information refer to the *RUCKUS FastIron Command Reference*

Viewing port statistics

Port statistics are polled by default every 10 seconds.

You can view statistics for ports by entering the following **show** commands:

- **show interfaces**
- **show configuration**
- **show statistics**

The Elapsed Timestamp information is displayed in the output of the following **show** commands:

- **show statistics**
- **show statistics brief**
- **show statistics ethernet**
- **show statistics management**

Refer to the *RUCKUS FastIron Command Reference* for more information.

The following displays elapsed timestamp information for a specified Ethernet interface.

```
device# show statistics ethernet 1/1/13
Port      Link      State Dupl Speed Trunk Tag Pvid Pri   MAC           Name
1/1/13    Up        Forward Full 1G   None No  1    0   748e.f893.065c

Port 1/1/13 Counters:
*Last time counter reset (Elapsed Timestamp): 1 hour(s) 21 minute(s) 12 second(s)
InOctets      50218819740      OutOctets      50216689676
InPkts        63180119         OutPkts        63428168
InBroadcastPkts      5      OutBroadcastPkts      3
InMulticastPkts      63180114      OutMulticastPkts      63428165
InUnicastPkts              OutUnicastPkts
InBadPkts
InFragments
InDiscards              OutErrors
CRC                    Collisions
InErrors              LateCollisions
InGiantPkts          0
InShortPkts
InJabber
InFlowCtrlPkts              OutFlowCtrlPkts
InBitsPerSec      97441855      OutBitsPerSec      97432612
InPktsPerSec      153280      OutPktsPerSec      153972
InUtilization      100.00%      OutUtilization      100.00%
```

Viewing STP Statistics

You can use various **show** commands to view information about Spanning Tree Protocol (STP) statistics for Layer 2 and Layer 3 switches.

Use one of the following commands to view STP statistics. The commands do not need to be entered in the specified order. Using these commands is optional and they can be entered in any order. For more information on the full list of **show span** commands, refer to the *RUCKUS FastIron Command Reference*.

1. Enter the **show span** command to display general STP information.

```
device> show span

VLAN 1 BPDU cam_index is 3 and the Master DMA Are(HEX)
STP instance owned by VLAN 1
Global STP (IEEE 802.1D) Parameters:
VLAN      Root      Root      Root      Prio      Max      He-      Ho-      Fwd      Last      Chg      Bridge
ID        ID          Cost      Port      rity      Age      llo      ld       dly      Chang   cnt
Address
1         800000e0804d4a00 0         Root      Hex      sec      sec      sec      sec      sec      689     1
00e0804d4a00
Port STP Parameters:
Port      Prio      Path      State      Fwd      Design  Designated  Designated
Num       rity      Cost      State      Trans   Cost    Root        Bridge
Hex
1         80        19        FORWARDING 1         0        800000e0804d4a00 800000e0804d4a00
2         80        0         DISABLED   0         0        0000000000000000 0000000000000000
3         80        0         DISABLED   0         0        0000000000000000 0000000000000000
4         80        0         DISABLED   0         0        0000000000000000 0000000000000000
5         80        19        FORWARDING 1         0        800000e0804d4a00 800000e0804d4a00
6         80        19        BLOCKING   0         0        800000e0804d4a00 800000e0804d4a00
7         80        0         DISABLED   0         0        0000000000000000 0000000000000000
<lines for remaining ports excluded for brevity>
```

2. Enter the **show span** command with the **detail** keyword to display detailed STP information for a port.

```
device# show span detail
=====
VLAN 1 - MULTIPLE SPANNING TREE (MSTP) ACTIVE
=====
Bridge identifier - 0x800000e0804d4a00
Active global timers - Hello: 0
Port 1/1/1 is FORWARDING
Port - Path cost: 19, Priority: 128, Root: 0x800000e052a9bb00
Designated - Bridge: 0x800000e052a9bb00, Interface: 1, Path cost: 0
Active Timers - None
BPDUs - Sent: 11, Received: 0
Port 1/1/2 is DISABLED
Port 1/1/3 is DISABLED
Port 1/1/4 is DISABLED <lines for remaining ports excluded for brevity>
```

3. Enter the **show span** command with the **vlan** keyword, specifying a VLAN, to display STP information for the VLAN.

```
device# show span vlan 100
STP instance owned by VLAN 100

Global STP (IEEE 802.1D) Parameters:

VLAN Root          Root Root      Prio Max He- Ho- Fwd Last   Chg Bridge
ID   ID              Cost Port          rity Age llo ld  dly Chang cnt Address

   100 8000cc4e24b46fcc 0    Root          Hex  sec sec sec sec sec
                        8000 20  2   1   15  11          1   cc4e24b46fcc

Port STP Parameters:

Port   Prio Path  State      Fwd  Design  Designated  Designated
Num    rity Cost  State      Trans Cost      Root         Bridge
      Hex

1/1/1   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
1/1/2   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
1/1/3   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
1/1/4   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
1/1/5   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
1/1/6   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
1/1/7   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
1/1/8   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
lg1     80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
lg256   80   4   FORWARDING  1     0     8000cc4e24b46fcc 8000cc4e24b46fcc
```

Clearing Statistics

You can clear statistics for many parameters using the **clear** command.

To determine the available **clear** commands for the system, enter the **clear** command from the privileged exec mode of the CLI.

```
device# clear ?
```

You also can enter **clear** at the command prompt, then press the TAB key.

Refer to the *RUCKUS FastIron Command Reference* for more information.

Viewing Egress Queue Counters

You can display the number of packets that were queued for each QoS priority (traffic class) and dropped because of congestion for a port. The egress queue counters are displayed in the command output, as shown in the following example. This command output also displays the total of unicast and multicast counters for any particular QOS priority.

```
device> show interface ethernet 1/1/1

10GigabitEthernet 1/1/1 is down, line protocol is down
Port down for 16 hours 16 minutes 48 seconds
Hardware is 10GigabitEthernet , address is 748e.f8f9.6280 (bia 748e.f8f9.6280)
Interface type is 40Gig Fiber
Configured speed 40Gbit, actual unknown, configured duplex fdx, actual unknown
Configured mdi mode AUTO, actual unknown
Member of L2 VLAN ID 1, port is untagged, port state is BLOCKING
```

```
BPDU guard is Disabled, ROOT protect is Disabled, Designated protect is Disabled
Link Error Dampening is Disabled
STP configured to ON, priority is level0, mac-learning is enabled
Flow Control is enabled
Mirror disabled, Monitor disabled
Mac-notification is disabled
Not member of any active trunks
Not member of any configured trunks
No port name
IPG MII 96 bits-time, IPG GMII 96 bits-time
MTU 1500 bytes
300 second input rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
300 second output rate: 0 bits/sec, 0 packets/sec, 0.00% utilization
0 packets input, 0 bytes, 0 no buffer
Received 0 broadcasts, 0 multicasts, 0 unicasts
0 input errors, 0 CRC, 0 frame, 0 ignored
0 runts, 0 giants
0 packets output, 0 bytes, 0 underruns
Transmitted 0 broadcasts, 0 multicasts, 0 unicasts
0 output errors, 0 collisions
Relay Agent Information option: Disabled
```

```
Egress queues:
Queue counters    Queued packets    Dropped Packets
0                 0                 0
1                 0                 0
2                 0                 0
3                 0                 0
4                 0                 0
5                 0                 0
6                 0                 0
7                 0                 0
```

Clearing the Egress Queue Counters

You can clear egress queue statistics (reset them to zero).

The following example clears the statistics for a specific Ethernet interface.

```
device(config)# clear statistics ethernet 1/1/1
```

The following example egress queue statistics globally.

```
device(config)# clear statistics
```

Collecting CPU Packet Statistics

You can collect statistics on packets destined for the CPU. These statistics can be used to help troubleshoot high CPU issues.

Packet statistics are collected for a set of specified fields, such as Layer 2 destination MAC address, Layer 2 MAC type, Layer 2 source address, and inbound Ethernet port. For packets matching the specified fields, the packet information is copied to a hash table.

CPU packet statistics can be collected for the units in a stacking system. Every 30 seconds the details for queues, ports, and packet statistics are synced. Every 60 seconds the queues and ports history is synced.

Complete the following steps to configure and display CPU packet statistics.

1. Enter the **pstat field-add** command to configure the fields for which CPU packet statistics will be collected. The command can be entered more than once to configure multiple fields.

```
device(config)# pstat field-add l2-dest-mac
device(config)# pstat field-add input-port
```

2. Enter the **pstat max** command to configure the maximum number fields to be used for collecting statistics.

```
device(config)# pstat max 3
```

3. Enter the **pstat start** command to initiate the collection of CPU packet statistics.

```
device(config)# pstat start
```

The **pstat stop** command can be used to stop the collecting of packet statistics.

4. Use the **show pstat** command to display the CPU packet statistics counters.

```
device(config)#show pstat 11
```

input-port	l2-dest-mac	l2-dest-mac-type	Count
mgmt1	0100.5e00.0002	Multicast	19
11/1/7	0180.c200.0000	Multicast	10
2/1/7	0180.c200.000e	Multicast	1
11/1/7	0180.c200.000e	Multicast	1
mgmt1	0180.c200.0000	Multicast	10
mgmt1	cf4e.2445.0400	Multicast	19
mgmt1	778e.f8d4.00c0	Multicast	63
mgmt1	ffff.ffff.ffff	Broadcast	23

```
Number of Entries = 8
```

The following commands can also be used to display CPU packet statistics:

- **show pstat hist**: Displays per-second CPU packet statistics for the specified period of time.
- **show pstat dump**: Displays the contents of the CPU queue and port status.
- **show pstat status**: Displays the fields specified for collecting CPU packet statistics and whether CPU packet statistics collection is enabled.

Refer to the *RUCKUS FastIron Command Reference* for more information.

5. Enter the **clear pstat** command to clear the CPU packet statistics counters.

Link Fault Signaling for 10Gbps Ethernet Devices

Link Fault Signaling (LFS) is a physical layer protocol that enables communication on a link between two 10 Gbps Ethernet devices. When configured on a RUCKUS 10 Gbps Ethernet port, the port can detect and report fault conditions on transmit and receive ports.

When LFS is enabled on an interface, the following Syslog messages are generated when the link goes up or down, or when the TX or RX fiber is removed from one or both sides of the link that has LFS enabled.

```
Interface ethernet 1/1/1, state down - link down  
Interface ethernet 1/1/1, state up
```

When a link fault occurs, the Link and Activity LEDs turn OFF.

The Link and Activity LEDs turn ON when there is traffic traversing the link after the fiber is installed.

Enabling Link Fault Signaling

The following task enables Link Fault Signaling (LFS) between two 10 Gbps Ethernet devices.

On RUCKUS FastIron devices, RX LFS is always enabled by default and cannot be disabled. The **[no] link-fault-signal** command only applies to enabling or disabling TX LFS.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **interface** command, specifying an interface, to enter interface configuration mode for the specified interface.

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

3. Enter the **link-fault-signal** command to enable Link Fault Signaling (LFS) between two 10 Gbps Ethernet devices.

```
device(config-if-e1000-1/1/1)# link-fault-signal
```

The following example enables Link Fault Signaling (LFS) between two 10 Gbps Ethernet devices.

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# link-fault-signal
```

Viewing the Status of LFS-enabled Links

You can use the **show interface** command to view the status of LFS-enabled links.

Use one of the following commands to view the status of LFS-enabled links. Using these commands is optional and they can be entered in any order. For more information refer to the *RUCKUS FastIron Command Reference*.

1. Enter the **show interface** command with the **ethernet** keyword, specifying an interface, to display general information about the status of LFS-enabled links.

```
device> show interface ethernet 1/1/10

10GigabitEthernet1/1/10 is down (remote fault), line protocol is down
  Hardware is 10GigabitEthernet, address is 0000.0027.79d8 (bia 0000.0027.79d8)
  Configured speed 10Gbit, actual unknown, configured duplex fdx, actual unknown
  Member of L2 VLAN ID 1, port is untagged, port state is BLOCKING
  BPDU guard is Disabled, ROOT protect is Disabled
  Link Fault Signaling is Enabled, Link Error Dampening is Disabled
  STP configured to ON, priority is level0
  Flow Control is disabled
  mirror disabled, monitor disabled
<Truncated for brevity...>
```

The above output shows that the LFS-enabled link (port 1/1/10) is down because of an error on the remote port.

2. Enter the **show span** command with the **brief** keyword to display summarized information about the status of LFS-enabled links.

```
device# show interfaces brief
Port Link      State Dupl Speed Trunk Tag Pvid Pri MAC Name
1/1/10 Err-LFS
None      None None None No 1 0 0000.0027.79d8
```

The command output indicates that there is an error on the LFS-enabled link on port 1/1/10 and the link is down.

Locating a Device Using Port LEDs

Locating a device in a rack of many interconnected devices can be a difficult task. The LED ON/OFF feature allows you to turn on all the port LEDs to display steady green, irrespective of the port status. When all the port LEDs are turned on to steady green, locating a specific device, a standalone unit, stacked unit, or PE unit, becomes easier.

In a stacking environment, the LED ON/OFF feature can be enabled from the active unit, standby unit, member unit, or PE unit. If the port LEDs are turned on from the active unit, all the port LEDs of the active unit and the standby units are turned on. To identify or locate a specific unit, you must specify the unit ID of the respective unit. After locating the device, you can either turn off all the port LEDs or set the port LEDs to the default status.

The default LED status is the status of the LEDs according to the current status mode, if the status mode is supported. The **led default** command turns on all port LEDs with the default status. If the status mode is not supported, the link status or port status is considered to be the default LED status. The default LED status behavior also applies after the device is rebooted when all port LEDs are turned on using the **led on** command.

NOTE

The LED ON/OFF feature is supported only on the platforms where LED status mode is available. The status mode is supported on the ICX 7150, ICX 7650, ICX 7850, and ICX 7550 devices.

The following example shows how to turn on port LEDs on a standalone unit or stack unit.

```
device# led on
```

The following example shows how to turn on port LEDs to locate a unit in a stack from the active unit.

```
device# led on unit 5
```

Blink LED

The **led blink** command sets all the green port LEDs blinking for a given time interval on active, standby, and member units. The port then goes to a state based on the mode button. This option can be used in SmartZone environment to identify any device. The LED state can be altered within the blink interval by external events such as mode button or cable/SFP plugin. The configurable time interval is from 15 seconds to a maximum of 120 seconds. The default time interval is 15 seconds.

NOTE

The blink LED feature is supported on the ICX 7150, ICX 7650, ICX 7850, and ICX 7550 devices.

You can use the **led blink stack all** command from an active to set the port LEDs of all units in a stack to blinking for 50 seconds. After the specified time interval, the device goes to the default status.

```
device# led blink stack all 50
```

Use the following command to set the port LEDs blinking for a given stack ID and time interval.

```
device# led blink stack 1 25
```

led blink stack <stack unit-id> <blink interval in seconds (15-120sec)>

led blink <seconds>

Limitations

The following details must be considered when the blink LED or the LED ON/OFF feature is used:

- The status mode feature cannot be used when the blink LED or the LED ON/OFF feature is used. . If the status mode is changed, the LED status will be changed to the respective mode.

- If the status of the port changes due to any events such as linkup or linkdown, plugging in or unplugging a USB, optical interruption, and so on, the LED status of that particular port LED is changed.

Hardware Component Monitoring

- Virtual Cable Testing..... 51
- Digital Optical Monitoring..... 54
- Syslog Messages for Optical Transceivers..... 58

Virtual Cable Testing

Virtual Cable Tester (VCT) is a cable diagnostic feature in the physical layer (PHY) used for fault detection and advanced cable performance monitoring.

VCT technology enables the diagnosis of a conductor (wire or cable) by sending a pulsed signal into the conductor, then examining the reflection of that pulse. This method of cable analysis is referred to as Time Domain Reflectometry (TDR). By examining the reflection, the RUCKUS device can detect and report cable statistics such as local and remote link pair, cable length, and link status.

VCT Configuration Notes

VCT is not an IEEE standard. VCT uses TDR (Time Domain Reflectometry) to send a signal to a remote partner that loops back to the same port. Different vendors and remote partners may have different techniques, terms, and implementations.

The VCT can be performed when the link partner is autonegotiating. VCT has to be run in autonegotiation and full duplex mode.

VCT supports:

- Copper ports with 2.5G/10G for Aquantia PHY and 1G for Broadcom PHY.
- Link UP ports only. Link DOWN ports are not used.
- 4 pairs per interface, per cable.
- Most types of RJ45 cables including Cat 3, 4, 5, 6 and 7.
- Only on-demand CLI runs for diagnostics. You must trigger the test for a port connected with the cable.

NOTE

Autonegotiation is where common transmission parameters between devices, such as speed, duplex mode, and flow control are chosen. The devices share their parameter capabilities and then choose the highest performance transmission mode they both support.

VCT Restrictions

- The port where the cable is connected must be enabled when you issue the command to diagnose the cable. If the port is disabled, the command is rejected.
- VCT cannot be executed when port speed downshift is configured on the port to downgrade speeds at 10M and 100M. VCT can only be run at default auto speed of the port.
- The length of cable measurement could be different due to different PHY used across the ICX products. The VCT feature must not be used for accurately measuring the length of the cable.
- The Ethernet port speed must be configured to Auto; VCT does not work on ports with fixed speeds.
- If the remote pair is set to forced 100 Mbps, any change in MDI/MDIX may cause the device to interpret the Multilevel Threshold-3 (MLT-3) as a reflected pulse, in this case, the device will report a faulty condition.

NOTE

In this scenario, we recommend that you run the TDR test a few times, clearing the registers before each test.

- You should not run VCT commands in a live network environment. VCT commands may impact port up, down, and network performance.
- Cat 3 or 4 with 2 or 3 pairs may fail in 1 or 2 pairs.
- We do not recommend that you run VCT commands when adjacent ports are up.
- Fiber ports and 1G transmit media are not supported.
- Link DOWN ports are not used.
- VCT commands do not apply to the management port
- There is no support for configuration.
- No trunk and system-wide support.

Crosstalk Between Ports

A maximum PGA gain setting and a disabled echo canceler are required for cable diagnostics, but the side effect is a high sensitivity to crosstalk. Reducing the gain and disabling the echo canceler remove the crosstalk sensitivity. However, this breaks the cable test, leading to invalid TDR results.

When there is crosstalk between the ports, running VCT on the port will provide invalid results. This can be detected using the **show cable tdr** command, as illustrated in the following example.

NOTE

In the following example, port 1/1/48 has 2 devices, the first has 24 ports (1/1/1 to 1/1/24) and the next has 24 ports (1/1/25 to 1/1/48).

```
device# show cable tdr 1/1/48
```

Port	Speed	Local pair	Pair Length	Remote pair	Pair status
1/1/48	1000M	Pair A	Unknown		Invalid
		Pair B	Unknown		Invalid
		Pair C	Unknown		Invalid
		Pair D	Unknown		Invalid

To avoid crosstalk and to run the VCT successfully with consistent results, it is recommended that you disable the adjacent port and stop the traffic on the port where line rate traffic is passed.

In the following example, when you disable port 1/1/47 and stop traffic on port 1/1/48 you get the following results:

```
device# show cable tdr 1/1/48
```

Port	Speed	Local pair	Pair Length	Remote pair	Pair status
1/1/48	1000M	Pair A	<50M	Pair B	terminated
		Pair B	<50M	Pair A	terminated
		Pair C	<50M	Pair D	terminated
		Pair D	<50M	Pair C	terminated

The pair status "terminated" indicates an active port.

In some cases, the Power over Ethernet (PoE) port in ICX 7450 and ICX 7250 devices encounter crosstalk in single ports. An example when crosstalk is seen in a single port shows the following:

```
device# show cable tdr 1/1/13
```

Port	Speed	Local pair	Pair Length	Remote pair	Pair status
1/1/13	1000M	Pair A	<50M	Pair B	terminated
		Pair B	<50M	Pair A	terminated
		Pair C	<=5M		crosstalk
		Pair D	<=5M		crosstalk

Diagnosing a Cable using Time Domain Reflectometry

The following task diagnoses a cable using Time Domain Reflectometry (TDR).

1. Clear any previous TDR test results on the specified port.

```
device# clear cable-diagnostics tdr 1/2/3
```

It is recommended that you clear the TDR test registers before each test.

2. Run the VCT TDR test on the specified port.

```
device# phy cable-diagnostics tdr 1/1/1
```

When you use the **phy cable-diagnostics** command, the command brings the port down for a second or two, and then immediately brings the port back up.

The following example clears previous TDR test results for a port and runs the the VCT TDR test on the specified port.

```
device# clear cable-diagnostics tdr 1/2/3
device# phy cable-diagnostics tdr 1/1/1
```

Viewing the Results of the Cable Analysis

You can use the **show cable-diagnostics** command to display the results of Virtual Cable Test (VCT) TDR cable diagnostic testing.

In the following example, the command displays TDR test results for port 1, slot 1 on device 1 in the stack. The results indicate that the port is down or the cable is not connected.

```
device# show cable-diagnostics tdr 1/1/1
Port      Speed Local pair Pair Length Remote pair Pair status
-----
01        UNKWN Pair A    <=3 M                Open
          Pair B    <=3 M                Open
          Pair C    <=3 M                Open
          Pair D    <=3 M                Open
```

In the following example, the TDR test results for the same port show details for an active port.

```
device# show cable-diagnostics tdr 1/1/1
Port      Speed Local pair Pair Length Remote pair Pair status
-----
01        1000M Pair A    <50M                Pair B    Terminated
          Pair B    <50M                Pair A    Terminated
          Pair C    <50M                Pair D    Terminated
          Pair D    <50M                Pair C    Terminated
```


Local pair indicates the assignment of wire pairs from left to right, where Pair A is the left-most pair. The following table shows the Local pair mapping to the T568A pin/pair and color assignment from the TIA/EIA-568-B standard.

TABLE 7 Local Pair Definition

Local Pair	T568A Pair and Color Assignment
Pair A	Pair 3 (green)
Pair B	Pair 2 (orange)
Pair C	Pair 1 (blue)
Pair D	Pair 4 (brown)

The following figure illustrates the T568A pin/pair assignment.

FIGURE 1 T568A Pin/Pair Assignment



Digital Optical Monitoring

Digital optical monitoring (DOM) provides a diagnostic monitoring interface for SFP and SFP+ optics. DOM supports monitoring of optical output power, optical input power, temperature, laser bias current, and transceiver voltage.

You can configure your RUCKUS device to monitor optical transceivers in the system, either globally or by specified ports. When DOM is enabled, the system monitors the temperature and signal power levels for the optical transceivers in the specified ports. Console messages and syslog messages are triggered when optical operating conditions fall below or rise above the SFP, SFP28, SFP+, QSFP, QSFP+, and QSFP28 manufacturer-recommended thresholds.

NOTE

DOM is supported on RUCKUS optics. DOM is supported on all the ICX switches.

For a list of supported media types, refer to the [Ruckus Ethernet Optics data sheet](#).

Beginning with FastIron release 08.0.95, new RUCKUS optic has been introduced for ICX 7650 and ICX 7550. The new 100G-ER4 (Extended Reach 4 Lanes) or 100G-ER4-Lite optic do not support stacking ports. They support 40km transmission with Forward-Error-Correction (FEC) and 30km transmission without FEC.

NOTE

The 100G-ER4 optic support uplink ports only.

NOTE

DOM cannot be enabled for user disabled ports. When DOM is configured globally but a port is disabled by the user, the following appears in the output: "Optical monitoring is in progress, please try later".

The **show media ethernet** command is used to display information about the media device installed in a port.

```
ICX7550-24F Switch# show media ethernet 1/1/1
Port 1/1/1: Type : 100GBASE-ER4 40km (QSFP28)
Vendor: RUCKUS Version: A
Part# : 57-1000489-01 Serial#: YFK119452000012
```

Display the optical monitoring information using the **show optic** command.

```
ICX7550-24F Switch# show optic 1/2/1

                100GBASE-ER4
                =====
Port  Temperature      Voltage      Tx Power      Rx Power      Tx Bias Current
+-----+-----+-----+-----+-----+-----+
1/2/1  31.7773 C      3.1911 volts  000.7309 dBm  -015.6224 dBm  59.932 mA
      Normal      Normal      Normal      Normal      Normal

Chan Rx Power #1  Rx Power #2  Rx Power #3  Rx Power #4
+-----+-----+-----+-----+
      -015.6224 dBm  -015.3165 dBm  -015.1427 dBm  -015.2143 dBm
      Normal      Normal      Normal      Normal

Chan Tx Bias #1  Tx Bias #2  Tx Bias #3  Tx Bias #4
+-----+-----+-----+-----+
      59.932 mA      57.544 mA      60.216 mA      59.932 mA
      Normal      Normal      Normal      Normal

Chan Tx Power #1  Tx Power #2  Tx Power #3  Tx Power #4
+-----+-----+-----+-----+
      000.7309 dBm  000.4575 dBm  000.5530 dBm  000.4626 dBm
      Normal      Normal      Normal      Normal
```

The 100G-PAM4 optic supports RUCKUS ICX 7850, ICX 7650 and ICX 7550 with new 100G or 200G or 400G standard. They support uplink ports only. The maximum transmission distance for 100G-DR or 100G-DR1 is 500m. DR1 represents one lane with 100G. The maximum transmission distance for 100G-FR or 100G-FR1 is 2km. The maximum transmission distance for 100G-LR or 100G-LR1 is 10km.

```
ICX7850-32Q Router# show media ethernet 1/3/5
Port 1/3/5: Type : 100GBASE-DR PAM4 500m (QSFP28)
Vendor: RUCKUS Version: A
Part# : 57-1000488-01 Serial#: YTK11950T000003
```

```
ICX7850-32Q Router# show optic 1/3/5

                100GBASE-DR
                =====
Port  Temperature      Voltage      Tx Power      Rx Power      Tx Bias Current
+-----+-----+-----+-----+-----+-----+
1/3/5  31.0742 C      3.1667 volts  001.0157 dBm  001.6725 dBm  72.000 mA
      Normal      Normal      Normal      Normal      Normal
```

DOM Show and Configuration Commands

The following commands are associated with DOM:

- **optical-monitor**: Allows users to configure all ports (system-wide), a range of ports, or a single port, for monitoring.
- **show lrm_adapter ethernet**: Allows users to display the LRM adapter parameters.
- **show media**: Displays information about the media devices installed per device, per slot, and per port.
- **show optic**: Displays the optical monitoring information.
- **show optic thresholds**: Displays the thresholds for a qualified optical transceiver in a particular port.
- **show optic-timer**: Displays the current DOM time interval setting.

For more information on these commands, refer to the *RUCKUS FastIron Command Reference*.

NOTE

The **show media** command and DOM features are supported on LRM adapters. Command output remains the same as that of regular optics. For more information on LRM adapters, refer to the hardware installation guide for the respective product family.

Enabling DOM

Complete the following steps to enable DOM.

NOTE

DOM is supported only on RUCKUS optics.

1. Use the **optical-monitor** command to enable digital optical monitoring, and specify the polling and alarm interval. Enter one of the following versions of the command:

- Enable digital optical monitoring and specify an alarm interval.

```
device(config)# optical-monitor 18
Enable optical monitoring and set alarm/warn interval to 18 minute(s)
```

- Enable digital optical monitoring, without specifying an alarm interval, to set the alarm interval to the default.

```
device(config)# optical-monitor
Enable optical monitoring and set alarm/warn interval to default(8 minutes)
```

The default alarm interval for the ICX 7450, ICX 7550, ICX 7650, and ICX 7850 is 8 minutes. The default alarm interval for the ICX 7250 and ICX 7150 is 3 minutes.

Ports that are down are not included for optical monitoring by default configuration. To allow the monitoring for the down ports on an ICX device, enter the **optical-monitor down-port-enable** command.

2. Use the **interface** command, specifying an interface, to enter interface configuration mode for the specified interface.

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

3. Use the **optical-monitor** command, without specifying a value, to enable optical monitoring on the specified port and set the default polling and alarm interval.

```
device(config-if-e10000-1/1/1)# optical-monitor
```

4. Use the **exit** command to return to global configuration mode.

```
device(config-if-e10000-1/1/1)# exit
```

5. Use the **interface** command, specifying a range of interfaces, to specify a range of interfaces.

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

6. Use the **optical-monitor** command, without specifying a value, to enable optical monitoring on the specified range of ports and set the default polling and alarm interval.

```
device(config-mif-e10000-1/1/1-1/1/2)# optical-monitor
```

7. Use the **exit** command to return to global configuration mode.

```
device(config-mif-e10000-1/1/1-1/1/2)# exit
```


8. Verify the alarm and warning interval.

```
device(config)# show optic-timer 1/1/4
Optical monitoring timer Interval for Port 1/1/4 is 8 mins
```

9. Verify the media device configuration.

a) Display information about the media devices installed per device, per stack, and per port.

```
device(config)# show media
Port 1/1/1:      Type : 1G M-C (Gig-Copper)
Port 1/1/2:      Type : 1G M-C (Gig-Copper)
Port 1/1/3:      Type : 1G M-C (Gig-Copper)
Port 1/1/4:      Type : 1G M-C (Gig-Copper)
Port 1/1/5:      Type : 1G M-C (Gig-Copper)
Port 1/1/6:      Type : 1G M-C (Gig-Copper)
Port 1/1/7:      Type : 1G M-C (Gig-Copper)
Port 1/1/8:      Type : 1G M-C (Gig-Copper)
Port 1/1/9:      Type : 1G M-C (Gig-Copper)
...
Port 1/2/1:      Type : 10GE SR 300m (SFP +)
Port 1/2/2:      Type : EMPTY
Port 1/2/3:      Type : 1G Twinax 1m (SFP)
Port 1/2/4:      Type : 1G Twinax 1m (SFP)
```

b) Display information about the media device installed in a port.

```
device(config)# show media ethernet 1/1/17
Port 1/1/17: Type : 1GE M-SX(SFP)
Vendor: Ruckus Networks. Version: A
Part# : 33210-100 Serial#: TAA11106M3GV
```

c) Verify if your optics are official RUCKUS optics or another brand.

```
device# show media validation ethernet 1/3/3
Port      Supported   Vendor
Type
-----
1/3/3     Yes             RUCKUS      Type : 10GE LR 10km (SFP+)
```

10. As a test, check the optical statistics for any enabled port.

```
device(config)# show optic 2/1/1
Port  Temperature      Tx Power      Rx Power      Tx Bias Current
+-----+-----+-----+-----+-----+
2/1/1  32.2578 C    -002.5157 dBm  -002.8091 dBm  5.966 mA
Normal      Normal      Normal      Normal
```

11. Verify the optic warning and alarm thresholds for any enabled port.

```
device(config)# show optic thresholds 2/1/1
```

DOM Configuration Example

The following example shows a complete configuration and verification of digital optical monitoring.

```
device(config)# optical-monitor 8 >>>> Global
Enable optical monitoring and set alarm/warn interval to 8 minute(s)
device(config)# interface ethernet 1/2/1 >>>> For a specific port
```

Hardware Component Monitoring

Syslog Messages for Optical Transceivers

```
device(config-if-e10000-1/2/1)# optical-monitor

device(config)# show optic-timer 1/1/4
Optical monitoring timer Interval for Port 1/1/4 is 8 mins

device(config)# show media >>>>>> Global
Port 1/1/1:      Type : 1G M-C (Gig-Copper)
Port 1/1/2:      Type : 1G M-C (Gig-Copper)
Port 1/1/3:      Type : 1G M-C (Gig-Copper)
...
Port 1/2/1:      Type : 10GE SR 300m (SFP +)
Port 1/2/2:      Type : 10GE      Twinax 1m (SFP +)
Port 1/2/3:      Type : 1G Twinax 1m (SFP)
Port 1/2/4:      Type : 1G Twinax 1m (SFP)

device(config)# show media ethernet 1/2/1 >>>>> For a specific port
Port 1/2/1: Type : 1GE M-SX(SFP)
Vendor: Ruckus Networks. Version: A
Part# : 33210-100 Serial#: TAA11106M3GV

device(config)# show media validation
Port      Supported      Vendor      Type
-----
1/2/1     Yes                FINISAR CORP.  1GE M-SX(SFP)
1/2/2     Yes                10GE      Twinax 1m (SFP +)
2/2/1     Yes                10GE      SR 300m (SFP +)
2/2/3     Yes                10GE      SR 300m (SFP +)

device(config)# show optic 2/1/1
Port      Temperature      Tx Power      Rx Power      Tx Bias Current
+-----+-----+-----+-----+-----+
2/1/1     32.2578 C      -002.5157 dBm  -002.8091 dBm  5.966 mA
Normal    Normal          Normal        Normal          Normal

device(config)# show optic thresholds 1/3/1
Port 1/3/1 sfp monitor thresholds:
Temperature High alarm      5d00      93.0000 C
Temperature Low alarm       f300      -13.0000 C
Temperature High warning    5800      88.0000 C
Temperature Low warning     f800      -8.0000 C
Supply Voltage High alarm   9088      3.7000 Volts
Supply Voltage Low alarm    7148      2.9000 Volts
Supply Voltage High warning 8ca0      3.6000 Volts
Supply Voltage Low warning  7530      3.0000 Volts
TX Bias High alarm         170c      11.0800 mA
TX Bias Low alarm          07d0      4.0000 mA
TX Bias High warning       1518      10.0800 mA
TX Bias Low warning        09c4      5.0000 mA
TX Power High alarm        207e      -000.7998 dBm
TX Power Low alarm         09d0      -005.9998 dBm
TX Power High warning     19cf      -001.7999 dBm
TX Power Low warning      0c5a      -005.0003 dBm
RX Power High alarm        2710      000.0000 dBm
RX Power Low alarm         0064      -020.0000 dBm
RX Power High warning     1f07      -001.0001 dBm
RX Power Low warning      009e      -018.0134 dBm
```

Syslog Messages for Optical Transceivers

The system generates syslog messages for optical transceivers in the following circumstances:

- The temperature, supply voltage, TX bias, TX power, or TX power value goes above or below the high or low warning or alarm threshold set by the manufacturer.
- The optical transceiver does not support digital optical monitoring.
- The optical transceiver is not qualified, and therefore not supported by RUCKUS.

For details about the above syslog messages, refer to [Syslog Message Descriptions](#) on page 113.

Port Mirroring and Monitoring

- Port Mirroring and Monitoring Overview.....61
- Port Mirroring and Monitoring Configuration..... 61
- Mirroring Configuration on a Traditional Stack..... 63
- ACL-based Inbound Mirroring..... 64
- MAC ACL Mirroring.....67
- VLAN-based Mirroring.....68
- Remote Switched Port Analyzer..... 69
- Encapsulated Remote Switched Port Analyzer (ERSPAN) 73

Port Mirroring and Monitoring Overview

Port mirroring is a method of monitoring network traffic that forwards a copy of each incoming or outgoing packet from one port on a network switch to another port where the packet can be analyzed. Port mirroring can be used as a diagnostic tool or debugging feature, especially for preventing attacks. Port mirroring can be managed locally or remotely.

You can configure port mirroring, by assigning a port (known as the Monitor port), from which the packets are copied and sent to a destination port (known as the Mirror port). All packets received on the Monitor port or issued from it, are forwarded to the second port. You next attach a protocol analyzer on the mirror port to monitor each segment separately. The analyzer captures and evaluates the data without affecting the client on the original port.

The mirror port may be a port on the same switch with an attached RMON probe, a port on a different switch in the same hub, or the switch processor.

Port Mirroring and Monitoring Configuration

When configuring port monitoring, the mirror port must be specified before enabling monitoring on the monitored port.

The *mirror port* is the port to which the monitored traffic is copied. Attach your protocol analyzer to the mirror port. The monitored port is the port with the traffic you want to monitor.

The following table lists the number of mirror and monitor ports supported on the RUCKUS devices.

TABLE 8 Number of Mirror and Monitor Ports Supported

Maximum number supported	
Port Type	ICX (7150, 7450, 7250, 7650, 7850 and 7550)
Ingress mirror ports	1 per port region
Egress mirror ports	1 per port region
Ingress monitored ports	No limit
Egress monitored ports	20

NOTE

Region refers to the number of devices (BCM chipset) in a unit. You can configure more than eight egress ports, although only the first eight are operational. This is also true for mirrored VLANs - more than eight can be configured, but only the first eight are operational.

Configuration Notes for Port Mirroring and Monitoring

Refer to the following guidelines when configuring port mirroring and monitoring:

- If you configure both ACL mirroring and ACL-based rate limiting on the same port, all packets that match are mirrored. This includes the packets that exceed the rate limit.
- sFlow and port monitoring are supported together on the same port.
- Mirror ports can be configured specifically as an ingress port, an egress port, or both.
- Mirror ports can run at any speed, and are not related to the speed of the ingress or egress monitored ports.
- The same port cannot be both a monitored port and the mirror port.
- The same port can be monitored by one mirror port for ingress traffic and another mirror port for egress traffic.
- The mirror port cannot be a trunk port.
- The monitored port and its mirror port do not need to belong to the same port-based VLAN:
 - If the mirror port is in a *different* VLAN from the monitored port, the packets are tagged with the monitor port VLAN ID.
 - If the mirror port is in the *same* VLAN as the monitored port, the packets are tagged or untagged, depending on the mirror port configuration.
- More than one monitored port can be assigned to the same mirror port.
- If the LAG virtual interface is enabled for monitoring, the entire LAG is monitored. You can also enable an individual member ports of a LAG for monitoring using the **monitor** command from the LAG configuration mode.
- For *stacked* devices, if the ingress and egress analyzer ports are always network ports on the local device, each device may configure the ingress and egress analyzer port independently. However, if you need to mirror to a remote port, then only one ingress and one egress analyzer port are supported for the entire system.
- For ingress ACL mirroring, the ingress rule for stacked devices also applies. The analyzer port setting command **acl-mirror-port** must be specified for each port, even though the hardware only supports one port per device. This applies whether the analyzer port is on the local device or on a remote device. For example, when port mirroring is set to a remote device, any mirroring-enabled ports (ACL, MAC address filter, or VLAN) enabled ports are set globally to a single analyzer port, as shown in the following example.

```
device(config)# mirror ethernet 1/1/24
device(config)# mirror ethernet 2/1/48
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# monitor ethernet 2/1/48 both
```

The analyzer port (2/1/48) is set to all devices in the system.

```
device(config)# interface ethernet 1/1/2
device(config-if-e1000-1/1/2)# ip access-group 101 in
device(config-if-e1000-1/1/2)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# acl-mirror-port ethernet 2/1/48
```

The previous command is required even though the analyzer port is already set globally by the port mirroring command.

```
device(config)# interface ethernet 1/1/3
device(config-if-e1000-1/1/3)# ip access-group 101 in
device(config-if-e1000-1/1/3)# acl-mirror-port ethernet 2/1/48
device(config-if-e1000-1/1/3)# ip access-group 102 in
```

Commands for Port Mirroring and Monitoring

This section describes how to configure port mirroring and monitoring.

Monitoring a Port

To configure port monitoring on an individual port on a device, complete the following steps.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter the **mirror-port** command, specifying an Interface, to configure the mirror port that functions as the destination port to which the packets that need to be monitored are copied and forwarded.

```
device(config)# mirror-port ethernet 1/2/4
```

3. Enter interface configuration mode to configure port monitoring.

```
device(config)# interface ethernet 1/2/11
```

4. Configure port monitoring to send the traffic from the monitored port to the specified mirror port.

```
device(config-if-e1000-1/2/11)# monitor ethernet 1/2/4 both
```

Monitoring an Individual LAG Port

You can monitor the traffic on an individual port of a static LAG group, and on an individual port of an LACP LAG group.

By default, when you monitor the LAG virtual interface, aggregated traffic for all the ports in the LAG is copied to the mirror port. You can configure the device to monitor individual ports in a LAG as well.

To configure port monitoring on an individual port in a LAG, enter commands such as the following.

```
device(config)# lag automation static id 1
device(config-lag-automation)# ports ethernet 1/1/2 to 1/1/9
device(config-lag-automation)# exit
device(config)# mirror-port ethernet 1/1/1
device(config)# lag automation
device(config-lag-automation)# monitor ethe-port-monitored 1/1/2 ethernet 1/1/1 both
```

```
device# show mirror
Mirror port 1/1/1
  Input monitoring      : (U1/M1)   1
  Output monitoring    : (U1/M1)   1
```

```
device# show mirror ethernet 1/1/1
Mirror port 1/1/1
  Input monitoring      : (U1/M1)   1
  Output monitoring    : (U1/M1)   1
```

```
device# show running-config | i mirror
mirror-port ethernet 1/1/1
```

```
device# show running-config | i monitor ethernet
monitor ethe-port-monitored 1/1/2 ethe 1/1/1 both
```

Traffic on LAG port e 1/1/2 is monitored, and the monitored traffic is copied to port e 1/1/1, the mirror port.

Mirroring Configuration on a Traditional Stack

You can configure mirroring on a RUCKUS traditional stack. A traditional stack consists of up to twelve FastIron devices of the same type. The stack operates as a chassis. The following examples show how to configure mirroring for ports that are on different members of a stack, and for ports that are on the same stack member as the mirror port.

Configuration Notes for Traditional Stack Mirroring

The following mirroring configuration information applies to FastIron devices connected in a traditional stack topology:

- The input or output mirroring port can be on different ports.
- All FastIron devices can have one mirroring port that monitors multiple ports, but cannot have multiple mirror ports for one monitored port.
- If the mirror port and the monitored ports are on different stack units, only one active mirror port is allowed for the entire traditional stack.
- If the mirror port and the monitored ports are on the same port region, multiple active mirror ports are allowed for the entire traditional stack. Devices in a traditional stack support 24 ports per port region.
- The maximum number of monitored VLANs on a traditional stack is 8.

Configuring Mirroring for Ports on Different Members in a Traditional Stack Example

In this example, although two ports are configured as active ports, only one active mirror port (port 1/1/24) is allowed for the entire stack because the mirror ports and the monitored ports are on different stack members.

```
device(config)# mirror-port ethernet 1/1/24
device(config)# mirror-port ethernet 2/1/24
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# monitor ethernet 1/1/24 both
device(config-if-e1000-1/1/1)# exit
device(config)# interface ethernet 2/1/1
device(config-if-e1000-2/1/1)# monitor ethernet 1/1/24 both
device(config-if-e1000-2/1/1)# exit
device(config)# interface ethernet 4/1/1
device(config-if-e1000-4/1/1)# monitor ethernet 1/1/24 both
```

Configuring Mirroring for Ports on the Same Stack Member in a Traditional Stack Example

In this example, the mirror ports are assigned to different monitor ports.

```
device(config)# mirror-port ethernet 1/1/24
device(config)# mirror-port ethernet 2/1/24
device(config)# mirror-port ethernet 3/1/24
device(config)# mirror-port ethernet 4/1/24
device(config)# interface ethernet 1/1/1
device(config-if-e1000-1/1/1)# monitor ethernet 1/1/24 both
device(config-if-e1000-1/1/1)# exit
device(config)# interface ethernet 2/1/1
device(config-if-e1000-2/1/1)# monitor ethernet 2/1/24 both
device(config-if-e1000-2/1/1)# exit
device(config)# interface ethernet 4/1/1
device(config-if-e1000-4/1/1)# monitor ethernet 4/1/24 both
```

ACL-based Inbound Mirroring

This section describes ACL-based inbound mirroring for FastIron devices.

Creating an ACL-based Inbound Mirror Clause

The following steps configure an ACL-based inbound mirror clause.

1. Configure the mirror port.

```
device(config)# mirror-port ethernet 1/1/2
```

2. Configure the ACL-based inbound mirror clause.

```
device(config)# ip access-list extended 101  
device(config-ext-ipacl-101)# permit ip any any mirror
```

3. Apply the ACL-based inbound clause to the monitor port.

```
device(config)# interface ethernet 1/1/5  
device(config-if-e1000-1/1/5)# ip access-group 101 in
```

4. Create the ACL mirror port.

```
device(config-if-e1000-1/1/5)# acl-mirror-port ethernet 1/1/2
```

5. (Optional) Use the **show access-list all** command to verify ACL mirror settings.

```
device(config-ext-ipacl-101)# show access-list all  
Extended IP access list 101: 1 entries  
10: permit ip any any mirror
```

Destination mirror port

You can specify physical ports or a trunk to mirror traffic. If you complete the rest of the configuration but do not specify a destination mirror port, the port-mirroring ACL is non-operational. This can be useful if you want to be able to mirror traffic by a set criteria on demand. With this configuration, you configure a destination mirror port whenever you want the port-mirroring ACL to become operational.

The following sections describe how to specify a destination port for a port or a trunk, as well as the special considerations required when mirroring traffic from a virtual interface.

Specifying the Destination Mirror Port for Physical Ports

When you want traffic that has been selected by ACL-based inbound mirroring to be mirrored, you must configure a destination mirror port. This configuration is performed from the interface configuration mode of the port with the traffic you are mirroring.

In the following example, ACL mirroring traffic from port 1/1/1 is mirrored to port 1/1/3.

```
device(config)# interface ethernet 1/1/1  
device(config-if-e1000-1/1)# acl-mirror-port ethernet 1/1/3
```

In a single PP device, if the same user ACL is applied on multiple ports and `acl-mirror-port` is applied on one port, the mirroring will be done on all ports where the same user ACL is applied.

In the case of a dual-PP device, if the user ACL is applied on port of both PP and `acl-mirror-port` is applied on one port of one PP, then mirroring will be done on only one PP.

If the same ACL is applied on multiple interfaces, it is recommended to configure the same `ACL-mirror-port` on all ports on interfaces where the user ACL is applied.

Specifying the Destination Mirror Port for LAG Ports

You can mirror the traffic that has been selected by ACL-based inbound mirroring from a LAG by configuring a destination port for the LAG virtual interface within the LAG configuration, as shown in the following example.

```
device(config)# lag blue static id 1
device(config-lag-blue)# ports ethernet 1/1/1 to 1/1/4
device(config)# interface lag 1
device(config-lag-if-lg1)# acl-mirror-port ethernet 1/1/8
```

Using this configuration, all LAG traffic is mirrored to port 1/1/8.

Limitations when Configuring ACL-based Mirroring with LAGs

If an individual port is configured for ACL-based mirroring, you cannot add it to a LAG. If you try to add a port that is configured for ACL-based mirroring to a LAG, the following message appears.

```
Note - ACL-mirror-port configuration is removed from port 2 in new trunk.
```

NOTE

If you want to add a port configured for ACL-based mirroring to a LAG, you must first remove the **acl-mirror-port** command from the port configuration. You can then add the port to a LAG that can then be configured for ACL-based LAG mirroring.

Behavior of ACL-based Mirroring when Deleting LAGs

If you delete a LAG, the ACL-based mirroring configuration applied on the LAG interface also will get removed.

Configuring ACL-based Mirroring for ACLs Bound to Virtual Interfaces

For configurations that have an ACL configured for ACL-based mirroring bound to a VLAN, you must use the **acl-mirror-port** command on a physical port that is a member of the same VLAN. Additionally, only traffic that arrives at ports that belong to the same port group as the physical port where this command has been used is mirrored. This follows the same rules described in [Specifying the Destination Mirror Port for Physical Ports](#) on page 65.

For example, in the following configuration, ports 1/1/7, 1/1/8, and 1/1/26 are in VLAN 222. Ports 1/1/7 and 1/1/8 belong to the same port group, while port 1/1/26 belongs to another port group.

```
device# configure terminal
device(config)# vlan 222
device(config-vlan-222)# tagged ethernet 1/1/7 to 1/1/8
device(config-vlan-222)# tagged ethernet 1/1/26
device(config)# interface ethernet 1/1/7
device(config-if-e10000-1/1/7)# acl-mirror-port ethernet 1/1/1
device(config-if-e10000-1/1/7)# exit

device(config)# ip access-list extended mirror-1
device(config-ext-ipacl-mirror-1)# permit ip any any mirror
device(config-ext-ipacl-mirror-1)# exit

device(config)# vlan 222
device(config-vif-222)# ip access-group mirror-1 in
Extended IP access list mirror-1: 1 entries
10: permit ip any any mirror
```

In this configuration, the **acl-mirror-port** command is applied to port 1/1/7, which is a member of VLAN 222. Because of this, ACL-based mirroring will only apply to VLAN 222 traffic that arrives on ports 1/1/7 and 1/1/8. It will not apply to VLAN 222 traffic that arrives on port 1/1/26 because that port belongs to a port group different from ports 1/1/7 and 1/1/8. This is because if you apply ACL-based mirroring on an entire VLAN, and enable mirroring in only one port region, traffic that is in the same VLAN but on a port in a different port region will not be mirrored.

To make the configuration apply ACL-based mirroring to VLAN 222 traffic arriving on port 1/1/26, you must add the following commands to the configuration.

```
device(config)# interface ethernet 1/5/3
device(config-if-e10000-1/5/3)# acl-mirror-port ethernet 1/1/1
```

If a port is in both mirrored and non-mirrored VLANs, only traffic on the port from the mirrored VLAN is mirrored. For example, the following configuration adds VLAN 225 to the previous configuration. In this example, ports 1/1/7 and 1/1/8 are in both VLAN 222 and VLAN 225. ACL-based mirroring is only applied to VLAN 222. Consequently, traffic that is on ports 1/1/7 and 1/1/8 that belongs to VLAN 225 will not be mirrored.

```
device(config)# vlan 222
device(config-vlan-222)# tagged ethernet 1/1/7 to 1/1/8
device(config-vlan-222)# tagged ethernet 1/1/26

device(config)# vlan 225
device(config-vlan-225)# tagged ethernet 1/1/7 to 1/1/8
device(config)# interface ethernet 1/1/7

device(config-if-e10000-1/1/7)# acl-mirror-port ethernet 1/1/1
device(config-if-e10000-1/1/7)# exit

device(config)# ip access-list extended mirror-1
device(config-ext-ipacl-mirror-1)# permit ip any any mirror
device(config-ext-ipacl-mirror-1)# exit

device(config)# vlan 222
device(config-vif-222)# ip access-group mirror-1 in

Extended IP access list mirror-1: 1 entries
10: permit ip any any mirror
```

MAC ACL Mirroring

Traffic entering an ingress port can be monitored from a mirror port connected to a data analyzer, based on specific source and destination MAC addresses. This feature supports mirroring of inbound traffic only. Outbound mirroring is not supported.

MAC ACL allows a user to specify a particular stream of data for mirroring using a filter. This eliminates the need to analyze all incoming data to the monitored port. To configure ACL mirroring, the user must perform three steps:

1. Define a mirror port
2. Create a MAC ACL with a mirroring clause
3. Apply the MAC ACL on an interface

MAC ACL Configuration Notes

- If there is no input mirror port configured, MAC ACL does not take effect. It remains in the configuration, but is not activated.
- MAC ACL can be enabled on a port at the same time as either port-based mirroring or VLAN-based mirroring. When port-based mirroring and MAC ACL mirroring are enabled on a port at the same time, the preference order is port-based mirroring followed by MAC ACL. When VLAN-based mirroring and MAC ACL mirroring are enabled on a port at the same time, the preference order is VLAN-based mirroring and MAC ACL mirroring.
- Port-based mirroring and MAC-filter-based mirroring are enabled on a port at the same time, the preference order is port-based mirroring followed by MAC-based filtering. When VLAN-based mirroring and MAC-filter-based mirroring are enabled on a port at the same time, the preference order is VLAN-based mirroring and MAC-filter-based mirroring.
- Port-based mirroring and VLAN-based mirroring can not be enabled on a port at the same time.

Configuring MAC ACL Mirroring

1. Enter global configuration mode.

```
device# configure terminal
```

2. Activate mirroring on the port by using the **mirror** command.

```
device(config)# mirror ethernet 1/1/1
```

3. Create MAC ACL with mirror option.

```
device(config)# permit 0000.0011.2222 ffff.ffff.ffff 0000.0022.3333 ffff.ffff.ffff mirror
```

The keyword is added to MAC address filter clauses to direct desired traffic to the mirror port. In the following example, the MAC address filter directs traffic to a mirror port. In this example, any flow matching the source address (SA) 0000.0011.2222 and the destination address (DA) 0000.0022.3333 is mirrored. Other flows are not mirrored.

4. Apply the MAC ACL to the interface.

```
device(config)# interface ethernet 1/1/11  
device(config-if-e10000-1/1/11)# mac access-group MAC_ACL in
```

5. Configure the monitor port to use the mirror port.

```
device(config)# interface ethernet 1/1/11  
device(config-if-e10000-1/1/11)# acl-mirror-port ethernet 1/1/1
```

VLAN-based Mirroring

The VLAN-based mirroring feature allows users to monitor all incoming traffic in one or more VLANs by sending a mirror image of that traffic to a configured mirror port. This feature meets the requirements of CALEA (Communications Assistance for Law Enforcement Act of 1994).

Configuration Notes for VLAN-based Mirroring

The following guidelines apply to VLAN-based mirroring configurations:

- A VLAN must have at least one port member configured before monitoring can be configured.
- Multiple VLANs can have monitoring enabled at the same time, and the maximum number of monitor-configured VLANs is 8.
- The mirror port is subject to the same scheduling and bandwidth management as the other ports in the system. If the amount of traffic being sent to the mirror port exceeds the available bandwidth, some of that traffic may be dropped.
- All incoming traffic (tagged and untagged) in the VLAN is mirrored. mirroring is "as-is", and is not affected by the configuration of the mirror port itself. Incoming tagged traffic is sent out tagged and incoming untagged traffic is sent out untagged, regardless of which VLANs the mirror port belongs to, and whether the mirror port is tagged or untagged.
- VLAN-based mirroring is supported on Layer 2 and Layer 3 images.

Configuring VLAN-based Mirroring

The following example enables mirroring on VLANs 10 and 20, to mirror port e 1/1/21.

```
device(config)# mirror-port ethernet 1/1/21 input  
device(config)# vlan 10  
device(config-VLAN-10)# monitor ethernet 1/1/21  
device(config-VLAN-10)# exit
```

```
device(config)# vlan 20
device(config-VLAN-20)# monitor ethernet 1/1/21
device(config-VLAN-20)# end
```

The following example disables mirroring on VLAN 20.

```
device(config)# vlan 20
device(config-VLAN-20)# no monitor ethernet 1/1/21
device(config-VLAN-20)# end
```

Displaying VLAN-based Mirroring Status

The following example displays the VLAN-based mirroring status.

```
device> show vlan


Total PORT-VLAN entries: 4
Maximum PORT-VLAN entries: 4060
Legend: [Stk=Stack-Unit, S=Slot]
PORT-VLAN 1, Name DEFAULT-VLAN, Priority level0, Spanning tree On
  Untagged Ports: (Stk0/S1)  3  4  5  6  7  8  9  10  11  12  13  14
  Untagged Ports: (Stk0/S1)  15 16 17 18 19 20 21 22 23 24 25 26
  Untagged Ports: (Stk0/S1)  27 28 29 30 31 32 33 34 35 36 37 38
  Untagged Ports: (Stk0/S1)  39 40 41 42 43 44 45 46 47 48
  Untagged Ports: (Stk0/S2)  1  2
  Tagged Ports: None
  Uplink Ports: None
  DualMode Ports: None
  Mac-Vlan Ports: None
  Monitoring: Disabled
PORT-VLAN 10, Name [None], Priority level0, Spanning tree On
  Untagged Ports: (Stk0/S1)  1
  Tagged Ports: None
  Uplink Ports: None
  DualMode Ports: None
  Mac-Vlan Ports: None
  Monitoring: Enabled
PORT-VLAN 20, Name [None], Priority level0, Spanning tree On
  Untagged Ports: (Stk0/S1)  2
  Tagged Ports: None
  Uplink Ports: None
  DualMode Ports: None
  Mac-Vlan Ports: None
  Monitoring: Disabled
```

Remote Switched Port Analyzer

Remote Switched Port Analyzer (RSPAN) enables remote monitoring of multiple switches across a network. When RSPAN is enabled, a copy of each incoming or outgoing packet from one port on a network switch is forwarded to another port on the same switch where the packet can be analyzed. RSPAN can be used as a diagnostic tool for preventing network attacks.

RSPAN monitors traffic from source ports distributed over multiple switches so that network capture devices can be centralized. The configured source port or ports is mirrored to the RSPAN VLAN, and the ports that are members of this VLAN receive the mirrored traffic. This VLAN is then trunked to other switches, allowing the RSPAN traffic to be transported across multiple switches to the destination port, as illustrated in the following figure. Transmitted, received, or both directions of traffic can be mirrored to the destination interface.


FIGURE 2 Traffic Monitoring using RSPAN



All participating devices must be connected by Layer 2 trunks, and the remote VLAN must be configured on all devices participating in the RSPAN session.

The following figure shows an RSPAN VLAN that is carrying mirrored traffic to the destination port. S1 is the host device, with interface Ethernet 1/1/1 configured as the source port on which incoming traffic is mirrored and tunneled in RSPAN VLAN 4000 through the intermediate device, S2, to S3 where interface Ethernet 1/1/2 is configured as the destination port. Refer to [Configuring RSPAN](#) on page 71 to view the steps for the configuring RSPAN.

FIGURE 3 Sample RSPAN Configuration



The monitored traffic can be configured to all directions of the monitor port. You can configure:

- Ingress traffic only
- Egress traffic only
- Both ingress and egress traffic

RSPAN Feature Limitations and Considerations

The following limitations and considerations apply when configuring RSPAN:

- All participating devices must be connected by Layer 2 trunks.
- Egress and ingress traffic mirroring is supported.
- 20 source ports are supported.

- A source port cannot be a member of the RSPAN VLAN.
- A destination port must be a member of the RSPAN VLAN.
- A source port cannot be configured unless a destination port is already configured.
- Management ports, stack ports, MCT ports, and PE ports are not supported.
- Only one RSPAN VLAN can be configured in a single network.
- There is no limitation on the number of member ports.
- STP and RSTP is supported on the RSPAN VLAN.
- Normal VLAN commands are not applicable.
- MAC learning is disabled on the RSPAN VLAN.
- The RSPAN VLAN must be a non-existent VLAN in a switch and must be the same across the network.
- Any VLAN can be configured as an RSPAN VLAN as long as all participating network devices support the configuration of RSPAN VLANs.
- You must configure the RSPAN VLAN on all source, intermediate, and destination network devices.
- If tagged, outgoing packets carry the RSPAN VLAN 802.1Q tag.
- There is no distinction between forwarded traffic and mirrored traffic at the destination.
- The RSPAN VLAN must be the same for the entire switched system for RSPAN forwarding rules to be followed to carry traffic to the analyzer port.
- Static mac configuration is not allowed based on the VLAN provided.
- The .1Q PRI field in the RSPAN header has a default of 0.
- All packets are HW forwarded with no effect on the CPU.
- Any logical ports checks do not take effect because traffic mirroring happens before all forwarding.
- If the RSPAN VLAN is also used as a forwarding VLAN, each switch in the RSPAN network receives two streams of traffic (one stream of flooded traffic and another stream of mirrored traffic).
- Any incoming packet with an RSPAN VLAN ID is forwarded within the network.
- RSPAN does not support double tagged packets at source.
- Mirrored L2 BPDU, UDLD, Stacking/ZTP/MRP/Ruckus Proprietary MACs are all suppressed at source.
- ISSU is not supported.

Configuring RSPAN

RSPAN enables remote monitoring of multiple devices across a network. The following configuration demonstrates an RSPAN for both ingress and egress traffic.

1. **On the source device**, enter the **configure terminal** command to access global configuration mode.

```
device1# configure terminal
```

2. Enter the **rspan-vlan** command, specifying a VLAN ID, to define an RSPAN VLAN on the source device.

```
device1(config)# rspan-vlan 4000
```

3. Enter the **tagged ethernet** command, specifying an interface, to add a member port.

```
device1(config-rspan-vlan-4000)# tagged ethernet 1/1/2
```

Port Mirroring and Monitoring

Remote Switched Port Analyzer

4. Enter the **rspan destination** command, specifying an interface, to configure the RSPAN destination port.

```
device1(config-rspan-vlan-4000)# rspan destination ethernet 1/1/2
```

5. Enter the **rspan source** command with the **monitor-both** keyword, specifying an interface, to configure the RSPAN source port and specify that both ingress and egress traffic is monitored.

```
device1(config-rspan-vlan-4000)# rspan source monitor-both ethernet 1/1/1
```

6. **On the intermediate device**, enter the **configure terminal** command to access global configuration mode.

```
device2# configure terminal
```

7. Enter the **rspan-vlan** command, specifying a VLAN ID, to define an RSPAN VLAN on the intermediate device.

```
device2(config)# rspan-vlan 4000
```

8. Enter the **tagged ethernet** command with the **ethernet** keyword, specifying an interface, to add member ports.

```
device2(config-rspan-vlan-4000)# tagged ethernet 1/1/1 ethernet 1/1/2
```

9. **On the destination device**, enter the **configure terminal** command to access global configuration mode.

```
device3# configure terminal
```

10. Enter the **rspan-vlan** command, specifying a VLAN ID, to define an RSPAN VLAN on the destination device.

```
device3(config)# rspan-vlan 4000
```

11. Enter the **tagged ethernet** command with the **ethernet** keyword, specifying an interface, to add member ports.

```
device3(config-rspan-vlan-4000)# tagged ethernet 1/1/1 ethernet 1/1/2
```

12. Enter the **rspan destination** command, specifying an interface, to specify the RSPAN destination port.

```
device3(config-rspan-vlan-4000)# rspan destination ethernet 1/1/2
```

The following example configures an RSPAN for ingress and egress traffic.

Source device:

```
device1# configure terminal
device1(config)# rspan-vlan 4000
device1(config-rspan-vlan-4000)# tagged ethernet 1/1/2
device1(config-rspan-vlan-4000)# rspan destination ethernet 1/1/2
device1(config-rspan-vlan-4000)# rspan source monitor-both ethernet 1/1/1
```

Intermediate device:

```
device2# configure terminal
device2(config)# rspan-vlan 4000
device2(config-rspan-vlan-4000)# tagged ethernet 1/1/1 ethernet 1/1/2
```

Destination device:

```
device3# configure terminal
device3(config)# rspan-vlan 4000
device3(config-rspan-vlan-4000)# tagged ethernet 1/1/1 ethernet 1/1/2
device3(config-rspan-vlan-4000)# rspan destination ethernet 1/1/2
```


Encapsulated Remote Switched Port Analyzer (ERSPAN)


ERSPAN allows mirroring of packets across a Layer 3 network. Using ERSPAN, you can encapsulate monitored traffic and send it to an analysis station not directly connected to the switch.

ERSPAN encapsulates mirrored packets using GRE with IP delivery. After a packet has been encapsulated, it is forwarded throughout the Layer 3-routed network across a special Layer 3 tunnel. The data section contains the original mirrored packet.

With ERSPAN, port mirroring, from any port to any port, is enabled regardless of the port type and the modularity of the device.

The following figure shows a typical ERSPAN data flow. In the figure, traffic going into and out of the monitor port (in this case, traffic between Host 2 and Host 3) is also sent to Host 1, across the ERSPAN tunnel.

FIGURE 4 ERSPAN data flow



Port Mirroring and Monitoring

Encapsulated Remote Switched Port Analyzer (ERSPAN)

The monitored traffic can be configured to all possible directions of the monitor port. You can configure ingress traffic only, egress traffic only, or both ingress and egress traffic.

ERSPAN is available only in Layer 3.

ERSPAN Configuration Steps

You must complete the following tasks to enable ERSPAN:

- Configure the ERSPAN profile.
- Configure the monitor port.

ERSPAN feature limitations

- The maximum number of mirroring sessions per device is four.
- VLAN mirroring is not supported.
- In some cases, speed mismatches may prevent mirroring of all traffic.
- You cannot terminate the Generic Routing Encapsulation (GRE) tunnel on an ICX switch. ERSPAN must be terminated on the host/analyzer.
- ERSPAN has not been tested against implementations by other vendors.

Configuring an ERSPAN Profile

An ERSPAN profile defines a tunnel over a Layer 3 network from a router to a remote host. Mirrored packets can then be sent to this remote host.

The router must have a configured IP on at least one of the interfaces.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Create an ERSPAN profile and assign it a number.

```
device(config)# monitor-profile 1 type erspan
```

This command puts you in monitor-profile mode.

3. Enter the IP address of the source router.

```
device(config-monitor-profile 1)# source-ip 10.1.1.1
```

The IP address can be any IP on the router.

4. Enter the IP address of the destination host.

```
device(config-monitor-profile 1)# destination-ip 1.1.1.1
```

The IP address is for the host that is collecting the mirrored traffic, not the device.

5. Exit monitor-profile mode.

```
device(config-monitor-profile 1)# exit
```

6. Verify the configuration.

```
device(config)# show erspan profile 1
Profile 1
Type                ERSPAN
Mirror destination  reachable.*/Error condition - Mirror destination Not reachable/*
Destination IP      10.1.1.100
Destination MAC     0000.0000.0000
Source IP           10.1.1.1
Source MAC          cc4e.0000.0000
Ports monitored:
  Input monitoring   : (U1/M1)  1
  Output monitoring  : (U1/M1)  1
HW destination id for each device:
stack_id/device:dest_id
```

If `Mirror destination Not reachable.` appears in the output, see the section [Troubleshooting ERSPAN reachability errors](#) on page 75.

ERSPAN profile configuration example

```
device# configure terminal
device(config)# monitor-profile 1 type erspan
device(config-monitor-profile 1)# source-ip 10.1.1.1
device(config-monitor-profile 1)# destination-ip 10.1.1.100
device(config-monitor-profile 1)# exit
device(config)# show erspan profile 1
```

Next, you need to configure the monitor port.

Troubleshooting ERSPAN reachability errors

Follow these examples to troubleshoot and resolve ERSPAN destination not reachable errors.

```
device(config)# show erspan profile 1
Profile 1
Type                ERSPAN
Mirror destination Not reachable.
Reason: *****
Destination IP      10.1.1.100
Destination MAC     0000.0000.0000
Source IP           10.1.1.1
Source MAC          cc4e.0000.0000
Ports monitored:
  Input monitoring   : (U1/M1)  1
  Output monitoring  : (U1/M1)  1
HW destination id for each device:
stack_id/device:dest_id
```

There are seven reasons why a `Mirror destination Not reachable` error occurs.

Case	Reason
1	ARP is not resolved.
2	Route does not exist.
3	Outgoing port is a management port.
4	Outgoing port is a loopback port.
5	Outgoing port is a GRE IP tunnel port.
6	Outgoing port is not known.
7	Outgoing port is a sink port.

Port Mirroring and Monitoring

Encapsulated Remote Switched Port Analyzer (ERSPAN)

Follow these examples to resolve the errors.

Case 1: Problem ARP not resolved

```
device(config-vif-66)# show erspan

Profile 1
Mirror destination Not reachable.
Reason: ARP not resolved
Destination IP    10.10.10.4
Destination MAC   0000.0000.0000
Source IP         10.10.10.1
Source MAC        0000.0000.0000
Ports monitored:
HW destination id for each device:
stack_id/device:dest_id
```

Solution - Configure ARP

```
device(config-vif-66)# show arp

Total number of ARP entries: 1
Entries in default routing instance:
No.   IP Address      MAC Address      Type      Age Port      Status
1     10.10.10.4       None             Dynamic   1    v66        Pend

device(config-vif-66)# arp 10.10.10.4 aa.bb.cc ethernet 1/1/5

ADD static arp 10.10.10.4 -> 00aa.00bb.00cc -> 1/1/5 (VRF: 0)

device(config-vif-66)# show erspan profile all

Profile 1
Type          ERSPAN
Mirror destination Reachable.
Destination IP    10.10.10.4
Destination MAC   00aa.00bb.00cc
Source IP         10.10.10.1
Source MAC        748e.f8f9.6d80
Outgoing port    1/1/5
Outgoing VLAN    66
Outgoing VE      66
Ports monitored:
HW destination id for each device:
stack_id/device:dest_id
```

NOTE

The ARP can also be obtained by using a **ping 10.10.10.4** command.

Case 2: Problem Route not exist

```
device(config-vif-66)# show ip route

Total number of IP routes: 1
Type Codes - B:BGP D:Connected O:OSPF R:RIP S:Static; Cost - Dist/Metric
BGP Codes - i:iBGP e:eBGP
OSPF Codes - i:Inter Area 1:External Type 1 2:External Type 2
Destination      Gateway          Port          Cost      Type Uptime
1                10.10.10.0/24   DIRECT        ve 66     0/0   D    0m48s

device(config-vif-66)# disable
SYSLOG: <14> Jan  1 00:02:40 SWDR_8 System: Interface ve 66, state down

device(config-vif-66)# show erspan profile all

Profile 1
Type          ERSPAN
```

```
Mirror destination Not reachable.  
Reason: Route not exist  
Destination IP 10.10.10.4  
Destination MAC 0000.0000.0000  
Source IP 10.10.10.1  
Source MAC 748e.f8f9.6d80  
Ports monitored:  
HW destination id for each device:  
stack_id/device:dest_id
```

Solution - Enable the interface (VIF)

```
device(config-vif-66)# enable  
  
SYSLOG: <14> Jan 1 00:03:07 SWDR_8 System: Interface ve 66, state up  
  
device(config-vif-66)# show erspan profile all  
  
Profile 1  
Type ERSPAN  
Mirror destination Reachable.  
Destination IP 10.10.10.4  
Destination MAC 00aa.00bb.00cc  
Source IP 10.10.10.1  
Source MAC 748e.f8f9.6d80  
Outgoing port 1/1/5  
Outgoing VLAN 66  
Outgoing VE 66  
Ports monitored:  
HW destination id for each device:  
stack_id/device:dest_id
```

Case 3: Problem Outgoing port is management port

```
device(config-vif-66)# show erspan profile 1  
  
Profile 1  
Type ERSPAN  
Mirror destination Not reachable.  
Reason: Outgoing port is management port  
Destination IP 10.10.10.4  
Destination MAC 0000.0000.0000  
Source IP 10.10.10.1  
Source MAC 748e.f8f9.6d80  
Ports monitored:  
HW destination id for each device:  
stack_id/device:dest_id  
  
device(config-vif-66)# show ip route  
  
Total number of IP routes: 2  
Type Codes - B:BGP D:Connected O:OSPF R:RIP S:Static; Cost - Dist/Metric  
BGP Codes - i:iBGP e:eBGP  
OSPF Codes - i:Inter Area 1:External Type 1 2:External Type 2  
Destination Gateway Port Cost Type Uptime  
1 10.10.10.0/24 DIRECT e mgmt1 0/0 D 0m2s  
2 40.40.40.0/24 DIRECT e mgmt1 0/0 D 0m2s
```

The router for an ERSPAN destination IP can be learned by routing protocols (RIP, OSPF, etc.) or you can configure it statistically using **ip route** command. The commands **disable** or **enable** when run on the port is one way to add or remove routes.

Solution

There is no solution to this problem if you continue to use a management port as the outgoing port. You cannot use an IP address reachable through a management port as an ERSPAN destination.

Port Mirroring and Monitoring

Encapsulated Remote Switched Port Analyzer (ERSPAN)

Case 4: Problem Outgoing port is loopback port

```
device# show erspan profile 1

Profile 1
Type          ERSPAN
Mirror destination Not reachable.
Reason: Outgoing port is loopback port
Destination IP 10.10.10.4
Destination MAC 0000.0000.0000
Source IP      10.10.10.1
Source MAC     0000.0000.0000
Ports monitored:
HW destination id for each device:
stack_id/device:dest_id

device# show ip route

Total number of IP routes: 1
Type Codes - B:BGP D:Connected O:OSPF R:RIP S:Static; Cost - Dist/Metric
BGP Codes - i:iBGP e:eBGP
OSPF Codes - i:Inter Area 1:External Type 1 2:External Type 2
  Destination      Gateway         Port           Cost    Type Uptime
1      10.10.10.0/24   DIRECT         loopback 1    0/0   D     0m39s
```

Solution

There is no solution to this problem if you continue to use an IP address reachable through a loop back interface. You cannot use an IP address reachable through a loop back interface as an ERSPAN destination.

Case 5: Problem Outgoing port is GRE IP tunnel port

```
device# show erspan profile 1

Profile 1
Type          ERSPAN
Mirror destination Not reachable.
Reason: Outgoing port is GRE IP tunnel port
Destination IP 11.1.1.2
Destination MAC 0000.0000.0000
Source IP      33.33.33.2
Source MAC     0000.0000.0000
Ports monitored:
HW destination id for each device:
stack_id/device:dest_id

device(config)# show running-config | begin erspan

monitor-profile 1 type erspan
destination-ip 11.1.1.2
source-ip 33.33.33.2

device# show ip in
Interface      IP-Address      OK?  Method      Status      Protocol  VRF
Eth mgmt1     10.37.78.91     YES  NVRAM       up          up        default-vrf
Ve 33         33.33.33.2     YES  NVRAM       down        down      default-vrf
Ve 44         10.10.10.1     YES  manual      up          up        default-vrf
Tunnel 1      11.1.1.1       YES  manual      up          up        default-vrf
```

Solution

There is no solution to this problem if you continue to use an IP address reachable through a tunnel interface. You cannot use an IP address reachable through a tunnel interface as an ERSPAN destination.

Case 6: Problem Outgoing port is not known.

```
device# show erspan profile 1

Profile 1
Type          ERSPAN
Mirror destination Not reachable.
Reason: Outgoing port is not known
Destination IP 11.1.1.2
Destination MAC 0000.0000.0000
Source IP      33.33.33.2
Source MAC     0000.0000.0000
Ports monitored:
HW destination id for each device:
stack_id/device:dest_id
```

Solution - The same as with Case 2: Route not exist

Case 7: Problem Outgoing port is a sink port

```
device# show erspan profile 1

Profile 1
Type          ERSPAN

Mirror destination Not reachable.
Reason: Outgoing port is a sink port
Destination IP 11.1.1.2
Destination MAC 0000.0000.0000
Source IP      33.33.33.2
Source MAC     0000.0000.0000
Ports monitored:
HW destination id for each device:
stack_id/device:dest_id
```

Solution - The same as with Case 2: Route not exist.

Configuring a Monitor Port for ERSPAN

An ERSPAN monitor port is the port on which traffic is captured and sent to a remote destination over a Layer 3 network.

Before you configure the monitor port, you must configure an ERSPAN profile.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter interface configuration mode.

```
device(config)# interface ethernet 1/2/3
```

3. Configure the mirror port for ERSPAN.

```
device(config-if-e1000-1/2/3)# monitor profile 1 both
```

Port Mirroring and Monitoring

Encapsulated Remote Switched Port Analyzer (ERSPAN)

4. Verify the configuration.

```
device(config-if-e1000-1/2/3)# show erspan profile 1
Profile 1
Type                ERSPAN
Mirror destination  Reachable.
Destination IP      10.1.1.100
Destination MAC     cc4e.0000.0001
Source IP           10.1.1.1
Source MAC          cc4e.0000.0000
Outgoing port       1/2/3
Outgoing VLAN       101
Outgoing VE         101
Ports monitored:
  Input monitoring   : (U1/M1)  1
  Output monitoring  : (U1/M1)  1
HW destination id for each device:
stack_id/device:dest_id  1/1:3c000000
```

Port mirroring is now enabled between the monitor port and the destination that was specified in the ERSPAN profile.

ERSPAN Monitor Port Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/2/3
device(config-if-e1000-1/2/3)# monitor profile 1 both
device(config-if-e1000-1/2/3)# show erspan profile 1
```


sFlow

- [sFlow Overview](#)..... 81
- [Configuring sFlow](#)..... 85
- [sFlow Version 5 Feature Configuration](#).....88
- [Configuring sFlow with Multi-VRF](#)..... 90
- [Displaying sFlow Information](#)..... 91
- [Clearing sFlow Statistics](#)..... 91

sFlow Overview

sFlow is a standards-based protocol that allows network traffic to be sampled at a user-defined rate for the purpose of monitoring traffic flow patterns and identifying packet transfer rates on user-specified interfaces.

sFlow is described in RFC 3176, "InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks".

RUCKUS ICX devices support sFlow version 5 by default. When sFlow is enabled on a Layer 2 or Layer 3 switch, the system performs the following sFlow-related tasks:

- Samples traffic flows by copying packet header information
- Identifies ingress and egress interfaces for the sampled flows
- Combines sFlow samples into UDP packets and forwards them to the sFlow collectors for analysis
- Forwards byte and packet count data, or counter samples, to sFlow collectors

sFlow Version 5

sFlow version 5 enhances and modifies the format of the data sent to the sFlow collector. sFlow version 5 introduces several new sFlow features and also defines a new datagram syntax used by the sFlow agent to report flow samples and interface counters to the sFlow collector.

sFlow version 5 supports the following modifications:

- sFlow version 5 datagrams
- Sub-agent support
- Configurable sFlow export packet size
- Support for the new data field and sample type length in flow samples
- Configurable interval for exporting RUCKUS-specific data structures

sFlow version 5 is backward-compatible with sFlow version 2. The sFlow agent exports sFlow version 5 flow samples by default, but you can configure the device to export the data in sFlow version 2 format. You can switch between the sFlow version 2 and sFlow version 5 formats. The sFlow collector automatically parses each incoming sample and decodes it based on the version number.

The configuration procedures for sFlow version 5 are the same as for sFlow version 2, except where explicitly noted. For configuration procedures for sFlow, refer to [Configuring sFlow](#). The features and CLI commands that are specific to sFlow version 5 are described in [sFlow Version 5 Feature Configuration](#) on page 88.

sFlow Support for IPv6 Packets

The RUCKUS implementation of sFlow features supports IPv6 packets. This support includes extended router information and extended gateway information in the sampled packet. Note that sFlow support for IPv6 packets exists only on devices running software that supports IPv6.

The configuration procedures for IPv6 are the same as for IPv4, except where the collector is a link-local address on a Layer 3 switch. For details, refer to step 2 in Configuring sFlow.

Extended Router Information

IPv6 sFlow sampled packets include the following extended router information:

- IP address of the next-hop router
- Outgoing VLAN ID
- Source IP address prefix length
- Destination IP address prefix length

Note that in IPv6 devices, the prefix lengths of the source and destination IP addresses are collected if BGP is configured and the route lookup is completed. In IPv4 devices, this information is collected only if BGP is configured on the devices.

Extended Gateway Information

If BGP is enabled, extended gateway information is included in IPv6 sFlow sampled packets, including the following BGP information about a packet destination route:

- Autonomous System (AS) number for the router
- Source IP Autonomous System number for the route
- Source peer Autonomous System number for the route
- Autonomous System path to the destination

NOTE

Autonomous System communities and local preferences are not included in the sampled packets.

To obtain extended gateway information, use "struct extended_gateway" as described in RFC 3176.

IPv6 Packet Sampling

IPv6 sampling is performed by the packet processor. The system uses the sampling rate setting to selectively mark the monitoring bit in the header of an incoming packet. Marked packets inform the CPU that the packets are subject to sFlow sampling.

sFlow Configuration Considerations

Following are the sFlow configuration considerations on RUCKUS ICX devices:

- On ICX devices, you can use QoS queue 1 for priority traffic, even when sFlow is enabled on the port.
- If ICX stacks are rebooted, sFlow is disabled on standby and member units until the configuration is synchronized between the active and standby Controllers.
- sFlow is not supported on PE ports on 802.1BR-enabled RUCKUS ICX devices.

sFlow and Hardware Support

- RUCKUS ICX devices support sFlow packet sampling of inbound traffic only. These devices do not sample outbound packets. However, RUCKUS ICX devices support byte and packet count statistics for both traffic directions.
- sFlow is supported on all Ethernet ports (10/100 Mbps, 1 Gbps, and 10 Gbps)

sFlow and CPU Utilization

Enabling sFlow may cause a slight and noticeable increase of up to 20 percent in CPU utilization. In typical scenarios, this is normal behavior for sFlow, and does not affect the functionality of other features on the switch.

sFlow and Agent Address

The sampled sFlow data sent to the collectors includes an `agent_address` field. This field identifies the IP address of the device that sent the data:

- On a Layer 2 switch, `agent_address` is the Layer 2 switch management IP address. You must configure the management IP address in order to export sFlow data from the device. If the switch has both an IPv4 address and IPv6 address, the `agent_address` is the IPv4 address. If the switch has an IPv6 address only, the `agent_address` is the global IPv6 address.
- On a Layer 3 switch with IPv6 interfaces only, sFlow looks for an IPv6 address in the following order, and uses the first address found:
 - The first IPv6 address on the lowest-numbered loopback interface
 - The first IPv6 address on the lowest-numbered virtual interface
 - The first IPv6 address on any interface
- On a Layer 3 switch with both IPv4 and IPv6 interfaces, or with IPv4 interfaces only, sFlow looks for an IP address in the following order, and uses the first address found:
 - The IPv4 router ID configured by the `ip router-id` command
 - The first IPv4 address on the lowest-numbered loopback interface
 - The first IPv4 address on the lowest-numbered virtual interface
 - The first IPv4 address on any interface

NOTE

A device uses the router ID only if the device also has an IP interface with the same address. The router ID is not supported on IPv6 devices.

NOTE

If an IP address is not already configured when you enable sFlow, the sFlow version 5 feature uses the source address 0.0.0.0. To display the `agent_address`, enable sFlow, and then enter the `show sflow` command. Refer to step 7 in Configuring sFlow and [Displaying sFlow Information](#) on page 91.

NOTE

In sFlow version 5, you can set an arbitrary IPv4 or IPv6 address as the sFlow agent IP address. Refer to step 3 in Configuring sFlow Version 5.

sFlow and Source IP Address

When the sFlow packet is sent to the sFlow collector, by default, the IP address of the outgoing interface is used in the sFlow datagram.

However, you can specify the source interface, from which the IP address is selected for the sFlow datagram, using the `sflow source` command. The Ethernet, VE, or loopback interface can be configured as the source interface for both IPv4 and IPv6 addresses.

sFlow Source IP Address Configuration Notes

- The first IP address in the interface IP address list is considered the source IP address.
- If the sFlow destination is IPv6, and the sFlow source is configured for an IPv6 address, then an IPv6 address is selected from the configured interface.
- If the sFlow destination is IPv4, and the sFlow source is configured for an IPv4 address, then an IPv4 address is selected from the configured interface.
- At any point in time, only one source of the Ethernet, VE, or loopback interface can be specified as the source interface.
- Upon configuring another source for an IPv4 or IPv6 address, any previously configured source for the IPv4 or IPv6 address is deleted.
- If the source IP address is not configured, by default, the IP address of the outgoing interface is used in the sFlow datagram.
- You can configure IPv4 and IPv6 source interfaces independently.
- A LAG virtual interface or any member ports of the LAG cannot be configured as an sFlow source.
- The sFlow source IP address configuration is supported on sFlow version 2 and sFlow version 5 and is valid only for the router build.
- Addition and deletion of IPv4 and IPv6 addresses on an sFlow source interface triggers the following events:
 - If the added IP address is the first IP address in the table, then it is considered as the source IP address.
 - If the added IP address is positioned on top of the IP table (due to IP address sequence order), then it is reassigned as the source IP address.
 - If the IP address that is used as the source IP address is deleted, the next IP address on the same interface is considered as the source IP address.
 - If all the IP addresses are deleted from the source interface, the IP address of the outgoing interface is used in the sFlow datagram.

sFlow and Source Port

By default, sFlow sends data to the collector out of UDP source port 8888, but you can specify a different source port. For more information, refer to step 6 in Configuring sFlow.

sFlow and Sampling Rate

The sampling rate is the average ratio of the number of packets incoming on an sFlow-enabled port to the number of flow samples taken from those packets. sFlow sampling can affect performance in some configurations.

Note that on the RUCKUS ICX devices, the configured sampling rate and the actual rate are the same. The software does not adjust the configured sampling rate as on other RUCKUS ICX devices.

NOTE

The value range for the sampling rate is from 256 through 16777215 packets. The default value is 4096.

For more information on sampling rate configuration considerations, refer to step 5 in Configuring sFlow.

Sampling Rate Configuration Notes

- Configured rate and actual rate: When you enter a sampling rate value, this value is the configured rate as well as the actual sampling rate.
- Change to global rate: If you change the global sampling rate, the change is applied to all sFlow-enabled ports except those ports on which you have already explicitly set the sampling rate.
- Module rate: While different ports on a module may be configured to have different sampling rates, the hardware for the module is programmed to take samples at a single rate (the module sampling rate). The module sampling rate is the highest sampling rate (that is, the lowest number) configured for any of the ports on the module.

- Sampling rate for new ports: When you enable sFlow on a port, the port's sampling rate is set to the global default sampling rate. This also applies to ports on which you disable and then re-enable sFlow. The port does not retain the sampling rate it had when you disabled sFlow on the port, even if you explicitly set the sampling rate on the port.

sFlow and Port Monitoring

ICX devices support sFlow and port monitoring together on the same port.

Configuring sFlow

You must first enable sFlow globally and then sFlow forwarding can be enabled on an individual port or LAG port or both, and sFlow parameters can be configured.

The commands explained in this procedure apply to both sFlow version 2 and sFlow version 5. For commands specific to sFlow version 5, refer to Configuring sFlow version 5.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enable sFlow.

```
device(config)# sflow enable
```

3. Specify an sFlow collector with an IPv4 or IPv6 destination address.

```
device(config)# sflow destination 10.10.10.1
```

```
device(config)# sflow destination ipv6 2001:DB8:0::0b:02a
```

In the first example, a collector with IPv4 address 10.10.10.1 is listening for sFlow data on UDP port 6343. In the second example, a collector with IPv6 address 2001:DB8::0b:02a is listening for sFlow data on UDP port 6343.

NOTE

You can specify up to four collectors. You can specify more than one collector with the same IP address if the UDP port numbers are unique. You can have up to four unique combinations of IP addresses and UDP port numbers.

4. (Optional) Change the poll interval.

```
device(config)# sflow polling-interval 30
```

The poll interval value ranges from 0 through 4294967295 seconds. The default polling interval is 20 seconds. If you set the polling interval to 0, counter data sampling is disabled.

NOTE

The interval value applies to all interfaces on which sFlow is enabled. If multiple ports are enabled for sFlow, the RUCKUS ICX device staggers transmission of the counter data to smooth performance.

sFlow

Configuring sFlow

5. (Optional) Change the sampling rate using one of the following methods:

- Sampling rate can be changed globally (default).
- Sampling rate can be changed on an individual port.
- Sampling rate can be changed on a static or dynamic LAG.

```
device(config)# sflow sample 2048

device(config)# interface 1/1/1
device(config-if-1/1/1)# sflow sample 8192

device(config)# lag test static id 1
device(config-lag-test)# ports ethernet 1/1/1 to 1/1/4
device(config-lag-test)# sflow sample 8192
```

The sampling rate is a fraction in the form 1/N, meaning that, on average, one out of every N packets is sampled. The **sflow sample** command in global configuration mode or port mode specifies N, the denominator of the fraction. The sampling rate must be set depending on the input traffic rate.

On RUCKUSICX 7250, ICX 7450, ICX 7550, ICX 7650, and ICX 7850 devices, the CPU-bound sFlow sample packets are rate-limited to 50 samples per second to avoid high CPU utilization.

NOTE

You can change a module sampling rate only by changing the sampling rate of a port on that module.

NOTE

You can configure an individual port or a static LAG or a dynamic LAG to use a different sampling rate than the global default sampling rate. This is useful in cases where ports have different bandwidths.

6. Change the sFlow source port.

```
device(config)# sflow source-port 8000
```

By default, sFlow sends data to the collector using UDP source port 8888, but you can change the source UDP port to any port number ranging from 1025 through 65535.

7. Enable sFlow forwarding.

To enable sFlow forwarding, you must first enable sFlow on a global basis using the **sflow enable** command, then enable on individual interfaces or LAG ports or both using the **sflow forwarding** command.

a) Enable sFlow forwarding on individual interfaces.

```
device(config)# interface ethernet 1/1/1 to 1/1/8
device(config-mif-1/1/1-1/1/8)# sflow forwarding
```

b) Enable sFlow forwarding on individual LAG ports.

```
device(config)# lag test static id 1
device(config-lag-test)# ports ethernet 1/1/1 to 1/1/8
device(config-lag-test)# sflow forwarding
```

NOTE

sFlow forwarding is supported on LAG and LACP LAG ports.

NOTE

When a management port is used, sFlow can be received only from active units in a stack (not from all units). However, if you use a management VLAN with a data port, sFlow is received normally. To receive sFlow from all units in a stack, you must use a data port.

NOTE

When you enable sFlow forwarding on an 802.1X-enabled interface, the samples taken from the interface include the username used to obtain access to either or both the inbound and outbound ports, if that information is available. For information about 802.1X, refer to the "Flexible Authentication" chapter in the *RUCKUS FastIron Security Configuration Guide*.

8. Configure sFlow version 5 features if your device supports sFlow version 5.

Refer to [Configuring sFlow Version 5](#) on page 89.

sFlow

sFlow Version 5 Feature Configuration

The following example shows how to enable sFlow globally with the destination as an IPv4 address. sFlow parameters are configured.

```
device# configure terminal
device(config)# sflow enable
device(config)# sflow destination 10.10.10.1
device(config)# sflow polling-interval 30
device(config)# sflow sample 2048
device(config)# sflow source-port 8000
```

The following example shows how to enable sFlow globally with the destination as an IPv6 address on Ethernet port 1/1/1, and then enable sFlow forwarding on Ethernet ports 1/1/1 through 1/1/8. sFlow parameters are configured.

```
device# configure terminal
device(config)# sflow enable
device(config)# sflow destination ipv6 2001:DB8:0::0b:02a
device(config)# sflow polling-interval 30
device(config)# interface 1/1/1
device(config-if-1/1/1)# sflow sample 8192
device(config-if-1/1/1)# sflow source-port 8000
device(config-if-1/1/1)# interface ethernet 1/1/1 to 1/1/8
device(config-mif-1/1/1-1/1/8)# sflow forwarding
```

The following example shows how to enable sFlow globally with the destination as an IPv4 address, and then enable sFlow forwarding on Ethernet ports 1/1/1 through 1/1/8, within the LAG. sFlow parameters are configured.

```
device# configure terminal
device(config)# sflow enable
device(config)# sflow destination 10.10.10.1
device(config)# sflow polling-interval 30
device(config)# sflow source-port 8000
device(config)# lag test static id 1
device(config-lag-test)# ports ethernet 1/1/1 to 1/1/8
device(config-lag-test)# sflow sample 8192
device(config-lag-test)# ports ethernet 1/1/1 to 1/1/8
device(config-lag-test)# sflow forwarding
```

sFlow Version 5 Feature Configuration

NOTE

The commands shown in Configuring sFlow Version 5 are supported when sFlow version 5 is enabled on the device. The commands are not supported with sFlow version 2. sFlow version 5 also supports all of the sFlow configuration commands in Configuring sFlow.

When sFlow version 5 is enabled on the device, perform the following actions:

- Specify the sFlow version (version 2 or version 5)
- Specify the sFlow agent IP address
- Specify the maximum flow sample size
- Export CPU and memory usage information to the sFlow collector
- Specify the polling interval for exporting CPU and memory usage information to the sFlow collector
- Export CPU-directed data (management traffic) to the sFlow collector

Specifying Maximum Packet Size Values

With sFlow version 5, you can specify the maximum size of the flow sample sent to the sFlow collector.

The following table provides information about the sFlow sample size sent to the sFlow collector when the **sflow max-packet-size** is configured with different values.

TABLE 9 Maximum Packet Size Values

Maximum Packet Size (in bytes)	sFlow Sample Size
0	Only the information about the packet is captured and no data from the packet is sent to the sFlow collector.
1	One byte from the packet is sent to the sFlow collector. However, the packet is padded with zeros to make it 4 bytes.
2	Two bytes from the packet is sent to the sFlow collector. However, the packet is padded with zeros to make it 4 bytes.
100	100 bytes from the packet are sent to the sFlow collector.
200	200 bytes from the packet are sent to the sFlow collector.
1200	1200 bytes from the packet are sent to the sFlow collector.

Configuring sFlow Version 5

You can configure sFlow version 5 features if your device supports sFlow version 5.

When broadcast and multicast packets are sampled, they are usually forwarded to more than one port. However, the output port field in an sFlow datagram supports the display of one egress interface ID only. Therefore, the sFlow version 5 agent always sets the output port ID to 0x80000000 for broadcast and multicast packets that are sampled.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Specify the sFlow version.

```
device(config)# sflow version 5
```

3. Specify the sFlow agent IP address (IPv4 or IPv6).

```
device(config)# sflow agent-ip 10.10.10.1
device(config)# sflow agent-ip FE80::240:D0FF:FE48:4672
```

4. Specify the maximum flow sample size.

```
device(config)# sflow max-packet-size 1024
```

If a packet is larger than the specified maximum size, only the data of the packet up to the specified maximum number of bytes is exported.

5. Export CPU and memory usage information to the sFlow collector using the **sflow export system-info** and specify the poll interval.

```
device(config)# sflow export system-info 30
```

Polling interval value ranges from 5 through 1800 seconds. The default value is 300 seconds.

6. Export CPU-directed data (management traffic) to the sFlow collector using the **sflow export cpu-traffic** and specify the sampling rate.

```
device(config)# sflow export cpu-traffic 2048
```

The default sampling rate depends on the RUCKUS ICX device being configured. Refer to step 5 in Configuring sFlow for the default sampling rate for each kind of RUCKUS ICX device.

sFlow

Configuring sFlow with Multi-VRF

The following example shows how to configure sFlow version 5, specifying an IPv4 address as the sFlow agent IP address. The sFlow parameters are configured.

```
device# configure terminal
device(config)# sflow version 5
device(config)# sflow agent-ip 10.10.10.1
device(config)# sflow max-packet-size 1024
device(config)# sflow export system-info 30
device(config)# sflow export cpu-traffic 2048
```

Configuring sFlow with Multi-VRF

sFlow is a traffic-monitoring protocol that supports VRFs. sFlow provides traffic sampling on configured ports, based on sample rate and port information, to a collector. By default, sFlow uses the management VRF to send the samples to the collector.

Collectors can be added to individual VRFs so that collectors can be spread out across different VRFs. The sFlow forwarding port can belong to a non-default VRF, and captured sFlow packets will contain the correct sample routing next-hop information.

sFlow forwarding ports can come from ports belonging to any VRF. The port is not required to be in the same VRF as the collector. sFlow collects packets from all sFlow forwarding ports (even if they do not belong to a VRF), compiles the packets into the sFlow samples, and sends the samples to the particular collector with no filtering for VRF membership. For counter samples, sample statistics from each port are sent to each specified collector, even if the port and collector do not belong to a VRF instance.

To distinguish collected packets from different VRFs, refer to the VLAN data fields for each captured ingress packet. For example, when two collected packets are from different VRFs but have the same source or destination IP address and the same incoming or outgoing port, the VLAN data fields differ in the two samples. A VLAN or VE can belong to only one VRF. The collector does not have any VRF knowledge, but, based on the VLAN fields, the collector can distinguish which packet came from which VLAN or VRF.

To configure an sFlow collector and specify a VRF, enter the following command.

```
device(config)# sflow destination 10.10.10.vrf customer1
```

To disable the management VRF in sFlow, enter the following command.

```
device(config)# sflow management-vrf disable
```

To display sFlow configuration and statistics, enter the following command.

```
device(config)# show sflow
sFlow version: 5
sFlow services are enabled.
sFlow management VRF is disabled.
sFlow agent IP address: 10.37.230.21
Collector IP 10.37.224.233, UDP 6343, Configured VRF: green
UDP source port: 8888 (Default)
Polling interval is 20 seconds.
Configured default sampling rate: 1 per 500 packets.
Actual default sampling rate: 1 per 500 packets.
The maximum sFlow sample size: 128.
sFlow exporting cpu-traffic is disabled.
100 UDP packets exported
80 sFlow flow samples collected.
sFlow ports: ethe 4/1/5
Module Sampling Rates
-----
Port Sampling Rates
-----
Port=4/1/5, configured rate=500, actual rate=500
```

Displaying sFlow Information

To display sFlow configuration information and statistics, enter the following command from any mode of the CLI.

```
device# show sflow
sFlow version:5
sFlow services are enabled.
sFlow agent IP address: 10.123.123.1
sFlow source IP address: 5.5.5.5
sFlow source IPv6 address: 4545::2
4 collector destinations configured:
Collector IP 192.168.4.204, UDP 6343
Collector IP 192.168.4.200, UDP 6333
Collector IP 192.168.4.202, UDP 6355
Collector IP 192.168.4.203, UDP 6565
Configured UDP source port: 33333
Polling interval is 0 seconds.
Configured default sampling rate: 1 per 512 packets
Actual default sampling rate: 1 per 512 packets
Sample mode: Non-dropped packets
The maximum sFlow sample size:512
exporting cpu-traffic is enabled
exporting cpu-traffic sample rate:16
exporting system-info is enabled
exporting system-info polling interval:20 seconds
10552 UDP packets exported
24127 sFlow samples collected.
sFlow ports: ethe 1/1/2 to 1/1/12 ethe 1/1/15 ethe 1/1/25 to 1/1/26 ethe 1/4/1 ethe 1/5/10 to
1/5/20 ethe 1/8/1 ethe 1/8/4
Module Sampling Rates
-----
Slot 1 configured rate=512, actual rate=512
Slot 3 configured rate=0, actual rate=0
Slot 4 configured rate=10000, actual rate=32768
Slot 5 configured rate=512, actual rate=512
Slot 7 configured rate=0, actual rate=0
Slot 8 configured rate=512, actual rate=512
Port Sampling Rates
-----
Port 1/8/4, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/8/1, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/5/20, configured rate=3000, actual rate=8192, Subsampling factor=16
Port 1/5/19, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/5/18, configured rate=512, actual rate=512, Subsampling factor=1
Port 1/5/17, configured rate=1500, actual rate=2048, Subsampling factor=4
...Output truncated...
```

Clearing sFlow Statistics

To clear the UDP packet and sFlow sample counters (as shown in the **show sflow** output), enter the following command.

```
device# clear statistics
```

NOTE

The **clear statistics** command also clears the statistics counters used by other features (port statistics, egress queue statistics, and management interface).

HMON - Health Monitor Service

- [HMON overview.....](#) 93
- [Troubleshooting HMON.....](#) 95

HMON overview

Hosted applications and processes that are registered with Health Monitor in the FastIron operating system are monitored as high-availability (HA) processes. Health Monitor (HMON) provides the following services on ICX devices for registered applications and processes:

- Monitoring at specified intervals and restarting the process on failure
- On-demand process start or stop
- Process stop or start based on the device role in a stack

HMON process registration

For each process registered with HMON (also referred to as client processes), a set of attributes can be configured. These attributes include the following parameters:

- Monitoring interval - specified in 10 second increments
- Process criticality - critical or non-critical
- Maximum number of restarts - configurable cap on restart and recovery
- Stack role mask - the stack roles under which the process can run

NOTE

When a process fails, related functionality is not available from the time the process fails until HMON recovers it. Recovery time depends on the monitoring interval registered for the client process.

Dynamic start and stop

When a process is registered with HMON as an "on demand" process, FastIron sends IPC messages to HMON to start and stop the process dynamically based on defined functionality. HMON monitors the process from the time it is started until it is stopped.

Process availability based on stack role

Some processes are dependent on the current role of the ICX device in a stack configuration. For example, a particular application may run on the device only while the device is acting as the active controller. Role-dependency can be configured as part of HMON client registration, and the client process can be started and stopped in relation to the role.

Clients marked as faulty

When a client process exceeds the maximum number of recovery attempts, it is marked as faulty and is no longer available. The failure is logged as a syslog message. Cyclic crashes are an indication of an issue that should be addressed. Contact Ruckus technical support for assistance.

Critical processes

If the faulty client has been registered with HMON as system critical, the ICX device reboots. In a stack configuration, the standby controller takes over as the active controller, which may restore the client process to normal operation.

To determine whether a registered HMON client is a critical or non-critical process, enter the **show hmon client configuration all-clients** command in Privileged EXEC mode. Refer to [Troubleshooting HMON](#) on page 95 for an example of **show hmon client configuration all-clients** command output.

Determining the administrative and operational state of an HMON client.

Enter the **show hmon client status all-clients** command in Privileged EXEC mode to display both the administrative and operational state of all HMON clients registered on the device. Administrative states are defined as follows:

- Enabled, Not Started, HA Disabled - The client process is enabled; however, it is not started, as it is not qualified to run on the current stack role. As a result, the process is not being monitored.
- Enabled, Started, HA Enabled - The application is enabled, started or running, and monitored for HA.
- Disabled, HA Disabled - The application is not enabled, not started or running, and not monitored for HA.

The operational state provides additional information for each client process and can be one of the following:

- Up - The client process is up and running.
- Down - The client process is not running.
- Recovering - HMON has initiated a recovery for the crashed or failed process, and the process is recovering (transient state).
- Recovery Failed - Recovery has failed.
- Faulty - Due to repeated failed recoveries, the maximum allowable recovery attempts has been exceeded, and the client process is marked as faulty.

Dependent process restart

Beginning with FastIron release 09.0.00, when a process monitored by HMON restarts for recovery, the dependent processes also gets restarted automatically.

Dependent processes are defined as those processes that require a restart when another process is going for an HA restart recovery. All the dependent processes are grouped together and each process in the group have a criticality defined as following:

- HMON_DEP_PRCSGRP_CRITICAL: Other processes in the group should be restarted if this process restarts.
- HMON_DEP_PRCSGRP_NON_CRITICAL: Other processes in the group should not be restarted if this process restarts.

A process group can have more than one HMON_DEP_PRCSGRP_CRITICAL member in it. Processes should be grouped for dependency only if there is at least one HMON_DEP_PRCSGRP_CRITICAL member in it.

NOTE

When the HMON_DEP_PRCSGRP_ROLE_CRITICAL process crashes or become non-functional it will be restarted first for recovery, followed by dependent processes. The order of restart of dependent processes are not defined.

The following syslog message is displayed when dependant process is restarted:

```
Application <process name> is being restarted as the process <HMON_GROUP_CRITICAL process name> which it depends on got restarted
```

When process restart fails and HMON exceeds the maximum number of restart recovery attempts, the following critical syslog message is displayed to notify the unavailability of functionality provided by this process.

```
Application <process name> failed recovery and functionality provided by it will not be available until  
failure reason is remedied (may require manual intervention)
```

Contact RUCKUS technical support if this type of failure occurs.

Refer to [Troubleshooting HMON](#) on page 95 for an example of `show hmon client status all-clients` command output.

Troubleshooting HMON

NOTE

HMON commands provide output specific to the device on which the command is executed.

NOTE

Refer to the *Ruckus FastIron Command Reference Guide* for more information on the commands.

When an HMON client is marked as faulty, a syslog message similar to the following is issued:

```
Application webserver failed recovery and functionality provided by it may not be available until FastIron  
acts on it
```

Work with RUCKUS technical support if this type of failure occurs. Before contacting technical support, perform the following steps to gather diagnostic information.

1. In Privileged EXEC mode, enter the `show hmon status` command and check the client list to confirm that the failed process is indeed being monitored by HMON. Capture the output.

```
device# show hmon status  
-----  
Health Monitor Status:  
-----  
Hmon's Stack Role is : Standalone  
Number of Clients    : 18  
  
Client Names (ID) :  
  nginx (4)  
  uwsgi (5)  
  PySzAgtSrv.py (6)  
  dhcpd (3)  
  dhcpdv4 (20)  
  snmpagentd (7)  
  logmgr.py (13)  
  poed (8)  
  netconfd_svc (9)  
  fcgiwrap (11)  
  netconfd-pro (10)  
  plugin_proc.py (15)  
  cfg_sync (12)  
  sshd (16)  
  fi_syslog (17)  
  slam_server.py (18)  
  PyWebFlaskAgent (19)  
  socat (28)
```

The previous example of the `hmon status` command displays information for eighteen registered HMON processes identified by their Client name and ID. The output also indicates that the device is not part of a stack.

2. Enter the **show hmon client status all-clients** command to check the administrative and operational state of the application. Capture the output.

```
device# show hmon client status all-clients
-----
Health Monitor Client Status:
-----

Status for client ID 4:
Process Name           : nginx
Valid                  : Yes
Admin. State           : Enabled, Started, HA Enabled
Oper. State            : Up

Status for client ID 5:
Process Name           : uwsgi
Valid                  : Yes
Admin. State           : Disabled, HA Disabled
Oper. State            : Down

Status for client ID 6:
Process Name           : PySzAgtSrv.py
Valid                  : Yes
Admin. State           : Disabled, HA Disabled
Oper. State            : Down

Status for client ID 3:
Process Name           : dhcpd
Valid                  : Yes
Admin. State           : Disabled, HA Disabled
Oper. State            : Down
```

The previous example of the **show hmon client status all-clients** shows that the dhcpd process is disabled.

3. Enter the **show log** command and check for hmond entries similar to those in the following example.

```
Dynamic Log Buffer (4000 lines):
Mar  4 22:22:19:E:hmond[392]: Client uwsgi-2.7 has reached/exceeded max funcmtr fail count: 2,
initiating recovery from state UtilReportedFail
Mar  4 22:22:19:E:hmond[392]: Client uwsgi-2.7 is not functional, fail count is: 2, funcmtr fail
count limit is: 2
Mar  4 22:22:09:E:hmond[392]: Client uwsgi-2.7 is not functional, fail count is: 1, funcmtr fail
count limit is: 2
```


4. Enter the **show hmon client statistics all-clients** command and capture the output.

```

device# show hmon client statistics all-clients
-----
Health Monitor Client Statistics:
-----

Statistics for client ID 4:
Process Name                : nginx
Most recent PID             : 1328
Func. Monitor fail counts   : 0
Total number of admin stops : 0
Total number of disallowed admin stops : 0
Total number of admin starts : 1
Total number of disallowed admin starts : 1
Total number of admin restarts : 0
Total number of restarts for recovery : 0
Code from latest func. fail recovery : 0x0
Status from latest Func. Monitor check : Invalid

Statistics for client ID 5:
Process Name                : uwsgi-2.7
Most recent PID             : 1343
Func. Monitor fail counts   : 0
Total number of admin stops : 0
Total number of disallowed admin stops : 0
Total number of admin starts : 1
Total number of disallowed admin starts : 1
Total number of admin restarts : 0
Total number of restarts for recovery : 0
Code from latest func. fail recovery : 0x0
Status from latest Func. Monitor check : Invalid

Statistics for client ID 6:
Process Name                : PySzAgtSrv.py
Most recent PID             : 1356
Func. Monitor fail counts   : 0
Total number of admin stops : 0
Total number of disallowed admin stops : 0
Total number of admin starts : 1
Total number of disallowed admin starts : 1
Total number of admin restarts : 0
Total number of restarts for recovery : 0
Code from latest func. fail recovery : 0x0
Status from latest Func. Monitor check : Access Issue

Statistics for client ID 3:
Process Name                : dhcpd
Most recent PID             : Not Available
Func. Monitor fail counts   : 0
Total number of admin stops : 0
Total number of disallowed admin stops : 0
Total number of admin starts : 0
Total number of disallowed admin starts : 0
Total number of admin restarts : 0
Total number of restarts for recovery : 0
Code from latest func. fail recovery : 0x0
Status from latest Func. Monitor check : Not Invoked

```

5. Enter the **show hmon client configuration all-clients** command and capture the output.

```
device# show hmon client configuration all-clients
-----
Health Monitor Client Configuration:
-----

Configuration attributes for client ID 4:
Process Name           : nginx
Startup Script         : nginx-service.sh
Stackrole mask         : 0x3
Starts on bootup       : No
Process restartable    : Yes
Criticality of the process : Non-Critical
Process restart count limit : 5
Heart-Beat monitoring reqd. : No
Functionality monitoring reqd.: Yes
Func. monitoring interval : 10 Secs
Func. fail count limit  : 2

Configuration attributes for client ID 5:
Process Name           : uwsgi-2.7
Startup Script         : uwsgi-service.sh
Stackrole mask         : 0x3
Starts on bootup       : No
Process restartable    : Yes
Criticality of the process : Non-Critical
Process restart count limit : 5
Heart-Beat monitoring reqd. : No
Functionality monitoring reqd.: Yes
Func. monitoring interval : 10 Secs
Func. fail count limit  : 2

Configuration attributes for client ID 6:
Process Name           : PySzAgtSrv.py
Startup Script         : pySzagent-service.sh
Stackrole mask         : 0x3
Starts on bootup       : No
Process restartable    : Yes
Criticality of the process : Non-Critical
Process restart count limit : 5
Heart-Beat monitoring reqd. : No
Functionality monitoring reqd.: Yes
Func. monitoring interval : 10 Secs
Func. fail count limit  : 2

Configuration attributes for client ID 3:
Process Name           : dhcpd
Startup Script         : dhcpd-script.sh
Stackrole mask         : 0x3
Starts on bootup       : No
Process restartable    : Yes
Criticality of the process : Non-Critical
Process restart count limit : 5
Heart-Beat monitoring reqd. : No
Functionality monitoring reqd.: Yes
Func. monitoring interval : 10 Secs
Func. fail count limit  : 2
```

The previous example of the **show hmon client configuration all-clients** command shows that all HMON clients running on the device are non-critical; that is, the processes become unavailable when marked faulty, and no unit reboot is attempted to initiate switchover to the standby controller.

6. Enter the **supportsave all** followed by the IP address of the tftp server where supportsave logs are to be uploaded as shown in the following example.

NOTE

For more **supportsave** command options, refer to the *RUCKUS FastIron Command Reference*.

```
ICX7650-48P Router# supportsave all 10.22.141.59
```

7. Collect the logs.
8. Contact RUCKUS technical support.

Syslog

- Syslog Messages..... 101
- Enabling Real-Time Display of Syslog Messages..... 102
- Disabling Real-Time Display of Syslog Messages..... 102
- Broadcast, Unknown Unicast, and Multicast Suppression Syslog and SNMP notification..... 103
- Displaying Syslog Messages..... 104
- Syslog Service Configuration..... 105

Syslog Messages

FastIron software can write syslog messages to provide information at the following severity levels:

- Emergencies
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

The device writes the messages to a local buffer.

You also can specify the IP address or host name of up to six syslog servers. When you specify a syslog server, the RUCKUS ICX device writes the messages to both the system log and the syslog server.

Using a syslog server ensures that the messages remain available even after a system reload. The RUCKUS local syslog buffer is cleared during a system reload or reboot, but the syslog messages sent to the syslog server remain on the server.

NOTE

To enable the RUCKUS device to retain syslog messages after a soft reboot (**reload** command), refer to Syslog Reboot Configuration Considerations.

The syslog service on a syslog server receives logging messages from applications on the local host or from devices such as a Layer 2 switch or Layer 3 switch. Syslog adds a timestamp to each received message and directs messages to a log file. Most Unix workstations come with syslog configured. Some third-party vendor products also provide syslog server running on Windows.

Syslog uses UDP port 514 and therefore each syslog message is sent with destination port 514. Each syslog message is one line within the syslog message format. The message is embedded in the text portion of the syslog format. There are several subfields in the format. Keywords are used to identify each subfield, and commas are delimiters. The subfields can appear in any order, except that the text subfield should be the last field in the message. All the subfields are optional.

Enabling Real-Time Display of Syslog Messages

You can enable real-time display of syslog messages from global configuration mode or from a Telnet or SSH session.

1. Enter global configuration mode.

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

2. Enable real-time display of syslog messages.

- Enable real-time display of syslog messages in the management console.

```
device(config)# logging console
```

- Enable real-time display of syslog messages in a Telnet or an SSH session from privileged EXEC mode.

```
telnet@device# terminal monitor  
Syslog trace was turned ON
```

You must be in a Telnet or SSH session before using the **terminal monitor** command.

When the logging console is not enabled on the device and you enter the **terminal monitor** command in a Telnet or SSH session, the console displays the following error message. You must enable the logging console first and then configure the terminal monitor.

```
telnet@device# terminal monitor  
Logging console disabled  
telnet@device# configure terminal  
telnet@device(config)# logging console  
telnet@device(config)# end  
telnet@device# terminal monitor  
Syslog trace was turned ON
```

The following example shows how to enable real-time display of syslog messages in the management console.

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

The following example shows how to enable real-time display of syslog messages in a Telnet or SSH session from privileged EXEC mode.

```
device(config)# configure terminal  
telnet@device# terminal monitor
```

Disabling Real-Time Display of Syslog Messages

To disable real-time display of syslog messages in the management console, enter the **no logging on** command from global configuration mode.

```
device(config)# no logging on
```

In a Telnet and SSH session, the **terminal monitor** command acts as a toggle. The first entry enables and the second entry disables the display of syslog messages.

```
telnet@device# terminal monitor  
Syslog trace was turned OFF
```

Broadcast, Unknown Unicast, and Multicast Suppression Syslog and SNMP notification

Rate limiting broadcast, unknown unicast, and multicast (BUM) traffic protects a switch, router node, or network from Denial of Service (DoS) attacks or unintentional traffic configurations. When an incoming packet exceeds the maximum number of bytes that you set with rate limiting, a syslog notification is generated.

BUM Suppression Restrictions and Limitations

- All of the restrictions that apply to configuring ACLs on an interface apply to BUM suppression. For the restrictions that apply to ACLs, refer to the *RUCKUS FastIron Security Configuration Guide*. The main restriction is that you cannot change VLAN membership of a port.
- By default, the syslog logs once a minute; however, you can configure syslog notifications so that they log at a maximum interval of every 10 minutes.

Enabling BUM Suppression Logging

Complete the following steps to enable BUM suppression logging.

NOTE

Rate limiting must be enabled.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter Ethernet interface configuration mode.

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

3. Enable rate limiting.

```
device(config-if-e10000-1/1/1)# broadcast limit 8388607
```

Broadcast is used in this example. The command syntax is same for multicast or unknown unicast with the corresponding command (**multicast** or **unknown-unicast**).

4. Enable logging when the specified limit is exceeded.

```
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps log
```

Broadcast is used in this example. The command syntax is same for multicast or unknown unicast with the corresponding command (**multicast** or **unknown-unicast**).

5. Use the **threshold** option for the broadcast limit to shut-down the port and generate the syslog.

```
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps threshold 1000 action port-shutdown
```

The following example shuts down the port for 300 seconds (default) when the packet drop threshold value exceeds 1000 Kbps.

```
device(config)# interface ethernet 1/2/1  
device(config-if-e10000-1/2/1)# unknown-unicast limit 100 kbps threshold 1000 action port-shutdown
```

Syslog

Displaying Syslog Messages

6. Globally configure the log interval.

```
device(config)# rate-limit-log 6
device(config)# exit
```

7. Verify the logging interval.

```
device(config)# show running-config | include rate-limit-log
rate-limit-log 6
```

8. Verify the configuration.

```
device# show logging | include 1/1/1
Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 1434 kB
are dropped
```

Enabling BUM Suppression Logging Configuration Example

```
device# configure terminal
device(config)# interface ethernet 1/1/1
device(config-if-e10000-1/1/1)# broadcast limit 8388607
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps log
device(config-if-e10000-1/1/1)# broadcast limit 100 kbps threshold 1000 action port-shutdown
device(config)# rate-limit-log 6
device(config)# show running-config | include rate-limit-log
device(config)# exit
device# show logging | include 1/1/1
```

Viewing BUM Suppression Syslog Notifications

Use the following commands to display BUM suppression syslog notification information.

Use the **show logging** command to view the BUM suppression syslog notifications for all interfaces.

```
device# show logging
Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 11620 kB are
dropped
Jan 13 12:14:23:I:Security: Interface ethernet 1/3/12 reached the Multicast traffic limit and 870 kB are
dropped
Jan 13 12:45:38:I:Security: Interface ethernet 3/2/14 reached the Unknown-Unicast traffic limit and 2321 kB
are dropped
```

The first section of the output is `mmm dd hh:mm:ss:Info:System`.

To view the BUM suppression syslog notifications for a specific interface, use the following command.

```
device# show logging | include 1/1/1
Jan 13 12:02:12:I:Security: Interface ethernet 1/1/1 reached the Broadcast traffic limit and 11620 kB are
dropped
```

Displaying Syslog Messages

To display the syslog messages in the device local buffer, enter the **show logging** command from any CLI mode.

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 9 overruns)
Buffer logging: level ACDMEINW, 50 messages logged
level code: A=alert C=critical D=debugging M=emergency E=error
             I=informational N=notification W=warning

Static Log Buffer:
```



```
Jan  1 00:00:56:I:System: Stack unit 1   Power supply 2   is up

Dynamic Log Buffer (50 lines):
Feb  5 01:22:17:D:DHCPC: TFTP unable to download running-configuration
Feb  5 01:22:16:D:DHCPC: sending TFTP request for bootfile name ruckus.cfg
Feb  5 01:22:15:D:DHCPC: sending TFTP request for bootfile name icx7450.cfg
Feb  5 01:22:14:D:DHCPC: sending TFTP request for bootfile name ICX7450-24-
Route
  r.cfg
Feb  5 01:22:13:D:DHCPC: sending TFTP request for bootfile name ICX7450-24-
Route
  rcc4e.248b.b068.cfg
Feb  5 01:22:12:I:DHCPC: ICX7450-24 Router configured with ip-address 10.10.10.10; subnet mask
255.255.255.0 on port mgmt1
Feb  5 01:22:12:D:DHCPC: sending TFTP request for bootfile name ICX7450-24-
Route
  rcc4e.248b.b068-config.cfg
Feb  5 01:22:12:I:DHCPC: Setting boot-image download to secondary
Feb  5 01:22:12:E:DHCPC: Failed to configure default gatewa
Feb  5 01:22:12:I:PORT: mgmt1, added ip address 10.10.10.10 by un-
authenticate
  d user from console session
Feb  5 01:22:12:I:PORT: mgmt1, removed ip address 10.10.10.10 by un-
authentica
  ted user from console session
Feb  4 01:19:59:D:DHCPC: TFTP unable to download running-configuration
Feb  4 01:19:58:D:DHCPC: sending TFTP request for bootfile name ruckus.cfg
Feb  4 01:19:57:D:DHCPC: sending TFTP request for bootfile name icx7450.cfg
...
```

For information about syslog configuration, timestamps, and dynamic and static buffers, refer to [Displaying the Syslog Configuration](#) on page 109.

Displaying Real-Time Syslog messages

Any terminal logged in to a RUCKUS ICX switch can receive real-time syslog messages when the **terminal monitor** command is used.

Syslog Service Configuration

Complete the following procedures to configure the syslog:

- Specify a syslog server. You can configure the RUCKUS ICX device to use up to six syslog servers. (Use of a syslog server is optional. The system can hold up to 1000 syslog messages in an internal buffer.)
- Change the level of messages the system logs.
- Change the number of messages the local syslog buffer can hold.
- Configure a threshold value for logging buffer based on which a syslog warning message is generated when the log entries exceed the specified threshold value.
- Display the syslog configuration.
- Clear the local syslog buffer.

Logging is enabled by default, with the following settings:

- Messages of all severity levels (Emergencies through Debugging) are logged.
- By default, up to 4000 messages are retained in the local syslog buffer (but the number can be changed).
- No syslog server is specified.

Static and Dynamic Buffers

FastIron software provides two buffers:

- Static: Logs power supply failures, fan failures, and temperature warning or shutdown messages
- Dynamic: Logs all other message types

In the static log, new messages replace older ones, so only the most recent message is displayed. For example, only the most recent temperature warning message will be present in the log. If multiple temperature warning messages are sent to the log, the latest one replaces the previous one. The static buffer is not configurable.

The message types that appear in the static buffer do not appear in the dynamic buffer. The dynamic buffer contains up to the maximum number of messages configured for the buffer (4000 by default), and then begins removing the oldest messages (at the bottom of the log) to make room for new ones.

The static and dynamic buffers are both displayed when you display the log.

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
    level code: A=alert C=critical D=debugging M=emergency E=error
                I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Notice that the static buffer contains two separate messages for fan failures. Each message of each type has its own buffer. Thus, if you replace fan 1, but that fan also fails, the software replaces the first message about the failure of fan 1 with the newer message. The software does not overwrite the message for fan 2, unless the software sends a newer message for fan 2.

Clearing Log Entries

When you clear log entries, you can selectively clear the static or dynamic buffer, or you can clear both. For example, to clear only the dynamic buffer, enter the following command from privileged EXEC mode.

```
device# clear logging dynamic-buffer
```

You can specify **dynamic-buffer** to clear the dynamic buffer or **static-buffer** to clear the static buffer. If you do not specify a buffer, both buffers are cleared.

Timestamps

The contents of the timestamp differ depending on whether you have set the time and date on the onboard system clock.

If you have set the time and date on the onboard system clock, the date and time are shown in the following format:

```
mmm dd hh:mm:ss
```

The format takes the following form:

- *mmm*: Three-letter abbreviation for the name of the month
- *dd*: Day of the month
- *hh*: Hours

- *mm*: Minutes
- *ss*: Seconds

For example, "Oct 15 17:38:03" means October 15 at 5:38 PM and 3 seconds.

Example of Syslog messages on a device with the onboard clock set

The following example shows the format of messages on a device where the onboard system clock is set. Each timestamp shows the month, the day, and the time of the system clock when the message was generated. For example, the system time when the most recent message (the one at the top) was generated was October 15 at 5:38 PM and 3 seconds.

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 38 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dec 15 19:00:14:A:Fan 2, fan on left connector, failed
Dynamic Log Buffer (50 entries):
Oct 15 17:38:03:warning:list 101 denied tcp 10.157.22.191(0) (Ethernet 18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
Oct 15 07:03:30:warning:list 101 denied tcp 10.157.22.26(0) (Ethernet 18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
Oct 15 06:58:30:warning:list 101 denied tcp 10.157.22.198(0) (Ethernet 18
0000.001f.77ed) -> 10.99.4.69(http), 1 event(s)
```

Configuring Syslog Service

By default, syslog logging is enabled and up to 4000 messages are retained in the local syslog buffer.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Specify a syslog server.

```
device(config)# logging host 10.0.0.99
```

3. Change the message level the system logs by disabling logging of specific message levels.

```
device(config)# no logging buffered debugging
device(config)# no logging buffered informational
```

You must disable the message levels on an individual basis. In the example, the system does not log debugging and informational messages.

4. Change the number of messages the local syslog buffer can hold by using the **logging buffered** command.

```
device(config)# logging buffered 1000
device(config)# write memory
device(config)# exit
device# reload
```

The syslog buffer limit ranges from 1 through 4000 entries. In the example, the syslog buffer can store up to 1000 entries.

You must save the configuration and reload the software to place any changes into effect.

Syslog

Syslog Service Configuration

5. Set a threshold value for logging buffer based on which a syslog warning message is generated when the log entries exceed the specified threshold value. The threshold value is in percentage.

```
device(config)# logging buffered threshold 50
```

6. Change the log facility by using the **logging facility** command.

```
device(config)# logging facility local0
```

In the example, the log facility is changed to local0. "User" is the default log facility for messages sent by the RUCKUS ICX device to the syslog server. You can specify only one facility. If you configure the RUCKUS ICX device to use two syslog servers, the device uses the same facility on both servers.

7. Display the interface names in syslog messages.

```
device(config)# ip show-portname
```

The **ip show-portname** command is applied globally to all interfaces on Layer 2 and Layer 3 switches. By default, syslog messages show the interface type. If a port name is configured and assigned a name, then the port name is displayed.

8. Display the syslog configuration from privileged EXEC mode.

```
device# show logging
```

9. Configure the device to save the syslog messages after a soft reboot.

```
device(config)# logging persistence
```

Using the **logging persistence** command does not save the syslog messages after a hard reboot.

10. Clear the local syslog buffer from privileged EXEC mode.

```
device# clear logging
```

The following example shows how to configure syslog messages where logging is disabled for debugging messages. Logging parameters are configured.

```
device(config)# configure terminal
device(config)# logging host 10.0.0.99
device(config)# no logging buffered debugging
device(config)# logging buffered 1000
device(config)# write memory
device(config)# exit
device# reload
device(config)# logging facility local0
device(config)# ip show-portname
device# show logging
device(config)# logging persistence
device# clear logging
```

Local Buffer Configuration Notes

- You must save the configuration and reload the software to place any changes into effect.
- The modified number of syslog messages remains persistent across reloads if the **logging persistence** command is configured.
- The number of syslog messages remains persistent once a reload is performed after the **logging buffered** command is configured.
- The syslog messages in the buffer remain persistent across reloads when **logging persistence** command is configured.
- The number of persistent log messages across soft reboots is the same as the number of dynamic syslog messages.
- If you decrease the size of the syslog buffer, the software clears the buffer before placing any changes into effect.
- If you increase the size of the syslog buffer, the software clears some of the older locally buffered syslog messages.

Syslog Reboot Configuration Considerations

- If the syslog buffer size was set to a different value using the **logging buffered** command, the syslog is cleared after a soft reboot, even when logging persistence is in effect. This occurs only with a soft reboot immediately following a syslog buffer size change. A soft reboot by itself does not clear the syslog. To prevent the system from clearing the syslog, leave the number of entries allowed in the syslog buffer unchanged.
- Logging persistence does not save syslog messages after a hard reboot. When the RUCKUS ICX device is power-cycled, the syslog messages are cleared.
- If logging persistence is enabled and you load a new software image on the device, you must first clear the log if you want to reload the device. (Refer to step 9 of Configuring Syslog Service.)
- To configure the device to save the syslog messages after a soft reboot, use the **logging persistence** command.

Disabling or Re-enabling Syslog

Syslog is enabled by default. To disable it, enter the **no logging on** command from global configuration mode.

```
device(config)# no logging on
```

To re-enable syslog logging, enter the **logging on** command.

```
device(config)# logging on
```

NOTE

An expected delay of over 30 seconds is observed for this command functionality to take effect.

Displaying the Syslog Configuration

To display the syslog parameters currently in effect on a RUCKUS ICX device, enter the following command from any CLI mode.

```
device# show logging
Syslog logging: enabled (0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 3 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
               I=informational N=notification W=warning
Static Log Buffer:
Dec 15 19:04:14:A:Fan 1, fan on right connector, failed
Dynamic Log Buffer (50 entries):
Dec 15 18:46:17:I:Interface ethernet 1/1/4, state up
Dec 15 18:45:21:I:Bridge topology change, vlan 4095, interface 4, changed
state to forwarding
Dec 15 18:45:15:I:Warm start
```

Generating the Syslog Specific to RFC 5424

By default, syslog is generated in accordance with RFC 3164. To provide the maximum amount of information in every syslog in a structured format, you can enable syslog logging specific to RFC 5424.

The syslog that conforms to RFC 5424 has an enhanced syslog header that helps to identify the type of syslog, filter the syslog message, identify the syslog generation time with the year and milliseconds with respect to the time zone, and other enhancements. The syslog specific to RFC 5424 can be enabled using the **logging enable rfc5424** command. The logging buffer must be cleared before enabling syslog specific to RFC 5424; otherwise, the system displays an error.

NOTE

If the **logging cli-command** command is present in the running configuration, switching between syslog functionality that follows the default RFC 3164 standard and syslog specific to RFC 5424 standard is not supported.

The following table provides a comparison of the syslog header information available in the RFC 3164 and RFC 5424 syslog logging.

TABLE 10 Syslog Headers Available for RFC 3164 and RFC 5424

Syslog RFC 3164	Syslog RFC 5424
PRIORITY	PRIORITY
	VERSION
TIMESTAMP	TIMESTAMP
HOSTNAME	HOSTNAME
	APP-NAME
	PROCID
	MSGID
	STRUCTURED-DATA
MSG	MSG

RFC 5424 provides the following syslog headers:

- **PRIORITY:** Represents both facility and severity of the messages as described in RFC 3164.
- **VERSION:** Denotes the version of the syslog protocol specification.
- **TIMESTAMP:** A formalized timestamp that denotes the date and time when the event is logged and includes the syslog generation time with the year and milliseconds with respect to the time zone.

The following example shows the date and time format in RFC 5424.

2020-08-13T22:14:15.003Z represents August 13, 2020 at 10:14 PM and 15 seconds, 3 milliseconds into the next second. The timestamp is in UTC. The timestamp provides millisecond resolution.

NOTE

The suffix "Z", when applied to a time, denotes a Coordinated Universal Time (UTC) offset of 00:00.

- **HOSTNAME:** Identifies the machine that originally sent the syslog message. The contents of the HOSTNAME field may have one of the following values and the field uses the following order of preference:
 - FQDN
 - Hostname
 - NILVALUE: A field used when the syslog application is incapable of obtaining its host name.
- **APP-NAME:** Identifies the device or application from which the message is originated. The APP-NAME is intended for filtering messages on a relay or collector. The NILVALUE is used when the syslog application is incapable of obtaining its APP-NAME.
- **PROCID:** Often used to provide the process name or process ID associated with a syslog system. The NILVALUE is used when a process ID is not available.
- **MSGID:** Identifies the type of message. The NILVALUE is used when the syslog application does not, or cannot, provide any value.
- **STRUCTURED-DATA:** Provides a mechanism to express information in a well-defined and interpretable data format as per RFC 5424. STRUCTURED-DATA can contain zero, one, or multiple structured-data elements. In case of zero structured-data elements, the STRUCTURED-DATA field uses NILVALUE.
- **MSG:** Contains a free-form message that provides information about the event.

Displaying Syslog Messages Specific to RFC 5424

If syslog logging specific to RFC 5424 is enabled, the **show logging** command displays the syslog messages generated in the format specific to RFC 5424.

```
device# show logging
Syslog logging: enabled (RFC: 5424, 0 messages dropped, 0 flushes, 0 overruns)
  Buffer logging: level ACDMEINW, 22 messages logged
  level code: A=alert C=critical D=debugging M=emergency E=error
              I=informational N=notification W=warning

Static Log Buffer:
2020-10-07T08:18:40.728Z:I: - ICX7450_Router - System [meta sequenceId=5] BOMSystem: Stack unit 1 POE
Power supply 1 with 748000 mwatts capacity is up

Dynamic Log Buffer (50 lines):
2012-12-19T01:36:40.798Z:I: ruckus - - - [meta sequenceId=23] BOMSystem: Interface ethernet 3/1/23, state up
2012-12-19T01:36:40.797Z:I: ruckus - - - [meta sequenceId=22] BOMSystem: Interface ethernet 3/1/13, state up
2012-12-19T01:36:40.796Z:I: ruckus - - - [meta sequenceId=21] BOMSystem: Interface ethernet 3/1/1, state up
2012-12-19T01:36:24.591Z:A: ruckus - - - [meta sequenceId=20] BOMStack unit 3 Power supply 2 is down
2012-12-19T01:36:24.591Z:I: ruckus - - - [meta sequenceId=18] BOMSystem: Stack unit 3 Power supply 1 with
748000 mwatts capacity is up
2012-12-19T01:36:23.406Z:I: ruckus - - - [meta sequenceId=16] BOMSystem: Interface ethernet 3/3/1, state up
2012-12-19T01:36:22.526Z:I: ruckus - - - [meta sequenceId=15] BOMStack: Stack unit 1 has been elected as
ACTIVE unit of the stack system
2012-12-19T01:36:21.297Z:I: ruckus - - - [meta sequenceId=14] BOMSystem: Interface ethernet 1/4/1, state up
2012-12-19T01:36:20.858Z:I: ruckus - - - [meta sequenceId=13] BOMStack: Stack unit 1 has been elected as
ACTIVE unit of the stack system
2012-12-19T01:36:20.822Z:I: ruckus - - - [meta sequenceId=12] BOMStack: Stack unit 3 has been added to the
stack system
2012-12-19T01:36:20.500Z:I: ruckus - - - [meta sequenceId=11] BOMSystem: Interface ethernet 1/4/1, state
down
2012-12-19T01:36:19.695Z:I: ruckus - - - [meta sequenceId=10] BOMSystem: Interface ethernet 1/4/1, state up
2012-12-19T01:36:18.509Z:I: ruckus - - - [meta sequenceId=9] BOMSystem: Stack unit 1 Power supply 1 is u
2012-12-19T01:36:17.865Z:I: ruckus - - - [meta sequenceId=7] BOMSystem: Interface ethernet 1/3/1, state up
2012-12-19T01:36:16.466Z:I: ruckus - - - [meta sequenceId=6] BOMSystem: Interface ethernet mgmt1, state up
2012-12-19T01:36:16.447Z:I: ruckus - - - [meta sequenceId=5] BOMSystem: Warm start
2012-12-19T01:36:16.260Z:D: ruckus - - - [meta sequenceId=4] BOMDHCPC: starting dhcp client service on 57
port(s)
2012-12-19T01:36:16.259Z:D: ruckus - - - [meta sequenceId=3] BOMDHCPC: Found static IP address 10.20.15.15
subnet mask 255.255.255.0 on port mgmt1
2012-12-19T01:36:16.259Z:D: ruckus - - - [meta sequenceId=2] BOMDHCPC: Found static IP address 20.20.20.3
subnet mask 255.255.255.0 on port 1/1/3
2012-12-19T01:36:16.259Z:D: ruckus - - - [meta sequenceId=1] BOMDHCPC: Found static IP address 10.10.10.2
subnet mask 255.255.255.0 on port 1/1/1
```


Syslog Messages

- [Syslog Message Descriptions.....](#) 113
- [Syslog messages IPsec and IKEv2.....](#) 143

NOTE

This chapter does not list syslog messages that can be displayed when a debug option is enabled.

The messages are listed by the message type:

- Emergencies (none)
- Alerts
- Critical
- Errors
- Warnings
- Notifications
- Informational
- Debugging

Syslog Message Descriptions

This section lists all of the syslog messages. Note that some of the messages apply only to Layer 3 switches.

Message SYSLOG: <12> Apr 27 07:42:51 device Warning! Buffer reached the threshold limit of 50%

Explanation Indicates that the log entries reached the configured threshold value of logging buffer.

Message Level Warning

Message Flexlink: Active link 1/1/3 failed, Backup link :2/1/5 transitioning to UP state by root user from console session

Explanation Indicates active link failover and backup link transitions to the UP state.

Message Level Informational

Message Flexlink: Backup link 2/1/5 failed, Active link :1/1/3 transitioning to UP state by root user from console session

Explanation Indicates backup link failover and active link transitions to the UP state.

Message Level Informational

Message Flexlink: Preempting Backup link : 2/1/5 by Active link :1/1/3 due to 'forced' mode

Explanation Indicates that the current forwarding link is preempted according to the preemption forced mode configuration.

Message Level Informational

Message Flexlink: Preempting Active link : 1/1/3(100Mbit) by Backup link :2/1/5(1Gbit) due to 'bandwidth' mode

Explanation Indicates that the current forwarding link is preempted according to the preemption bandwidth mode configuration.

Message Level Informational

Message Failed to add route <ipv4 route> VRF:<num> (Max-Routes: <value>) in unit <id> due to route-table full.

Syslog Messages

Syslog Message Descriptions

Explanation	Indicates the IPv4 routing table is full.
	NOTE The usage of the Layer 3 route table can be checked using the show hardware route status command.
Message Level	Notification
Message	Failed to add route <ipv6 route> VRF:<num> (Max-Routes:<value>) in unit <id> due to route-table full.
Explanation	Indicates the IPv6 routing table is full.
	NOTE The usage of the Layer 3 route table can be checked using the show hardware ipv6-route status command.
Message Level	Notification
Message	System: Critical Temperature : <c-temp> and MAC Temperature : <mac-temp>, on Unit on unit 11.
Explanation	Indicates the critical temperature and MAC temperature on each unit. The temperature values are output as integers.
Message Level	Informational
Message	MVRP: VLAN <vlan-id> dynamically added.
Explanation	Indicates that a VLAN is dynamically added.
Message Level	Informational
Message	MVRP: VLAN <vlan-id> dynamically removed.
Explanation	Indicates that a VLAN is dynamically removed.
Message Level	Informational
Message	MVRP: Port(s) <interfaces> dynamically added to VLAN <vlan-id>.
Explanation	Indicates that a port or range of ports are added to a VLAN.
Message Level	Informational
Message	MVRP: Port(s) <interfaces> dynamically removed from VLAN <vlan-id>.
Explanation	Indicates that a port or range of ports are removed from a VLAN.
Message Level	Informational
Message	MVRP: VLAN <vlan-id> changed to static.
Explanation	Indicates that a VLAN is changed to static VLAN.
Message Level	Informational
Message	MVRP: Port(s) <interfaces> changed to static member of VLAN <vlan-id>.
Explanation	Indicates that a port or range of ports changed to static member of a VLAN.
Message Level	Informational
Message	MVRP: Auto removed port(s) <interfaces> from VLAN <vlan-id>.
Explanation	Indicates that a port or range of ports are removed from a VLAN.
Message Level	Informational
Message	MVRP: Auto added port(s) <interfaces> to VLAN <vlan-id>.
Explanation	Indicates that a port or range of ports are added to a VLAN.
Message Level	Informational
Message	MVRP: System maximum vlan reached, cannot add vlan <vlan-id>.
Explanation	Indicates that the System has reached the maximum VLAN limit and more VLANs cannot be added.
Message Level	Error

Message	Security: Port Security violation protect activated on interface <if-name>
Explanation	Indicates that a specific PMS port entered protection mode.
Message Level	Informational
Message	Security: Port Security violation protect de-activated on interface <if-name>
Explanation	Indicates that a specific PMS port is no longer in protection mode.
Message Level	Informational
Message	num-modules modules and 1 power supply, need more power supply!!
Explanation	Indicates that the chassis needs more power supplies to run the modules in the chassis.
Message Level	The num-modules parameter indicates the number of modules in the chassis. Alert
Message	<i>Stack: system upgrade completed</i>
Explanation	Indicates that the system upgrade is completed successfully.
Message Level	Informational
Message	<i>Stack: system upgrade failed</i>
Explanation	Indicates that the system upgrade failed.
Message Level	Alert
Message	<i>Stack: stack unit <unit_id> completed upgrade</i>
Explanation	Indicates that the stack unit with a particular stack id completed system upgrade.
Message Level	Informational
Message	<i>Stack: system upgrade failed, stack unit <unit_id> is in <failure_state></i>
Explanation	Indicates that the system upgrade for stack unit with a particular stack id failed and is in the failure state as specified in the message.
Message Level	Alert
Message	<i>Stack: system upgrade started and most of user interfaces are blocked</i>
Explanation	Indicates that the system upgrade started and most of user interfaces are blocked.
Message Level	Alert
Message	Fan num , location , failed
Explanation	A fan has failed.
Message Level	The num is the fan number. The location describes where the failed fan is in the chassis. Alert
Message	MAC Authentication failed for mac-address on portnum
Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum ; however, the VLAN returned in the RADIUS Access-Accept message did not refer to a valid VLAN or VLAN ID on the RUCKUS device. This is treated as an authentication failure.
Message Level	Alert
Message	MAC Authentication failed for mac-address on portnum (Invalid User)
Explanation	RADIUS authentication failed for the specified mac-address on the specified portnum because the MAC address sent to the RADIUS server was not found in the RADIUS server users database.
Message Level	Alert
Message	MAC Authentication failed for mac-address on portnum (No VLAN Info received from RADIUS server)

Syslog Messages

Syslog Message Descriptions

Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum ; however, dynamic VLAN assignment was enabled for the port, but the RADIUS Access-Accept message did not include VLAN information. This is treated as an authentication failure.
Message Level	Alert
Message	MAC Authentication failed for mac-address on portnum (Port is already in another radius given vlan)
Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum ; however, the RADIUS Access-Accept message specified a VLAN ID, although the port had previously been moved to a different RADIUS-assigned VLAN. This is treated as an authentication failure.
Message Level	Alert
Message	MAC Authentication failed for mac-address on portnum (RADIUS given vlan does not exist)
Explanation	RADIUS authentication was successful for the specified mac-address on the specified portnum ; however, the RADIUS Access-Accept message specified a VLAN that does not exist in the RUCKUS configuration. This is treated as an authentication failure.
Message Level	Alert
Message	MAC Authentication failed for mac-address on portnum (RADIUS given VLAN does not match with TAGGED vlan)
Explanation	Multi-device port authentication failed for the mac-address on a tagged port because the packet with this MAC address as the source was tagged with a VLAN ID different from the RADIUS-supplied VLAN ID.
Message Level	Alert
Message	Management module at slot slot-num state changed from module-state to module-state.
Explanation	Indicates a state change in a management module. The slot-num indicates the chassis slot containing the module. The module-state can be one of the following: <ul style="list-style-type: none">● active● standby● crashed● coming-up● unknown
Message Level	Alert
Message	OSPF LSA Overflow, LSA Type = lsa-type
Explanation	Indicates an LSA database overflow. The lsa-type parameter indicates the type of LSA that experienced the overflow condition. The LSA type is one of the following: <ul style="list-style-type: none">● 1 - Router● 2 - Network● 3 - Summary● 4 - Summary● 5 - External
Message Level	Alert
Message	OSPF Memory Overflow
Explanation	OSPF has run out of memory.

Message Level	Alert
Message	System: Module in slot slot-num encountered PCI config read error: Bus PCI-bus-number , Dev PCI-device-number , Reg Offset PCI-config-register-offse t .
Explanation	The module encountered a hardware configuration read error.
Message Level	Alert
Message	System: Module in slot slot-num encountered PCI config write error: Bus PCI-bus-number , Dev PCI-device-number , Reg Offset PCI-config-register-offset .
Explanation	The module encountered a hardware configuration write error.
Message Level	Alert
Message	System: Module in slot slot-num encountered PCI memory read error: Mem Addr memory-address
Explanation	The module encountered a hardware memory read error.
	The memory-address is in hexadecimal format.
Message Level	Alert
Message	System: Module in slot slot-num encountered PCI memory write error: Mem Addr memory-address .
Explanation	The module encountered a hardware memory write error.
	The memory-address is in hexadecimal format.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI bridge validation failure. Module will be deleted.
Explanation	The module encountered an unrecoverable (hardware) bridge validation failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI config read failure. Module will be deleted.
Explanation	The module encountered an unrecoverable hardware configuration read failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI config write failure. Module will be deleted.
Explanation	The module encountered an unrecoverable hardware configuration write failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI device validation failure. Module will be deleted.
Explanation	The module encountered an unrecoverable (hardware) device validation failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: Module in slot slot-num encountered unrecoverable PCI memory read failure. Module will be deleted.
Explanation	The module encountered an unrecoverable hardware memory read failure. The module will be disabled or powered down.
Message Level	Alert

Syslog Messages

Syslog Message Descriptions

Message	System: Module in slot slot-num encountered unrecoverable PCI memory write failure. Module will be deleted.
Explanation	The module encountered an unrecoverable hardware memory write failure. The module will be disabled or powered down.
Message Level	Alert
Message	System: No Free Tcam Entry available. System will be unstable
Explanation	You must reboot the device.
Message Level	Alert
Message	System: Temperature is over shutdown level, system is going to be reset in num seconds
Explanation	The chassis temperature has risen above shutdown level. The system will be shut down in the amount of time indicated.
Message Level	Alert
Message	Temperature degrees C degrees, warning level warn-degrees C degrees, shutdown level shutdown-degrees C degrees
Explanation	Indicates an over temperature condition on the active module. The degrees value indicates the temperature of the module. The warn-degrees value is the warning threshold temperature configured for the module. The shutdown-degrees value is the shutdown temperature configured for the module.
Message Level	Alert
Message	Authentication shut down portnum due to DOS attack
Explanation	Denial of Service (DoS) attack protection was enabled for multi-device port authentication on the specified portnum , and the per-second rate of RADIUS authentication attempts for the port exceeded the configured limit. The RUCKUS device considers this to be a DoS attack and disables the port.
Message Level	Critical
Message	BGP4: Not enough memory available to run BGP4
Explanation	The device could not start the BGP4 routing protocol because there is not enough memory available.
Message Level	Debug
Message	DOT1X: Not enough memory
Explanation	There is not enough system memory for 802.1X authentication to take place. Contact RUCKUS Technical Support.
Message Level	Debug
Message	No of prefixes received from BGP peer ip-addr exceeds maximum prefix-limit...shutdown
Explanation	The Layer 3 switch has received more than the specified maximum number of prefixes from the neighbor, and the Layer 3 switch is therefore shutting down its BGP4 session with the neighbor.
Message Level	Error
Message	IPv6: IPv6 protocol disabled on the device from session-id
Explanation	IPv6 protocol was disabled on the device during the specified session.
Message Level	Informational
Message	IPv6: IPv6 protocol enabled on the device from session-id
Explanation	IPv6 protocol was enabled on the device during the specified session.
Message Level	Informational
Message	MAC ACL applied to port port-id by username from session-id (acl id= ids)
Explanation	Indicates a MAC ACL was applied to the specified port by the specified user during the specified session. session-id can be console, telnet, ssh, or snmp.

Message Level	acl-ids is a list of the MAC ACLs that were applied. Informational
Message Explanation	MAC ACL removed from port port-id by username from session-id (ACL id=acl-ids) Indicates a MAC ACL was removed from the specified port by the specified user during the specified session. session-id can be console, telnet, ssh, or snmp.
Message Level	acl-ids is a list of the MAC ACLs that were removed. Informational
Message Explanation	Security: Password has been changed for user username from session-id Password of the specified user has been changed during the specified session ID or type. session-id can be console, telnet, ssh, or snmp.
Message Level	Informational
Message Explanation	device-name : Logical link on interface ethernet slot#/port# is down. The specified ports were logically brought down while singleton was configured on the port.
Message Level	Informational
Message Explanation	device-name : Logical link on interface ethernet slot#/port# is up. The specified ports were logically brought up while singleton was configured on the port.
Message Level	Informational
Message Explanation	user-name login to PRIVILEGED mode A user has logged into the Privileged exec mode of the CLI. The user-name is the user name.
Message Level	Informational
Message Explanation	user-name login to USER EXEC mode A user has logged into the User exec mode of the CLI. The user-name is the user name.
Message Level	Informational
Message Explanation	user-name logout from PRIVILEGED mode A user has logged out of Privileged exec mode of the CLI. The user-name is the user name.
Message Level	Informational
Message Explanation	user-name logout from USER EXEC mode A user has logged out of the User exec mode of the CLI. The user-name is the user name.
Message Level	Informational
Message Explanation	ACL ACL id added deleted modified from console telnet ssh snmp session A user created, modified, deleted, or applied an ACL through an SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message Explanation	Bridge is new root, vlan vlan-id , root ID root-id A Spanning Tree Protocol (STP) topology change has occurred, resulting in the RUCKUS device becoming the root bridge. The vlan-id is the ID of the VLAN in which the STP topology change occurred. The root-id is the STP bridge root ID.
Message Level	Informational

Syslog Messages

Syslog Message Descriptions

Message	Bridge root changed, vlan vlan-id , new root ID string , root interface portnum
Explanation	A Spanning Tree Protocol (STP) topology change has occurred. The vlan-id is the ID of the VLAN in which the STP topology change occurred. The root-id is the STP bridge root ID. The portnum is the number of the port connected to the new root bridge.
Message Level	Informational
Message	Bridge topology change, vlan vlan-id , interface portnum , changed state to stp-state
Explanation	A Spanning Tree Protocol (STP) topology change has occurred on a port. The vlan-id is the ID of the VLAN in which the STP topology change occurred. The portnum is the port number. The stp-state is the new STP state and can be one of the following: <ul style="list-style-type: none">• disabled• blocking• listening• learning• forwarding• unknown
Message Level	Informational
Message	Cold start
Explanation	The device has been powered on.
Message Level	Informational
Message	DHCP: snooping on untrusted port portnum , type number, drop
Explanation	The device has indicated that the DHCP client receives DHCP server reply packets on untrusted ports, and packets are dropped.
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address cannot remove inbound ACL
Explanation	An error occurred while removing the inbound ACL.
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address Downloading an IP ACL, but IP ACL have no effect on a switch port
Explanation	The RADIUS server returned an IP ACL, but the portnum is a switch port (no IP address).
Message Level	Informational
Message	DOT1X: port portnum - MAC mac address is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC ACLs
Explanation	802.1X authentication failed for the Client with the specified mac address on the specified portnum either due to insufficient system resources on the device, or due to invalid IP ACL or MAC ACL information returned by the RADIUS server.
Message Level	Informational
Message	DOT1X: Port portnum is unauthorized because system resource is not enough or the invalid information to set the dynamic assigned IP ACLs or MAC ACLs

Explanation	802.1X authentication could not take place on the port. This happened because strict security mode was enabled and one of the following occurred: <ul style="list-style-type: none"> • Insufficient system resources were available on the device to apply an IP ACL or MAC ACL to the port • Invalid information was received from the RADIUS server (for example, the Filter-ID attribute did not refer to an existing IP ACL or MAC ACL)
Message Level	Informational
Message	<code>DOT1X: Port portnum currently used vlan-id changes to vlan-id due to dot1x-RADIUS vlan assignment</code>
Explanation	A user has completed 802.1X authentication. The profile received from the RADIUS server specifies a VLAN ID for the user. The port to which the user is connected has been moved to the VLAN indicated by vlan-id .
Message Level	Informational
Message	<code>DOT1X: Port portnum currently used vlan-id is set back to port default vlan-id vlan-id</code>
Explanation	The user connected to portnum has disconnected, causing the port to be moved back into its default VLAN, vlan-id .
Message Level	Informational
Message	<code>DOT1X: Port portnum , AuthControlledPortStatus change: authorized</code>
Explanation	The status of the interface controlled port has changed from unauthorized to authorized.
Message Level	Informational
Message	<code>DOT1X: Port portnum , AuthControlledPortStatus change: unauthorized</code>
Explanation	The status of the interface controlled port has changed from authorized to unauthorized.
Message Level	Informational
Message	<code>Enable super port-config read-only password deleted added modified from console telnet ssh snmp OR Line password deleted added modified from console telnet ssh snmp</code>
Explanation	A user created, re-configured, or deleted an Enable or Line password through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	<code>ERR_DISABLE: Interface ethernet portnum err-disable recovery timeout</code>
Explanation	Errdisable recovery timer expired and the port has been reenabled.
Message Level	Informational
Message	<code>ERR_DISABLE: Interface ethernet 16, err-disable recovery timeout</code>
Explanation	If the wait time (port is down and is waiting to come up) expires and the port is brought up the following message is displayed.
Message Level	Informational
Message	<code>ERR_DISABLE: Link flaps on port ethernet 16 exceeded threshold; port in err-disable state</code>
Explanation	The threshold for the number of times that a port link toggles from "up" to "down" and "down" to "up" has been exceeded.
Message Level	Informational
Message	<code>Interface portnum , line protocol down</code>
Explanation	The line protocol on a port has gone down.
Message Level	The portnum is the port number. Informational
Message	<code>Interface portnum , line protocol up</code>
Explanation	The line protocol on a port has come up.

Syslog Messages

Syslog Message Descriptions

Message Level	The portnum is the port number. Informational
Message Explanation	System: Interface portnum , state down A port has gone down.
Message Level	The portnum is the port number. Informational
Message Explanation	Interface portnum , state up A port has come up.
Message Level	The portnum is the port number. Informational
Message Explanation	MAC Based Vlan Disabled on port port id A MAC Based VLAN has been disabled on a port
Message Level	Informational
Message Explanation	MAC Based Vlan Enabled on port port id A MAC Based VLAN has been enabled on a port.
Message Level	Informational
Message Explanation	MAC ACL added deleted modified from console telnet ssh snmp session filter id = MAC ACL ID , src MAC = Source MAC address any, dst MAC = Destination MAC address any A user created, modified, deleted, or applied this MAC ACL through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message Explanation	MSTP: BPDU-guard interface ethernet port-number detect (Received BPDU), putting into err-disable state. BPDU guard violation occurred in MSTP.
Message Level	Informational
Message Explanation	OPTICAL MONITORING: port port-number is not capable. The optical transceiver is qualified by RUCKUS, but the transceiver does not support digital optical performance monitoring.
Message Level	Informational
Message Explanation	Port p priority changed to n A port priority has changed.
Message Level	Informational
Message Explanation	Port portnum , srcip-security max-ipaddr-per-int reached.Last IP= ipaddr The address limit specified by the srcip-security max-ipaddr-per-interface command has been reached for the port.
Message Level	Informational
Message Explanation	Port portnum , srcip-security max-ipaddr-per-int reached.Last IP= ipaddr The address limit specified by the srcip-security max-ipaddr-per-interface command has been reached for the port.
Message Level	Informational
Message Explanation	Security: console login by username to USER PRIVILEGE EXEC mode The specified user logged into the device console into the specified exec mode.
Message Level	Informational
Message Explanation	Security: console logout by username The specified user logged out of the device console.
Message Level	Informational

Message	Security: telnet SSH login by username from src IP ip-address , src MAC mac-address to USER PRIVILEGE EXEC mode
Explanation	The specified user logged into the device using Telnet or SSH from either or both the specified IP address and MAC address. The user logged into the specified exec mode.
Message Level	Informational
Message	Security: telnet SSH logout by username from src IP ip-address, src MAC mac-address to USER PRIVILEGE EXEC mode
Explanation	The specified user logged out of the device. The user was using Telnet or SSH to access the device from either or both the specified IP address and MAC address. The user logged out of the specified exec mode.
Message Level	Informational
Message	SNMP read-only community read-write community contact location user group view engineid trap [host] [value -str] deleted added modified from console telnet ssh snmp session
Explanation	A user made SNMP configuration changes through the SNMP, console, SSH, or Telnet session. [value-str] does not appear in the message if SNMP community or engineid is specified.
Message Level	Informational
Message	SNMP Auth. failure, intruder IP: ip-addr
Explanation	A user has tried to open a management session with the device using an invalid SNMP community string. The ip-addr is the IP address of the host that sent the invalid community string.
Message Level	Informational
Message	SSH telnet server enabled disabled from console telnet ssh snmp session [by user username]
Explanation	A user enabled or disabled an SSH or Telnet session, or changed the SSH enable/disable configuration through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	startup-config was changed or startup-config was changed by user-name
Explanation	A configuration change was saved to the startup-config file. The user-name is the user ID, if they entered a user ID to log in.
Message Level	Informational
Message	STP: Root Guard Port port-number, VLAN vlan-ID consistent (Timeout).
Explanation	Root guard unblocks a port.
Message Level	Informational
Message	STP: Root Guard Port port-number , VLAN vlan-ID inconsistent (Received superior BPDU) .
Explanation	Root guard blocked a port.
Message Level	Informational
Message	STP: VLAN vlan id BPDU-Guard on Port port id triggered (Received BPDU), putting into err-disable state
Explanation	The BPDU guard feature has detected an incoming BPDU on {vlan-id, port-id}
Message Level	Informational
Message	STP: VLAN vlan id Root-Protect Port port id , Consistent (Timeout)
Explanation	The root protect feature goes back to the consistent state.
Message Level	Informational

Syslog Messages

Syslog Message Descriptions

Message	STP: VLAN vlan id Root-Protect Port port id , Inconsistent (Received superior BPDU)
Explanation	The root protect feature has detected a superior BPDU and goes into the inconsistent state on { vlan-id , port-id }.
Message Level	Informational
Message	STP: VLAN vlan-id BPDU-guard port port-number detect (Received BPDU), putting into err-disable state
Explanation	STP placed a port into an errdisable state for BPDU guard.
Message Level	Informational
Message	STP: VLAN 1 BPDU-guard port port-number detect (Received BPDU), putting into err-disable state.
Explanation	BPDU guard violation in occurred in STP or RSTP.
Message Level	Informational
Message	Syslog server IP-address deleted added modified from console telnet ssh snmp OR Syslog operation enabled disabled from console telnet ssh snmp
Explanation	A user made Syslog configuration changes to the specified Syslog server address, or enabled or disabled a Syslog operation through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	SYSTEM: Optic is not Ruckus-qualified (port-number)
Explanation	RUCKUS does not support the optical transceiver.
Message Level	Informational
Message	System: Fan fan id (from left when facing right side), ok
Explanation	The fan status has changed from fail to normal.
Message Level	Informational
Message	System: Fan speed changed automatically to fan speed
Explanation	The system automatically changed the fan speed to the speed specified in this message.
Message Level	Informational
Message	System: No free TCAM entry. System will be unstable
Explanation	There are no TCAM entries available.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is added from the unit / slot / port to unit / slot / port on VLANs vlan-id to vlan-id
Explanation	A MAC address is added to a range of interfaces, which are members of the specified VLAN range.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is added to the unit / slot / port to unit / slot / port on vlan-id
Explanation	A MAC address is added to a range of interfaces, which are members of the specified VLAN.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is added to portnumber unit / slot / port on VLAN vlan-id
Explanation	A MAC address is added to an interface and the interface is a member of the specified VLAN.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is deleted from the unit/slot/ port to unit / slot / port on vlan-id
Explanation	A MAC address is deleted from a range of interfaces, which are members of the specified VLAN.
Message Level	Informational

Message	System: Static MAC entry with MAC Address mac-address is deleted from et he unit / slot / port to unit / slot / port on VLANs vlan-id to vlan-id
Explanation	A MAC address is deleted from a range of interfaces, which are members of the specified VLAN range.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is deleted from portnumber unit / slot / port on vlan-id
Explanation	A MAC address is deleted from an interface and the interface is a member of the specified VLAN.
Message Level	Informational
Message	System: Static MAC entry with MAC Address mac-address is deleted from portnumber unit / slot / port on VLANs vlan-id to vlan-id
Explanation	A MAC address is deleted from an interface and the interface is a member of the specified VLAN range.
Message Level	Informational
Message	telnet SSH access [by username] from src IP source ip address , src MAC source MAC address rejected, n attempts
Explanation	There were failed SSH, or Telnet login access attempts from the specified source IP and MAC address. <ul style="list-style-type: none"> • [by user username] does not appear if telnet or SSH clients are specified. • n is the number of times this SNMP trap occurred in the last five minutes, or other configured number of minutes.
Message Level	Informational
Message	Trunk group (ports) created by 802.3ad link-aggregation module.
Explanation	802.3ad link aggregation is configured on the device, and the feature has dynamically created a trunk group (aggregate link). The ports variable is a list of the ports that were aggregated to make the trunk group.
Message Level	Informational
Message	user username added deleted modified from console telnet ssh snmp
Explanation	A user created, modified, or deleted a local user account through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	vlan vlan id added deleted modified from console telnet ssh snmp session
Explanation	A user created, modified, or deleted a VLAN through the SNMP, console, SSH, or Telnet session.
Message Level	Informational
Message	Warm start
Explanation	The system software (flash code) has been reloaded.
Message Level	Informational
Message	Stack: Stack unit unit# has been deleted to the stack system
Explanation	The specified unit has been deleted from the stacking system.
Message Level	Informational
Message	Stack unit unitNumber has been elected as ACTIVE unit of the stack system
Explanation	The specified unit in a stack has been elected as the Master unit for the stacking system.
Message Level	Informational
Message	Stack: Stack unit unit# has been added to the stack system
Explanation	The specified unit has been added to the stacking system.
Message Level	Informational
Message	System: Management MAC address changed to mac_address
Explanation	The management MAC address of a stacking system has been changed
Message Level	Informational

Syslog Messages

Syslog Message Descriptions

Message	System: Stack unit unit# Fan fan# (description), failed
Explanation	The operational status of a fan in the specified unit in a stack changed from normal to failure.
Message Level	Informational
Message	System: Stack unit unit# Power supply power-supply# is down
Explanation	The operational status of a power supply of the specified unit in a stack changed from normal to failure.
Message Level	Informational
Message	System: Stack unit unit# Power supply power-supply# is up
Explanation	The operational status of a power supply of the specified unit in a stack changed from failure to normal.
Message Level	Informational
Message	System: Stack unit unit# Fan fan# (description), ok
Explanation	The operational status of a fan in the specified unit in a stack changed from failure to normal.
Message Level	Informational
Message	System: Stack unit unitNumber Temperature actual-temp C degrees, warning level warning-temp C degrees, shutdown level shutdown-temp C degrees
Explanation	The actual temperature reading for a unit in a stack is above the warning temperature threshold.
Message Level	Informational
Message	vlan vlan-id Bridge is RootBridge mac-address (MgmtPriChg)
Explanation	802.1W changed the current bridge to be the root bridge of the given topology due to administrative change in bridge priority.
Message Level	Informational
Message	vlan vlan-id Bridge is RootBridge mac-address (MsgAgeExpiry)
Explanation	The message age expired on the Root port so 802.1W changed the current bridge to be the root bridge of the topology.
Message Level	Informational
Message	vlan vlan-id interface portnum Bridge TC Event (DOT1wTransition)
Explanation	802.1W recognized a topology change event in the bridge. The topology change event is the forwarding action that started on a non-edge Designated port or Root port.
Message Level	Informational
Message	vlan vlan-id interface portnum STP state - state (DOT1wTransition)
Explanation	802.1W changed the state of a port to a new state: forwarding, learning, blocking. If the port changes to blocking, the bridge port is in discarding state.
Message Level	Informational
Message	vlan vlan-id New RootBridge mac-address RootPort portnum (BpduRcvd)
Explanation	802.1W selected a new root bridge as a result of the BPDUs received on a bridge port.
Message Level	Informational
Message	vlan vlan-id New RootPort portnum (RootSelection)
Explanation	802.1W changed the port role to Root port, using the root selection computation.
Message Level	Informational
Message	ACL exceed max DMA L4 cam resource, using flow based ACL instead
Explanation	The port does not have enough Layer 4 CAM entries for the ACL. To correct this condition, allocate more Layer 4 CAM entries. To allocate more Layer 4 CAM entries, enter the following command at the CLI configuration level for the interface:
Message Level	ip access-group max-l4-cam num Notification

Message	ACL insufficient L4 cam resource, using flow based ACL instead
Explanation	The port does not have a large enough CAM partition for the ACLs
Message Level	Notification
Message	ACL insufficient L4 session resource, using flow based ACL instead
Explanation	The device does not have enough Layer 4 session entries. To correct this condition, allocate more memory for sessions. To allocate more memory, enter the following command from the global configuration mode of the CLI interface: system-max session-limit num
Message Level	Notification
Message	ACL port fragment packet inspect rate rate exceeded on port portnum
Explanation	The fragment rate allowed on an individual interface has been exceeded. The <i>rate</i> indicates the maximum rate allowed. The <i>portnum</i> indicates the port. This message can occur if fragment throttling is enabled.
Message Level	Notification
Message	ACL system fragment packet inspect rate rate exceeded
Explanation	The fragment rate allowed on the device has been exceeded. The <i>rate</i> indicates the maximum rate allowed. This message can occur if fragment throttling is enabled.
Message Level	Notification
Message	Authentication Disabled on portnum
Explanation	The multi-device port authentication feature was disabled on the on the specified portnum .
Message Level	Notification
Message	Authentication Enabled on portnum
Explanation	The multi-device port authentication feature was enabled on the on the specified portnum .
Message Level	Notification
Message	BGP Peer ip-addr DOWN (IDLE)
Explanation	Indicates that a BGP4 neighbor has gone down. The <i>ip-addr</i> is the IP address of the neighbor BGP4 interface with the RUCKUS device.
Message Level	Notification
Message	BGP Peer ip-addr UP (ESTABLISHED)
Explanation	Indicates that a BGP4 neighbor has come up. The <i>ip-addr</i> is the IP address of the neighbor BGP4 interface with the RUCKUS device.
Message Level	Notification
Message	DHCP: snooping on untrusted port portnum , type number, drop
Explanation	Indicates that the DHCP client receives DHCP server reply packets on untrusted ports, and packets are dropped.
Message Level	Critical
Message	DOT1X issues software but not physical port down indication of Port portnum to other software applications
Explanation	The device has indicated that the specified is no longer authorized, but the actual port may still be active.
Message Level	Notification

Syslog Messages

Syslog Message Descriptions

Message	<code>DOT1X issues software but not physical port up indication of Port portnum to other software applications</code>
Explanation	The device has indicated that the specified port has been authenticated, but the actual port may not be active.
Message Level	Notification
Message	<code>DOT1X: Port port_id Mac mac_address -user user_id - RADIUS timeout for authentication</code>
Explanation	The RADIUS session has timed out for this 802.1x port.
Message Level	Notification
Message	<code>ISIS L1 ADJACENCY DOWN system-id on circuit circuit-id</code>
Explanation	The Layer 3 switch adjacency with this Level-1 IS-IS has gone down. The system- <i>i d</i> is the system ID of the IS-IS. The circuit- <i>id</i> is the ID of the circuit over which the adjacency was established.
Message Level	Notification
Message	<code>ISIS L1 ADJACENCY UP system-id on circuit circuit-id</code>
Explanation	The Layer 3 switch adjacency with this Level-1 IS-IS has come up. The system- <i>id</i> is the system ID of the IS-IS. The circuit- <i>id</i> is the ID of the circuit over which the adjacency was established.
Message Level	Notification
Message	<code>ISIS L2 ADJACENCY DOWN system-id on circuit circuit-id</code>
Explanation	The Layer 3 switch adjacency with this Level-2 IS-IS has gone down. The system- <i>id</i> is the system ID of the IS-IS. The circuit- <i>id</i> is the ID of the circuit over which the adjacency was established.
Message Level	Notification
Message	<code>ISIS L2 ADJACENCY UP system-id on circuit circuit-id</code>
Explanation	The Layer 3 switch adjacency with this Level-2 IS-IS has come up. The system- <i>id</i> is the system ID of the IS-IS. The circuit- <i>id</i> is the ID of the circuit over which the adjacency was established.
Message Level	Notification
Message	<code>Local ICMP exceeds burst-max burst packets, stopping for lockup seconds!!</code>
Explanation	The number of ICMP packets exceeds the burst-max threshold set by the ip icmp burst command. The RUCKUS device may be the victim of a Denial of Service (DoS) attack. All ICMP packets will be dropped for the number of seconds specified by the lockup value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Message Level	Notification
Message	<code>Local TCP exceeds burst-max burst packets, stopping for lockup seconds!!</code>
Explanation	The number of TCP SYN packets exceeds the burst-max threshold set by the ip tcp burst command. The RUCKUS device may be the victim of a TCP SYN DoS attack. All TCP SYN packets will be dropped for the number of seconds specified by the lockup value. When the lockup period expires, the packet counter is reset and measurement is restarted.
Message Level	Notification
Message	<code>Local TCP exceeds num burst packets, stopping for num seconds!!</code>

Explanation Threshold parameters for local TCP traffic on the device have been configured, and the maximum burst size for TCP packets has been exceeded.

The first num is the maximum burst size (maximum number of packets allowed).

The second num is the number of seconds during which additional TCP packets will be blocked on the device.

NOTE

This message can occur in response to an attempted TCP SYN attack.

Message Level Notification

Message MAC Authentication RADIUS timeout for mac_address on port port_id

Explanation The RADIUS session has timed out for the MAC address for this port.

Message Level Notification

Message MAC Authentication succeeded for mac-address on portnum

Explanation RADIUS authentication was successful for the specified mac-address on the specified portnum .

Message Level Notification

Message Module was inserted to slot slot-num

Explanation Indicates that a module was inserted into a chassis slot.

The slot-num is the number of the chassis slot into which the module was inserted.

Message Level Notification

Message Module was removed from slot slot-num

Explanation Indicates that a module was removed from a chassis slot.

The slot-num is the number of the chassis slot from which the module was removed.

Message Level Notification

Message OSPF interface state changed,rid router-id , intf addr ip-addr , state ospf-state

Explanation Indicates that the state of an OSPF interface has changed.

The router-id is the router ID of the RUCKUS device.

The ip-addr is the interface IP address.

The ospf-state indicates the state to which the interface has changed and can be one of the following:

- down
- loopback
- waiting
- point-to-point
- designated router
- backup designated router
- other designated router
- unknown

Message Level Notification

Message OSPF intf authen failure, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type

Explanation Indicates that an OSPF interface authentication failure has occurred.

The *router-id* is the router ID of the RUCKUS device.

Syslog Messages

Syslog Message Descriptions

The ip-addr is the IP address of the interface on the RUCKUS device.

The src-ip-addr is the IP address of the interface from which the RUCKUS device received the authentication failure.

The error-type can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message

`OSPF intf config error, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type`

Explanation

Indicates that an OSPF interface configuration error has occurred.

The router-id is the router ID of the RUCKUS device.

The ip-addr is the IP address of the interface on the RUCKUS device.

The src-ip-addr is the IP address of the interface from which the RUCKUS device received the error packet.

The error-type can be one of the following:

- bad version
- area mismatch
- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch

- option mismatch
- unknown

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message

OSPF intf rcvd bad pkt, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , pkt type pkt-type

Explanation

Indicates that an OSPF interface received a bad packet.

The router-id is the router ID of the RUCKUS device.

The ip-addr is the IP address of the interface on the RUCKUS device.

The src-ip-addr is the IP address of the interface from which the RUCKUS device received the authentication failure.

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message

OSPF intf rcvd bad pkt: Bad Checksum, rid ip-addr , intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type

Explanation

The device received an OSPF packet that had an invalid checksum.

The rid ip-addr is the RUCKUS router ID.

The intf addr ip-addr is the IP address of the RUCKUS interface that received the packet.

The pkt size num is the number of bytes in the packet.

The checksum num is the checksum value for the packet.

The pkt src addr ip-addr is the IP address of the neighbor that sent the packet.

The pkt type type is the OSPF packet type and can be one of the following:

- hello
- database description
- link state request
- link state update
- link state acknowledgement

Syslog Messages

Syslog Message Descriptions

	<ul style="list-style-type: none">● unknown (indicates an invalid packet type)
Message Level	Notification
Message	OSPF intf rcvd bad pkt: Bad Packet type, rid ip-addr, intf addr ip-addr , pkt size num , checksum num , pkt src addr ip-addr , pkt type type
Explanation	The device received an OSPF packet with an invalid type. The parameters are the same as for the Bad Checksum message. The pkt type type value is "unknown", indicating that the packet type is invalid.
Message Level	Notification
Message	OSPF intf rcvd bad pkt: Invalid packet size, rid ip-addr, intf addr ip-addr, pkt size num , checksum num , pkt src addr ip-addr , pkt type type
Explanation	The device received an OSPF packet with an invalid packet size. The parameters are the same as for the Bad Checksum message.
Message Level	Notification
Message	OSPF intf rcvd bad pkt: Unable to find associated neighbor, rid ip-addr, intf addr ip-addr, pkt size num , checksum num , pkt src addr ip-addr , pkt type type
Explanation	The neighbor IP address in the packet is not in the list of OSPF neighbors in the RUCKUS device. The parameters are the same as for the Bad Checksum message.
Message Level	Notification
Message	OSPF intf retransmit, rid router-id, intf addr ip-addr, nbr rid nbr- router-id , pkt type is pkt-type, LSA type lsa-type , LSA id lsa-id, LSA rid lsa-router-id
Explanation	An OSPF interface on the RUCKUS device has retransmitted a Link State Advertisement (LSA). The router-id is the router ID of the RUCKUS device. The ip-addr is the IP address of the interface on the RUCKUS device. The nbr-router-id is the router ID of the neighbor router. The packet-type can be one of the following: <ul style="list-style-type: none">● hello● database description● link state request● link state update● link state ack● unknown The lsa-type is the type of LSA. The lsa-id is the LSA ID. The lsa-router-id is the LSA router ID.
Message Level	Notification
Message	OSPF LSDB approaching overflow, rid router-id , limit num
Explanation	The software is close to an LSDB condition. The router-id is the router ID of the RUCKUS device. The num is the number of LSAs.

Message Level	Notification
Message	OSPF LSDB overflow, rid router-id, limit num
Explanation	<p>A Link State Database Overflow (LSDB) condition has occurred.</p> <p>The router-id is the router ID of the RUCKUS device.</p> <p>The num is the number of LSAs.</p>
Message Level	Notification
Message	OSPF max age LSA, rid router-id , area area-id , LSA type lsa-type , LSA id lsa-id , LSA rid lsa-router-id
Explanation	<p>An LSA has reached its maximum age.</p> <p>The router-id is the router ID of the RUCKUS device.</p> <p>The area-id is the OSPF area.</p> <p>The lsa-type is the type of LSA.</p> <p>The lsa-id is the LSA ID.</p> <p>The lsa-router-id is the LSA router ID.</p>
Message Level	Notification
Message	OSPF nbr state changed, rid router-id , nbr addr ip-addr , nbr rid nbr-router-Id , state ospf-state
Explanation	<p>Indicates that the state of an OSPF neighbor has changed.</p> <p>The router-id is the router ID of the RUCKUS device.</p> <p>The ip-addr is the IP address of the neighbor.</p> <p>The nbr-router-id is the router ID of the neighbor.</p> <p>The ospf-state indicates the state to which the interface has changed and can be one of the following:</p> <ul style="list-style-type: none">• down• attempt• initializing• 2-way• exchange start• exchange• loading• full• unknown
Message Level	Notification
Message	OSPF originate LSA, rid router-id , area area-id , LSA type lsa-type , LSA id lsa-id , LSA router id lsa-router-id
Explanation	<p>An OSPF interface has originated an LSA.</p> <p>The router-id is the router ID of the RUCKUS device.</p> <p>The area-id is the OSPF area.</p> <p>The lsa-type is the type of LSA.</p>

Syslog Messages

Syslog Message Descriptions

	The lsa-id is the LSA ID.
	The lsa-router-id is the LSA router ID.
Message Level	Notification
Message	OSPF virtual intf authen failure, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type
Explanation	Indicates that an OSPF virtual routing interface authentication failure has occurred. The router-id is the router ID of the RUCKUS device. The ip-addr is the IP address of the interface on the RUCKUS device. The src-ip-addr is the IP address of the interface from which the RUCKUS device received the authentication failure. The error-type can be one of the following: <ul style="list-style-type: none">• bad version• area mismatch• unknown NBMA neighbor• unknown virtual neighbor• authentication type mismatch• authentication failure• network mask mismatch• hello interval mismatch• dead interval mismatch• option mismatch• unknown The packet-type can be one of the following: <ul style="list-style-type: none">• hello• database description• link state request• link state update• link state ack• unknown
Message Level	Notification
Message	OSPF virtual intf config error, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , error type error-type , pkt type pkt-type
Explanation	Indicates that an OSPF virtual routing interface configuration error has occurred. The router-id is the router ID of the RUCKUS device. The ip-addr is the IP address of the interface on the RUCKUS device. The src-ip-addr is the IP address of the interface from which the RUCKUS device received the error packet. The error-type can be one of the following: <ul style="list-style-type: none">• bad version• area mismatch

- unknown NBMA neighbor
- unknown virtual neighbor
- authentication type mismatch
- authentication failure
- network mask mismatch
- hello interval mismatch
- dead interval mismatch
- option mismatch
- unknown

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message

OSPF virtual intf rcvd bad pkt, rid router-id , intf addr ip-addr , pkt src addr src-ip-addr , pkt type pkt-type

Explanation

Indicates that an OSPF interface received a bad packet.

The router-id is the router ID of the RUCKUS device.

The ip-addr is the IP address of the interface on the RUCKUS device.

The src-ip-addr is the IP address of the interface from which the RUCKUS device received the authentication failure.

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

Message Level

Notification

Message

OSPF virtual intf retransmit, rid router-id , intf addr ip-addr , nbr rid nbr-router-id , pkt type is pkt-type , LSA type lsa-type , LSA id lsa-id , LSA rid lsa-router-id

Explanation

An OSPF interface on the RUCKUS device has retransmitted a Link State Advertisement (LSA).

The router-id is the router ID of the RUCKUS device.

The ip-addr is the IP address of the interface on the RUCKUS device.

The nbr-router-id is the router ID of the neighbor router.

Syslog Messages

Syslog Message Descriptions

The packet-type can be one of the following:

- hello
- database description
- link state request
- link state update
- link state ack
- unknown

The lsa-type is the type of LSA.

The lsa-id is the LSA ID.

The lsa-router-id is the LSA router ID.

Message Level

Notification

Message

OSPF virtual intf state changed, rid router-id , area area-id , nbr ip-addr , state ospf-state

Explanation

Indicates that the state of an OSPF virtual routing interface has changed.

The router-id is the router ID of the router the interface is on.

The area-id is the area the interface is in.

The ip-addr is the IP address of the OSPF neighbor.

The ospf-state indicates the state to which the interface has changed and can be one of the following:

- down
- loopback
- waiting
- point-to-point
- designated router
- backup designated router
- other designated router
- unknown

Message Level

Notification

Message

OSPF virtual nbr state changed, rid router-id , nbr addr ip-addr , nbr rid nbr-router-id , state ospf-state

Explanation

Indicates that the state of an OSPF virtual neighbor has changed.

The router-id is the router ID of the RUCKUS device.

The ip-addr is the IP address of the neighbor.

The nbr-router-id is the router ID of the neighbor.

The ospf-state indicates the state to which the interface has changed and can be one of the following:

- down
- attempt
- initializing
- 2-way

- exchange start
- exchange
- loading
- full
- unknown

Message Level

Notification

Message

Transit ICMP in interface portnum exceeds num burst packets, stopping for num seconds!!

Explanation

Threshold parameters for ICMP transit (through) traffic have been configured on an interface, and the maximum burst size for ICMP packets on the interface has been exceeded.

The portnum is the port number.

The first num is the maximum burst size (maximum number of packets allowed).

The second num is the number of seconds during which additional ICMP packets will be blocked on the interface.

NOTE

This message can occur in response to an attempted Smurf attack.

Message Level

Notification

Message

Transit TCP in interface portnum exceeds num burst packets, stopping for num seconds!

Explanation

Threshold parameters for TCP transit (through) traffic have been configured on an interface, and the maximum burst size for TCP packets on the interface has been exceeded.

The portnum is the port number.

The first num is the maximum burst size (maximum number of packets allowed).

The second num is the number of seconds during which additional TCP packets will be blocked on the interface.

NOTE

This message can occur in response to an attempted TCP SYN attack.

Message Level

Notification

Message

VRRP intf state changed, intf portnum , vrid virtual-router-id , state vrrp-state
VRRP (IPv6) intf state changed, intf portnum , vrid virtual-router-id , state vrrp-state

Explanation

A state change has occurred in a Virtual Router Redundancy Protocol (VRRP) or VRRP-E IPv4 or IPv6 interface.

The portnum is the port or interface where VRRP or VRRP-E is configured.

The virtual-router-id is the virtual router ID (VRID) configured on the interface.

The vrrp-state can be one of the following:

- init
- master
- backup
- unknown

Message Level

Notification

Syslog Messages

Syslog Message Descriptions

Message	<code>DOT1X security violation at port portnum , malicious MAC address detected: mac-address</code>
Explanation	A security violation was encountered at the specified port number.
Message Level	Warning
Message	<code>Dup IP ip-addr detected, sent from MAC mac-addr interface portnum</code>
Explanation	Indicates that the RUCKUS device received a packet from another device on the network with an IP address that is also configured on the RUCKUS device. The ip-addr is the duplicate IP address. The mac-addr is the MAC address of the device with the duplicate IP address. alert The portnum is the RUCKUS port that received the packet with the duplicate IP address. The address is the packet source IP address.
Message Level	Warning
Message	<code>IGMP/MLD no hardware vidx, broadcast to the entire vlan. rated limited number</code>
Explanation	IGMP or MLD snooping has run out of hardware application VLANs. There are 4096 application VLANs per device. Traffic streams for snooping entries without an application VLAN are switched to the entire VLAN and to the CPU to be dropped. This message is rate-limited to appear a maximum of once every 10 minutes. The rate-limited number shows the number on non-printed warnings.
Message Level	Warning
Message	<code>IGMP/MLD: vlanId(portId) is V1 but rcvd V2 from nbr ipAddr</code>
Explanation	Port has received a query with a MLD version that does not match the port MLD version. This message is rated-limited to appear a maximum of once every 10 hours.
Message Level	Warning
Message	<code>Latched low RX Power TX Power TX Bias Current Supply Voltage Temperature warning alarm warning, port port-number</code>
Explanation	The optical transceiver on the given port has risen above or fallen below the alarm or warning threshold.
Message Level	Warning
Message	<code>list ACL-num denied ip-proto src-ip-addr (src-tcp / udp-port) (Ethernet portnum mac-addr) - dst-ip-addr (dst-tcp / udp-port), 1 event(s)</code>
Explanation	Indicates that an Access Control List (ACL) denied (dropped) packets. The ACL-num indicates the ACL number. Numbers 1 - 99 indicate standard ACLs. Numbers 100 - 199 indicate extended ACLs. The ip-proto indicates the IP protocol of the denied packets. The src-ip-addr is the source IP address of the denied packets. The src-tcp / udp-port is the source TCP or UDP port, if applicable, of the denied packets. The portnum indicates the port number on which the packet was denied. The mac-addr indicates the source MAC address of the denied packets. The dst-ip-addr indicates the destination IP address of the denied packets. The dst-tcp / udp-port indicates the destination TCP or UDP port number, if applicable, of the denied packets.
Message Level	Warning
Message	<code>MAC ACL denied packets on port portnum, src macaddr mac-addr , num packets</code>
Explanation	Indicates that a MAC address filtergroup configured on a port has denied packets.

The portnum is the port on which the packets were denied.

The mac-addr is the source MAC address of the denied packets.

The num indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.

Message Level

Warning

Message

`multicast no software resource: resource-name , rate-limited number`

Explanation

IGMP or MLD snooping has run out of software resources. This message is rate-limited to appear a maximum of once every 10 minutes. The rate-limited number shows the number of non-printed warnings.

Message Level

Warning

Message

`No global IP! cannot send IGMP msg.`

Explanation

The device is configured for **ip multicast active** but there is no configured IP address and the device cannot send out IGMP queries.

Message Level

Warning

Message

`No of prefixes received from BGP peer ip-addr exceeds warning limit num`

Explanation

The Layer 3 switch has received more than the allowed percentage of prefixes from the neighbor.

The ip-addr is the IP address of the neighbor.

The num is the number of prefixes that matches the percentage you specified. For example, if you specified a threshold of 100 prefixes and 75 percent as the warning threshold, this message is generated if the Layer 3 switch receives a 76th prefix from the neighbor.

Message Level

Warning

Message

`rip filter list list-num direction V1 | V2 denied ip-addr , num packets`

Explanation

Indicates that a RIP route filter denied (dropped) packets.

The list-num is the ID of the filter list.

The direction indicates whether the filter was applied to incoming packets or outgoing packets. The value can be one of the following:

- in
- out

The V1 or V2 value specifies the RIP version (RIPv1 or RIPv2).

The ip-addr indicates the network number in the denied updates.

The num indicates how many packets matching the values above were dropped during the five-minute interval represented by the log entry.

Message Level

Warning

Message

`Temperature is over warning level.`

Explanation

The chassis temperature has risen above the warning level.

Message Level

Warning

Message

`ZTP: zero-touch-enable detects total <num of chains> chains (<num of units> units). unstable=<num of units>`

Explanation

Zero-touch-enable detects units.

The *num of chains* is the total number of detected chains

The *num of units* is the total number of detected units

Syslog Messages

Syslog Message Descriptions

Message Level	The <i>num of units</i> is the total number of unstable units Informational
Message	ZTP: Detect an invalid chain, aborts. <port_id> links to an invalid chain.
Explanation	Zero-touch-enable detects an invalid chain.
Message Level	<i>Port_id</i> is the id of a port links to the chain. Informational
Message	ZTP: Send reload to chain <chain_id>
Explanation	Zero-touch-enable sends reload to detected chain.
Message Level	<i>Chain_id</i> is the chain to be reloaded. Informational
Message	ZTP: Recv ZTP request, not qualify reason= <reason>
Explanation	Unit receive ztp request, but is not qualify to join.
Message Level	<i>Reason</i> is the reason for not being qualified. Informational
Message	INTERACTIVE SETUP: Detect an invalid chain, aborts. <port_id> links to an invalid chain.
Explanation	INTERACTIVE SETUP detects an invalid chain.
Message Level	<i>Port_id</i> is the ID of the port links to the chain. Informational
Message	INTERACTIVE SETUP: Send reload to chain <chain_id>.
Explanation	INTERACTIVE SETUP sends reload to detected chain.
Message Level	<i>Chain_id</i> is the ID of the chain to be reloaded. Informational
Message	<i>License: Self-Authenticated Upgrade license %s is applied to unit %d\n</i>
Explanation	The command update-license is successfully entered.
Message Level	Informational
Message	PoE: Power disabled on port 2/1/5 because of overdrive mode change
Explanation	PD is turned off due to overdrive configuration change.
Message Level	Informational
Message	PoE Severe Error: Hardware Fault with ports <number> to <number>. Remove PDs and then configure "no inline power" on these ports.
Explanation	PoE functionality on some ports will not be available when the device fails during operation.
Message Level	Critical
Message	PoE Severe Error: Internal Device supplying power to port <number> is hot. "Distribute the load so that each of the 8 ports group (ports 1-8, 9-16 etc) have equal power consumption."
Explanation	If high power consuming PDs are connected in consecutive ports and the ambient temperature is high, the device gets heated up.
Message Level	Critical
Message	PoE Severe Error: Power being injected on port <number>. No new PDs can get powered on this unit. Configure "no inline power" on all Switch to Switch connected ports of this unit and peer unit(s) to resolve the issue.
Explanation	Voltage applied from ext src is detected from POE port.

Message Level	Error
Message	PoE Severe Error: PD on port <number> cannot be powered due to power being injected on another port of this unit. Configure "no inline power" on all Switch to Switch connected ports of this unit and peer unit(s) to resolve the issue.
Explanation	Misconfiguration or the unit/PSU require RMA .
Message Level	Error
Message	PoE Info: EEPROM Read on slot <number> failed.
Explanation	Software failed to read the vendor ID (EEPROM).
Message Level	Informational
Message	PoE Info: Image <string> Not Supported on slot <number>. Minimum required image version <string>.
Explanation	Unsupported image version detected and suggests the minimum required image version.
Message Level	Informational
Message	PoE Warning: Upgrading firmware in slot <number>....DO NOT HOTSWAP OR POWER DOWN THE MODULE.
Explanation	Indicates that the firmware upgrade process is under way and module should not be powered down or hotswapped.
Message Level	Warning
Message	U%d-MSG: PoE Info: Firmware Download on slot <number>.....<number> percent completed.
Explanation	Indicates the status of firmware download.
Message Level	Informational
Message	PoE: Port <port-number> lost non-PD, so enabling PD detection.
Explanation	Indicates that PD detection is enabled after losing non-PD on a specific port.
Message Level	Informational
Message	PoE: Power enabled on port <port-number>.
Explanation	Indicates that power is enabled on a port.
Message Level	Informational
Message	PoE: Power disabled on port <port-number> because of detection of non-PD. PD detection will be disabled on port.
Explanation	Indicates that power is disabled upon detection of non-PD and PD detection will be disabled.
Message Level	Informational
Message	PoE: Power disabled on port <port-number> because of admin initiated firmware upgrade.
Explanation	Indicates that power is disabled upon admin-initiated firmware upgrade.
Message Level	Alert
Message	PoE: Power disabled on port <port-number> because of admin off.
Explanation	Indicates that power is disabled after admin off.
Message Level	Informational
Message	PoE: Power disabled on port <port-number> because of power management.
Explanation	Indicates that power is disabled because of power management.
Message Level	Critical
Message	PoE: Power disabled on port <port-number> because of PD disconnection.
Explanation	Indicates that power is disabled upon PD disconnection.
Message Level	Informational

Syslog Messages

Syslog Message Descriptions

Message	PoE: Power disabled on port <port-number> because of PD overload.
Explanation	Indicates that power is disabled because of PD overload.
Message Level	Informational
Message	PoE: Power disabled on port <port-number> because of PD fault.
Explanation	Indicates that power is disabled because of PD fault.
Message Level	Alert
Message	PoE: Power disabled on port <port-number> because of internal fault.
Explanation	Indicates that power is disabled because of internal fault.
Message Level	Alert
Message	PoE: Power disabled on port <port-number> because of PSU fault.
Explanation	Indicates that power is disabled because of PSU fault.
Message Level	Alert
Message	PoE: Power disabled on port <port-number> because of h/w pin assertion.
Explanation	Indicates that power is disabled because of hardware pin assertion.
Message Level	Alert
Message	PoE: Power disabled on port <port-number> because of unknown reason.
Explanation	Indicates that power is disabled because of unknown reason.
Message Level	Alert
Message	PoE: Power adjustment failed: insufficient free power for extra allocation of <decimal> mwatts to port <port-number>.
Explanation	Indicates that power adjustment failed due to insufficient free power.
Message Level	Alert
Message	PoE: Unexpected reset of controller on slot/unit <number> device <number> occurred. Recovery started.
Explanation	Indicates that recovery has started due to unexpected reset of controller.
Message Level	Informational
Message	PoE: Controller on slot/unit <number> device <number> restarted and power recovered on ports.
Explanation	Indicates that controller restarted power recovered on the ports.
Message Level	Informational
Message	PoE Info: PoE module <number> of Unit <number> on ports %d/1/1 to %d/1/%d detected. Initializing....
Explanation	Indicates that PoE module is detected on the ports and initializing started.
Message Level	Informational
Message	PoE Info: PoE module <number> of Unit <number> initialization is done.
Explanation	Indicates that PoE module initialization has completed.
Message Level	Informational
Message	PoE Error: PoE controller error on module%s in slot <number>.
Explanation	Indicates PoE controller error.
Message Level	Error
Message	PoE Error: General internal error when starting PoE module%s in slot <number>.
Explanation	Indicates general internal error when starting PoE module.
Message Level	Error

Message	PoE Error: Device 0 failed to start on PoE module%s in slot <number>.
Explanation	Indicates that device failed to start on PoE module.
Message Level	Error
Message	PoE Error: Device 0 disconnected all ports because of high temp when starting PoE module%s in slot <number>.
Explanation	Indicates that device disconnected all ports due to high temperature when starting PoE module.
Message Level	Error
Message	PoE Alarm: Device 0 has high temp alarm when starting PoE module%s in slot <number>.
Explanation	Indicates high temperature on device when starting PoE module.
Message Level	Alert
Message	PoE Error: Device 0 failed on PoE module%s in slot <number>.
Explanation	Indicates that device has failed on PoE module.
Message Level	Error
Message	PoE Severe: Device 0 disconnected all ports because high temp exceeded disconnection limit on PoE module%s in slot <number>.
Explanation	Indicates that device has disconnected all ports because high temp exceeded disconnection limit on PoE module.
Message Level	Critical
Message	PoE Severe Error: Lost communication link with the PoE controller in slot <number>. Shutting down and restarting the PoE module to recover.
Explanation	Indicates that PoE module restarts to recover after shutdown after communication link break with the PoE controller.
Message Level	Error
Message	PoE Error: Incompatible firmware for PoE module %d. Please install latest firmware.
Explanation	Indicates that the firmware is not compatible after performing PoE Hardware - firmware compatibility check during boot up and firmware upgrade.
Message Level	Error
Message	PoE Alarm: VOP Test Reported Error on device <number>, no af/at detection possible for ports %s, But PDs would get powered.
Explanation	Indicates that VOP Test Reported Error on device but PDs may get powered.
Message Level	Alert

Syslog messages IPsec and IKEv2

Message	IKEv2: Maximum IKE Peers Limit Reached
Explanation	The maximum IKEv2 peer limit is reached on the device.
Message Level	Warning
Message	IKEv2: Recovered from Maximum IKE Peers Limit Condition
Explanation	The device is recovered after reaching the maximum IKEv2 peer limit.
Message Level	Informational
Message	IKEv2: IKEv2 session <up_down> source <source_address> Destination <destination_address> VRF <vrf_id> SPI <spi_id>
Explanation	A state change has occurred for an IKEv2 session.
	The <i>up_down</i> is the state of the interface.
	The <i>source_address</i> is the source IP address of the IKEv2 session.

Syslog Messages

Syslog messages IPsec and IKEv2

The *destination address* is the destination IP address in the packet.

The *vrf_id* is the tunnel base VRF.

The *spi_id* is the security parameter index (SPI) in the packet.

Message Level

Informational

Message

IKEv2: Invalid Message Type Received with Source <source_address> Destination <destination_address> SPI <spi_id> MessageType <x>

Explanation

An invalid IKEv2 message type is received.

The *source_address* is the source IP address in the packet.

The *destination address* is the destination IP address in the packet.

The *spi_id* is the security parameter index (SPI) in the packet.

The *x* is the value of the unsupported message type in the IKEv2 packet.

Message Level

Informational

Message

IKEv2: Invalid Payload Type Received with Source <source_address> Destination <destination_address> SPI <spi_id> PayloadType <x>

Explanation

A state change has occurred for an IPsec session.

The *source_address* is the source IP address in the packet.

The *destination address* is the destination IP address in the packet.

The *spi_id* is the security parameter index (SPI) in the packet.

The *x* is the value of the unsupported payload type.

Message Level

Informational

Message

IPsec: IPsec module <module_id> on unit <unit_id> is <up_down>

Explanation

A state change has occurred for an IPsec module. Hot swapping of the module is not supported, but an IPsec module is marked as down when the IPsec module is not usable.

The *module_id* is the module ID.

The *unit_id* is the unit ID.

The *up_down* is the state of the module.

Message Level

Alert

Message

IPsec: IPsec session <up_down> source <source_address> Destination <destination_address> VRF <vrf_id> SPI <spi_id> Direction <direction>

Explanation

A state change has occurred for an IPsec session.

The *up_down* is the state of the interface.

The *source_address* is the source IP address of the IPsec session.

The *destination address* is the destination IP address in the packet.

The *vrf_id* is the tunnel base VRF.

The *spi_id* is the security parameter index (SPI) in the packet.

The *direction* is ingress or egress.

Message Level

Informational

Syslog messages system

Message	System: Interface ipsec_tnnl <tunnel_id>, state up
Explanation	The IPsec tunnel interface has come up . The <i>tunnel_id</i> is the tunnel ID.
Message Level	Informational
Message	System: Interface ipsec_tnnl <tunnel_id>, state down <reason>
Explanation	The IPsec tunnel interface has gone down. The <i>tunnel_id</i> is the tunnel ID. The <i>reason</i> variable can be one of the following: <ul style="list-style-type: none">• clear IKE SA• clear IPSEC SA• IKE session down• IPSEC session down• tunnel source interface down• tunnel no destination route• Administratively brought down• IPSEC card down• Switchover and Failover
Message Level	Informational

